7450 Ethernet Service Switch
7750 Service Router
7950 Extensible Routing System

Releases Up To 23.7.R2

MD-CLI Advanced Configuration Guide - Part II

3HE 19550 AAAB TQZZA
Edition: 01
September 2023

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Nokia is committed to diversity and inclusion. We are continuously reviewing our customer documentation and consulting with standards bodies to ensure that terminology is inclusive and aligned with the industry. Our future customer documentation will be updated accordingly.

This document includes Nokia proprietary and confidential information, which may not be distributed or disclosed to any third parties without the prior written consent of Nokia.

This document is intended for use by Nokia's customers ("You"/"Your") in connection with a product purchased or licensed from any company within Nokia Group of Companies. Use this document as agreed. You agree to notify Nokia of any errors you may find in this document; however, should you elect to use this document for any purpose(s) for which it is not intended, You understand and warrant that any determinations You may make or actions You may take will be based upon Your independent judgment and analysis of the content of this document.

Nokia reserves the right to make changes to this document without notice. At all times, the controlling version is the one available on Nokia's site.

No part of this document may be modified.

NO WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY OF AVAILABILITY, ACCURACY, RELIABILITY, TITLE, NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE, IS MADE IN RELATION TO THE CONTENT OF THIS DOCUMENT. IN NO EVENT WILL NOKIA BE LIABLE FOR ANY DAMAGES, INCLUDING BUT NOT LIMITED TO SPECIAL, DIRECT, INDIRECT, INCIDENTAL OR CONSEQUENTIAL OR ANY LOSSES, SUCH AS BUT NOT LIMITED TO LOSS OF PROFIT, REVENUE, BUSINESS INTERRUPTION, BUSINESS OPPORTUNITY OR DATA THAT MAY ARISE FROM THE USE OF THIS DOCUMENT OR THE INFORMATION IN IT, EVEN IN THE CASE OF ERRORS IN OR OMISSIONS FROM THIS DOCUMENT OR ITS CONTENT.

Copyright and trademark: Nokia is a registered trademark of Nokia Corporation. Other product names mentioned in this document may be trademarks of their respective owners.

© 2023 Nokia.

Table of contents

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

List of tables

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

List of figures

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Preface

About This Guide

The Advanced Configuration Guide is divided into three volumes, the Part I Guide, the Part II Guide, and the Part III Guide.

- Part I provides advanced configurations for basic systems, system management, interface configuration, router configuration, unicast routing protocols, MPLS, OAM and diagnostics, and VSR Installation and Setup.
- Part II provides advanced configurations for services overview, Layer 2 and EVPN services, Layer 3 services, and Quality of Service.
- Part III provides advanced configurations for Multi-Service Integrated Service Adapter (MS-ISA) – Extended Services Appliance (ESA), and Triple Play Service Delivery Architecture (TPSDA).

The MD-CLI Advanced Configuration Guide is divided into two volumes, the Part I Guide and the Part II Guide.

- Part I provides advanced configurations for basic systems, system management, interface configuration, router configuration, unicast routing protocols, MPLS, OAM and diagnostics, and VSR Installation and Setup.
- Part II provides advanced configurations for services overview, Layer 2 and EVPN services, Layer 3 services, Multi-Service Integrated Service Adapter (MS-ISA) – Extended Services Appliance (ESA), and Triple Play Service Delivery Architecture (TPSDA).

The guide is organized alphabetically within each category and provides feature and configuration explanations, CLI descriptions and overall solutions. The chapters in the Advanced Configuration Guide are written for and based on several Releases, up to 23.7.R2. The Applicability section in each chapter specifies on which release the configuration is based.

The Advanced Configuration Guide supplements the user configuration guides listed in the 7450 ESS, 7750 SR, and 7950 XRS Guide to Documentation.

Audience

This manual is intended for network administrators who are responsible for configuring the routers. It is assumed that the network administrators have a detailed understanding of networking principles and configurations.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Services Overview

This section provides configuration information for the following topics:

- BGP Selective Label-IPv4 Route Installation
- G.8032 Ethernet Ring Protection Multiple Ring Topology
- G.8032 Ethernet Ring Protection Single Ring Topology
- GRE Tunnel Origination and Termination Using Non-system IP Addresses
- Network Group Encryption Helper
- Seamless BFD Application — Auto-bind tunnel

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

BGP Selective Label-IPv4 Route Installation

This chapter provides information about BGP selective label-IPv4 route installation.

Topics in this chapter include:

Applicability

The information and configuration in this chapter are based on SR OS Release 23.3.R1. BGP selective label-IPv4 route installation is supported in SR OS Release 19.10.R2, and later.

Overview

Many service providers use BGP label-unicast (BGP-LU) to build network designs that connect multiple domains into unified and scalable network fabrics. However, the number of BGP-LU IPv4 routes that are distributed in the control plane can exceed the capacity of the Forwarding Information Base (FIB) and Label Forwarding Information Base (LFIB) of small access routers.

One solution is to apply import policies on the access router to limit the number of BGP-LU IPv4 routes accepted in the RIB-IN, but this is labor-intensive and prone to errors. A better solution is selective BGP-LU IPv4 route installation in the base routing instance, which addresses these issues.

When the **selective-label-ipv4-install** command is configured in the **bgp** context of the base router, BGP-LU IPv4 routes in the RIB-IN are made invalid if they are received from a base router BGP peer and not needed by any eligible service. When a BGP-LU IPv4 route is invalid in the RIB-IN, the BGP decision process prefers any valid route over this route, and the invalid BGP-LU IPv4 route is not programmed as a next-hop (primary next-hop, ECMP next-hop, or backup next-hop) of any IP route or tunnel.

The **selective-label-ipv4-install** command can be configured in the **bgp** context of the base router: in the global **bgp** context, the group context, or the neighbor context, as follows:

```
[/]
A:admin@PE-1# tree flat detail | match selective-label-ipv4-install
configure groups group <string> router <string> bgp group <string> selective-label-ipv4-install
 <boolean>
configure groups group <string> router <string> bgp neighbor <string | ipv4-address-with-zone
 | ipv4-address | ipv6-address-linklocal-with-zone | ipv6-address | ipv6-address-with-zone>
 selective-label-ipv4-install <boolean>
configure groups group <string> router <string> bgp selective-label-ipv4-install <boolean>
configure router <string> bgp group <string> selective-label-ipv4-install <boolean>
configure router <string> bgp neighbor <ipv4-address-with-zone | ipv4-address | ipv6-address-
linklocal-with-zone | ipv6-address | ipv6-address-with-zone> selective-label-ipv4-install
 <boolean>
configure router <string> bgp selective-label-ipv4-install <boolean>
```

When a BGP-LU IPv4 route is invalid in the RIB-IN, it is marked with the flag Label-Unicast-No-Svc and the invalid route is handled as follows:

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- No route for the IPv4 prefix is added to the route table from the BGP-LU RIB.
- No BGP tunnel for the /32 IPv4 prefix is added to the tunnel table.
- No RIB-OUT is generated for the invalid BGP-LU route, so this invalid route does not trigger a label-swap (incoming label map - ILM) entry to be programmed.

Note:
Configuring the **selective-label-ipv4-install** command on a BGP session unconditionally invalidates all non-/32 BGP-LU IPv4 routes received on that session, because those non-/32 routes are never used to resolve service endpoints.

Table 1: Selective BGP-LU installation logic by service type shows how BGP-LU IPv4 routes are handled when the selective-label-ipv4-install command is configured.

Table 1: Selective BGP-LU installation logic by service type

Service type	Logic marks BGP label-IPv4 routes as invalid except
L2 services with user-provisioned SDPs	When the user-provisioned SDP has a BGP tunnel as transport and the far end matches a /32 BGP-LU IPv4 route, that route is not marked as invalid, regardless of the operational state of the SDP.
L2 services with auto-created SDPs (BGP-AD, BGP-VPLS, BGP-EVPN)	If an L2 service imports a BGP-AD, BGP-VPLS, or BGP-EVPN route, /32 BGP-LU IPv4 routes matching the BGP next-hop address of this BGP route are not marked as invalid.
EVPN next-hop-self route reflector or model-B ASBR	If the base router BGP instance is configured as a next-hop-self RR or a model-B ASBR, BGP-LU IPv4 routes matching any IPv4 address in the BGP next-hop field of a received EVPN route are not marked as invalid, regardless of whether the transport-tunnel resolution filter allows BGP tunnels.
VPRN with explicitly configured SDP	BGP-LU IPv4 routes matching the SDP far-end address are not marked as invalid, regardless of the operational state of the SDP.
VPRN with auto-bind-tunnel	If the auto-bind VPRN service imports VPN-IPv4 or VPN-IPv6 routes where the BGP next-hop matches a BGP-LU IPv4 route, that route is not marked as invalid, regardless of whether the auto-bind-tunnel resolution filter allows BGP tunnels.
VPN-IP next-hop-self RR or model-B ASBR	If the base router BGP instance is configured as a next-hop-self RR or a model-B ASBR, BGP-LU IPv4 routes matching any IPv4 address in the BGP next-hop field of a received VPN-IP route are not marked as invalid, regardless of whether the transport-tunnel resolution filter allows BGP tunnels.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Configuration

Figure 1: Example topology shows the example topology with two PEs with the services that are configured.

Figure 1: Example topology

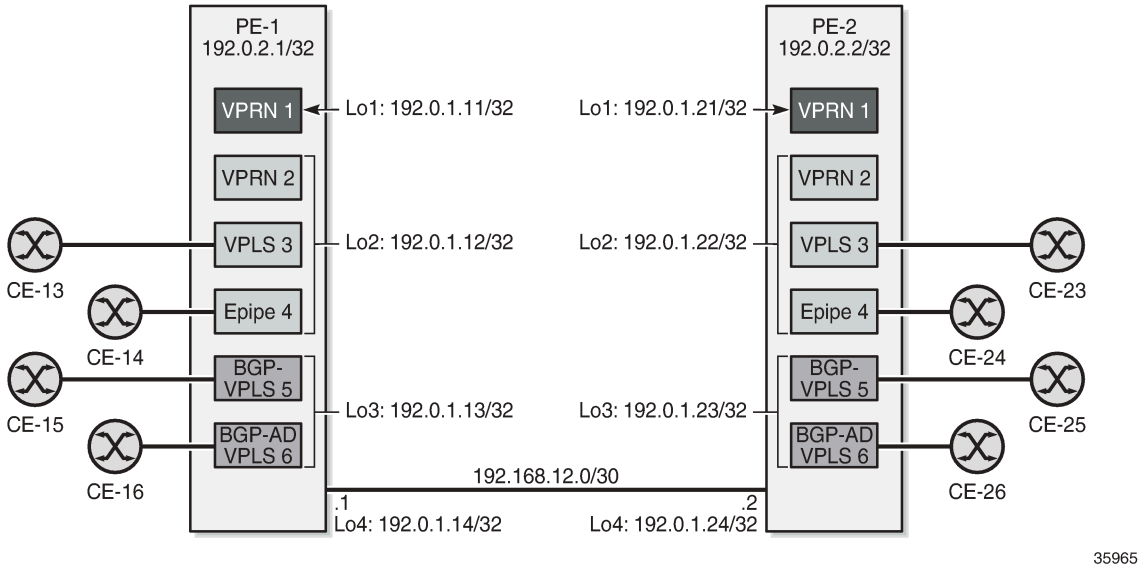

Initial configuration

The initial configuration on the PEs includes:

- Cards, MDAs, ports
- Router interfaces
- SR-ISIS

On PE-2, four loopback interfaces are configured in the base router context with /32 IPv4 addresses: 192.0.1.21/32, 192.0.1.22/32, 192.0.1.23/32, and 192.0.1.24/32. The list of router interfaces on PE-2 is as follows:

```
[/]
A:admin@PE-2# show router interface

===============================================================================
Interface Table (Router: Base)
===============================================================================
Interface-Name                   Adm       Opr(v4/v6)  Mode    Port/SapId
   IP-Address                                                  PfxState
-------------------------------------------------------------------------------
int-PE-2-PE-1                    Up        Up/Down     Network 1/1/c1/2:100
   192.168.12.2/30                                             n/a
lo1                              Up        Up/Down     Network loopback
   192.0.1.21/32                                               n/a
lo2                              Up        Up/Down     Network loopback
   192.0.1.22/32                                               n/a
lo3                              Up        Up/Down     Network loopback
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
   192.0.1.23/32                                                  n/a
lo4                            Up        Up/Down     Network loopback
   192.0.1.24/32                                                  n/a
system                         Up        Up/Down     Network system
   192.0.2.2/32                                                   n/a
-------------------------------------------------------------------------------
Interfaces : 6
===============================================================================
```

These prefixes are exported as BGP-LU routes and the next-hop resolution filter for label-IPv4 routes is configured with SR-ISIS. The configuration on PE-2 is as follows:

```
# on PE-2:
configure {
    policy-options {
        prefix-list "192.0.1.0/24" {
            prefix 192.0.1.0/24 type range {
                start-length 32
                end-length 32
            }
        }
        policy-statement "export-svc-lu-bgp" {
            entry 10 {
                from {
                    prefix-list ["192.0.1.0/24"]
                }
                action {
                    action-type accept
                }
            }
        }
    }
    router "Base" {
        bgp {
            split-horizon true
            ebgp-default-reject-policy {
                import false
                export false
            }
            next-hop-resolution {
                labeled-routes {
                    transport-tunnel {
                        family label-ipv4 {
                            resolution-filter {
                                ldp false
                                sr-isis true
                            }
                        }
                    }
                }
            }
            group "iBGPv4" {
                peer-as 64500
                family {
                    vpn-ipv4 true
                    label-ipv4 true
                }
            }
            neighbor "192.0.2.1" {
                group "iBGPv4"
                export {
                    policy ["export-svc-lu-bgp"]
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
        }
    }
```

PE-1 receives four valid label-IPv4 routes, as follows:

```
[/]
A:admin@PE-1# show router bgp routes label-ipv4
===============================================================================
 BGP Router ID:192.0.2.1        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP LABEL-IPV4 Routes
===============================================================================
Flag  Network                                           LocalPref   MED
      Nexthop (Router)                                  Path-Id     IGP Cost
      As-Path                                                       Label
-------------------------------------------------------------------------------
u*>i  192.0.1.21/32                                     100         None
      192.0.2.2                                         None        10
      No As-Path                                                    524286
u*>i  192.0.1.22/32                                     100         None
      192.0.2.2                                         None        10
      No As-Path                                                    524286
u*>i  192.0.1.23/32                                     100         None
      192.0.2.2                                         None        10
      No As-Path                                                    524286
u*>i  192.0.1.24/32                                     100         None
      192.0.2.2                                         None        10
      No As-Path                                                    524286
-------------------------------------------------------------------------------
Routes : 4
===============================================================================
```

The tunnel table on PE-1 includes four BGP tunnels toward the loopback interfaces on PE-2:

```
[/]
A:admin@PE-1# show router tunnel-table protocol bgp

===============================================================================
IPv4 Tunnel Table (Router: Base)
===============================================================================
Destination           Owner     Encap TunnelId  Pref   Nexthop        Metric
   Color
-------------------------------------------------------------------------------
192.0.1.21/32         bgp       MPLS  262148    12     192.0.2.2      1000
192.0.1.22/32         bgp       MPLS  262147    12     192.0.2.2      1000
192.0.1.23/32         bgp       MPLS  262146    12     192.0.2.2      1000
192.0.1.24/32         bgp       MPLS  262145    12     192.0.2.2      1000
-------------------------------------------------------------------------------
Flags: B = BGP or MPLS backup hop available
       L = Loop-Free Alternate (LFA) hop available
       E = Inactive best-external BGP route
       k = RIB-API or Forwarding Policy backup hop
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The route table on PE-1 shows four BGP-LU IPv4 routes toward the loopback interfaces on PE-2, with next-hop resolved via an SR-ISIS tunnel:

```
[/]
A:admin@PE-1# show router route-table protocol bgp-label

===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
192.0.1.21/32                                 Remote  BGP_LABEL 00h01m12s  170
       192.0.2.2 (tunneled:SR-ISIS:524290)                          10
192.0.1.22/32                                 Remote  BGP_LABEL 00h01m12s  170
       192.0.2.2 (tunneled:SR-ISIS:524290)                          10
192.0.1.23/32                                 Remote  BGP_LABEL 00h01m12s  170
       192.0.2.2 (tunneled:SR-ISIS:524290)                          10
192.0.1.24/32                                 Remote  BGP_LABEL 00h01m12s  170
       192.0.2.2 (tunneled:SR-ISIS:524290)                          10
-------------------------------------------------------------------------------
No. of Routes: 4
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

The tunnel toward destination 192.0.2.2 is the following SR-ISIS tunnel:

```
[/]
A:admin@PE-1# show router tunnel-table 192.0.2.2

===============================================================================
IPv4 Tunnel Table (Router: Base)
===============================================================================
Destination           Owner     Encap TunnelId  Pref   Nexthop        Metric
   Color
-------------------------------------------------------------------------------
192.0.2.2/32          isis (0)  MPLS  524290    11     192.168.12.2   10
-------------------------------------------------------------------------------
Flags: B = BGP or MPLS backup hop available
       L = Loop-Free Alternate (LFA) hop available
       E = Inactive best-external BGP route
       k = RIB-API or Forwarding Policy backup hop
===============================================================================
```

In the following examples, services that use these BGP tunnels are configured .

VPRN 1 with auto-bind-tunnel

VPRN 1 in Figure 2: VPRN 1 uses a BGP transport tunnel with endpoint 192.0.1.21 on PE-2 uses the BGP transport tunnel between loopback interfaces "lo1" with IP address 192.0.1.11/32 on PE-1 and 192.0.1.21/32 on PE-2.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 2: VPRN 1 uses a BGP transport tunnel with endpoint 192.0.1.21 on PE-2

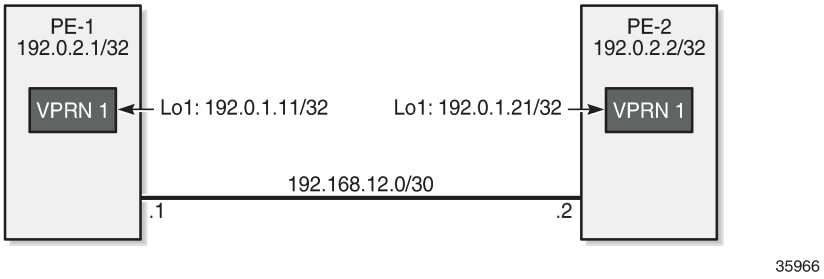

VPRN 1 is configured with an auto-bind-tunnel and the next-hop must be resolved using a BGP tunnel. On PE-2, the policy "export-VPRN1" sets the next-hop to 192.0.1.21 and adds the community "target:64500:1", which matches the vrf-target of VPRN 1.

```
# on PE-2:
configure {
    policy-options {
        community "target:64500:1" {
            member "target:64500:1" { }
        }
        policy-statement "export-VPRN1" {
            entry 10 {
                action {
                    action-type accept
                    next-hop 192.0.1.21
                    community {
                        add ["target:64500:1"]
                    }
                }
            }
        }
    }
    service {
        vprn "VPRN 1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64500:1"
                    vrf-target {
                        community "target:64500:1"
                    }
                    vrf-export {
                        policy ["export-VPRN1"]
                    }
                    auto-bind-tunnel {
                        resolution filter
                    }
                }
            }
            interface "lo1" {
                loopback true
                ipv4 {
                    primary {
                        address 172.31.1.2
                        prefix-length 32
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    }
                }
            }
        }
    }
```

The configuration is similar on PE-1, but the IP addresses are different.

VPRN 1 on PE-1 receives a BGP VPN-IPv4 route for prefix 172.31.1.2/32 from PE-2. The next-hop of this BGP-VPN route is 192.0.1.21:

```
[/]
A:admin@PE-1# show router bgp routes vpn-ipv4
===============================================================================
 BGP Router ID:192.0.2.1        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv4 Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop (Router)                                 Path-Id     IGP Cost
      As-Path                                                      Label
-------------------------------------------------------------------------------
u*>i  64500:1:172.31.1.2/32                            100         None
      192.0.1.21                                       None        0
      No As-Path                                                   524285
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

VPRN 1 on PE-1 uses the BGP tunnel toward 192.0.1.21/32 while the other BGP tunnels are not required on PE-1. When BGP is configured with the **selective-label-ipv4-install** command, only the BGP-LU IPv4 route for 192.0.1.21/32 remains valid. The command can be configured in the global BGP context (as in the following configuration), per **group**, or per **neighbor**:

```
# on PE-1:
configure {
    router "Base" {
        bgp {
            selective-label-ipv4-install true
        }
```

From the four BGP transport tunnels on PE-1, only the BGP tunnel with endpoint 192.0.1.21/32 is used by a service, so it remains valid, as follows:

```
[/]
A:admin@PE-1# show router bgp routes label-ipv4
===============================================================================
 BGP Router ID:192.0.2.1        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
BGP LABEL-IPV4 Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop (Router)                                 Path-Id     IGP Cost
      As-Path                                                      Label
-------------------------------------------------------------------------------
u*>i  192.0.1.21/32                                    100         None
      192.0.2.2                                        None        10
      No As-Path                                                   524286
i     192.0.1.22/32                                    100         None
      192.0.2.2                                        None        10
      No As-Path                                                   524286
i     192.0.1.23/32                                    100         None
      192.0.2.2                                        None        10
      No As-Path                                                   524286
i     192.0.1.24/32                                    100         None
      192.0.2.2                                        None        10
      No As-Path                                                   524286
-------------------------------------------------------------------------------
Routes : 4
===============================================================================
```

The first label-IPv4 route is valid; the other three label-IPv4 routes are marked invalid with flag Label-Unicast-No-Svc:

```
[/]
A:admin@PE-1# show router bgp routes label-ipv4 hunt | match Flags
Flags          : Used Valid Best IGP In-TTM In-RTM
Flags          : Invalid IGP Label-Unicast-No-Svc
Flags          : Invalid IGP Label-Unicast-No-Svc
Flags          : Invalid IGP Label-Unicast-No-Svc
```

In the route table on PE-1, only one BGP-LU IPv4 route remains:

```
[/]
A:admin@PE-1# show router route-table protocol bgp-label

===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
192.0.1.21/32                                 Remote  BGP_LABEL 00h02m05s  170
       192.0.2.2 (tunneled:SR-ISIS:524290)                          10
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

L2 and L3 services with user-provisioned SDP

When SDPs are configured to use a BGP transport tunnel, the corresponding BGP label-IPv4 route is not marked as invalid. The following TLDP-signaled SDP is configured with a BGP transport tunnel between the loopback interfaces "lo2" with IP address 192.0.1.12 on PE-1 and 192.0.1.22 on PE-2:

```
# on PE-2:
configure {
    router "Base" {
        ldp {
            targeted-session {
                peer 192.0.1.12 {
                    local-lsr-id {
                        interface-name "lo2"
                    }
                }
            }
        }
    }
    service {
        sdp 1 {
            admin-state enable
            delivery-type mpls
            bgp-tunnel true
            far-end {
                ip-address 192.0.1.12
            }
        }
    }
```

The configuration is similar on PE-1; only the far-end and peer address is now 192.0.1.22:

```
[/]
A:admin@PE-1# show service sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId  AdmMTU  OprMTU  Far End          Adm  Opr         Del     LSP   Sig
-------------------------------------------------------------------------------
1      0       8970    192.0.1.22       Up   Up          MPLS    B     TLDP
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
Legend: R = RSVP, L = LDP, B = BGP, M = MPLS-TP, n/a = Not Applicable
        I = SR-ISIS, O = SR-OSPF, T = SR-TE, F = FPE
===============================================================================
```

When an SDP uses a BGP transport tunnel, the corresponding BGP label-IPv4 route is not marked as invalid, regardless of the operational state of the SDP. The following command shows that the second BGP label-IPv4 route is now valid:

```
[/]
A:admin@PE-1# show router bgp routes label-ipv4
===============================================================================
 BGP Router ID:192.0.2.1        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
BGP LABEL-IPV4 Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop (Router)                                 Path-Id     IGP Cost
      As-Path                                                      Label
-------------------------------------------------------------------------------
u*>i  192.0.1.21/32                                    100         None
      192.0.2.2                                        None        10
      No As-Path                                                   524286
u*>i  192.0.1.22/32                                    100         None
      192.0.2.2                                        None        10
      No As-Path                                                   524286
i     192.0.1.23/32                                    100         None
      192.0.2.2                                        None        10
      No As-Path                                                   524286
i     192.0.1.24/32                                    100         None
      192.0.2.2                                        None        10
      No As-Path                                                   524286
-------------------------------------------------------------------------------
Routes : 4
===============================================================================
```

This SDP can be used by L2 and L3 services. Figure 3: VPRN 2, VPLS 3, and Epipe 4 use user-provisioned SDP 1 with BGP tunnel shows three services that use SDP 1: VPRN 2, VPLS 3, and Epipe 4.

Figure 3: VPRN 2, VPLS 3, and Epipe 4 use user-provisioned SDP 1 with BGP tunnel

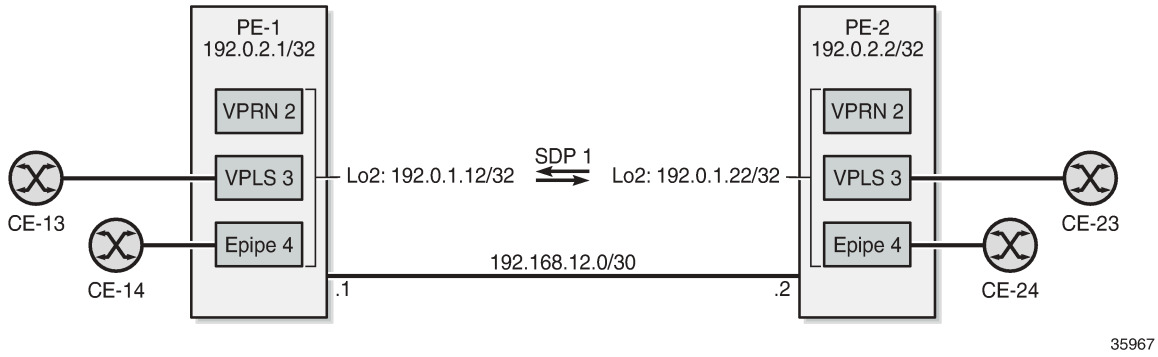

VPRN 2 is similar to VPRN 1, but a spoke-SDP is configured instead of the auto-bind-tunnel. The configuration is as follows:

```
# on PE-1:
configure {
    policy-options {
        community "target:64500:2" {
            member "target:64500:2" { }
        }
        policy-statement "export-VPRN2" {
            entry 10 {
                action {
                    action-type accept
                    next-hop 192.0.1.12
                    community {
                        add ["target:64500:2"]
                    }
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
        }
    }
    service {
        vprn "VPRN 2" {
            admin-state enable
            service-id 2
            customer "1"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64500:2"
                    vrf-target {
                        community "target:64500:2"
                    }
                    vrf-export {
                        policy ["export-VPRN2"]
                    }
                }
            }
            interface "lo1" {
                loopback true
                ipv4 {
                    primary {
                        address 172.31.2.1
                        prefix-length 32
                    }
                }
            }
            spoke-sdp 1:2 {
            }
        }
    }
```

VPLS 3 and Epipe 4 only have a spoke-SDP and a SAP, as follows:

```
# on PE-1:
configure {
    service {
        vpls "VPLS 3" {
            admin-state enable
            service-id 3
            customer "1"
            spoke-sdp 1:3 {
            }
            sap 1/1/c2/1:3 {
            }
        }
        epipe "Epipe 4" {
            admin-state enable
            service-id 4
            customer "1"
            spoke-sdp 1:4 {
            }
            sap 1/1/c2/1:4 {
            }
        }
    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

L2 services with auto-created SDPs

Figure 4: PE-1 receives BGP-VPLS and BGP-AD routes with next-hop 192.0.1.23 shows two VPLS services where the SDPs are auto-created between the loopback interfaces "lo3" on the PEs: BGP-VPLS 5 and BGP-AD VPLS 6.

Figure 4: PE-1 receives BGP-VPLS and BGP-AD routes with next-hop 192.0.1.23

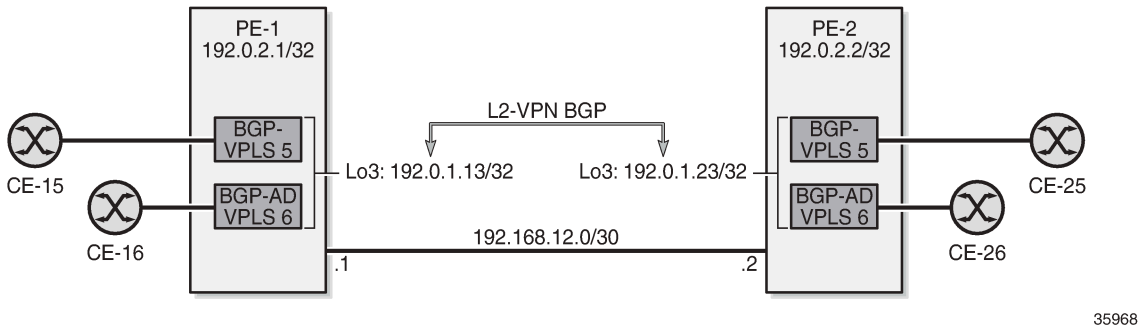

For BGP-VPLS and BGP-AD, a BGP session is established for the L2-VPN address family between the loopback interfaces "lo3" on both PEs:

```
# on PE-2:
configure {
    router "Base" {
        bgp {
            group "iBGP-L2" {
                type internal
                local-address 192.0.1.23
                family {
                    l2-vpn true
                }
            }
            neighbor "192.0.1.13" {
                group "iBGP-L2"
            }
```

For BGP-AD, T-LDP signaling is used, so the following T-LDP session is established:

```
# on PE-2:
configure {
    router "Base" {
        ldp {
            targeted-session {
                peer 192.0.1.13 {
                    local-lsr-id {
                        interface-name "lo3"
                    }
```

The service configuration is as follows:

```
# on PE-2:
configure {
    service {
        vpls "BGP-VPLS 5" {
            admin-state enable
            service-id 5
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            customer "1"
            bgp 1 {
                route-distinguisher "64500:5"
                route-target {
                    export "target:64500:5"
                    import "target:64500:5"
                }
                pw-template-binding "PW1" {
                    import-rt ["target:64500:5"]
                }
            }
            bgp-vpls {
                admin-state enable
                maximum-ve-id 100
                ve {
                    name "PE-2"
                    id 2
                }
            }
            sap 1/1/c2/1:5 {
            }
        }
        vpls "BGP-AD VPLS 6" {
            admin-state enable
            service-id 6
            customer "1"
            bgp 1 {
                route-distinguisher "64500:6"
                route-target {
                    export "target:64500:6"
                    import "target:64500:6"
                }
                pw-template-binding "PW1" {
                }
            }
            bgp-ad {
                admin-state enable
                vpls-id "64500:6"
                vsi-id-prefix 192.0.1.23
            }
            sap 1/1/c2/1:6 {
            }
        }
```

On PE-1, the received L2-VPN BGP routes have next-hop 192.0.1.23:

```
[/]
A:admin@PE-1# show router bgp routes l2-vpn
===============================================================================
 BGP Router ID:192.0.2.1        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP L2VPN Routes
===============================================================================
Flag  RouteType                  Prefix                           MED
      RD                         SiteId                           Label
      Nexthop                    VeId                  BlockSize  LocalPref
      As-Path                    BaseOffset            vplsLabelBa
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                                                          se
-------------------------------------------------------------------------------
u*>i  VPLS                      -                       -          0
      64500:5                   -                       -
      192.0.1.23                2                       8          100
      No As-Path                1                       524276
u*>i  AutoDiscovery             192.0.1.23              -          0
      64500:6                   -                       -
      192.0.1.23                -                       -          100
      No As-Path                -                       -
-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

On PE-1, the following SDPs with far-end address 192.0.1.23 are auto-created in BGP-VPLS 5 and BGP-AD VPLS 6:

```
[/]
A:admin@PE-1# show service id 5 sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
32766:4294967294 BgpVpls  192.0.1.23      Up      Up        524277    524276
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
```

```
[/]
A:admin@PE-1# show service id 6 sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
32767:4294967295 BgpAd    192.0.1.23      Up      Up        524273    524263
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
```

BGP-VPLS 5 and BGP-AD VPLS 6 use a BGP transport tunnel between the "lo3" interfaces, so the corresponding BGP label-IPv4 route is valid, as follows:

```
[/]
A:admin@PE-1# show router bgp routes label-ipv4
===============================================================================
 BGP Router ID:192.0.2.1        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP LABEL-IPV4 Routes
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Flag  Network                                             LocalPref   MED
      Nexthop (Router)                                    Path-Id     IGP Cost
      As-Path                                                         Label
-------------------------------------------------------------------------------
u*>i  192.0.1.21/32                                       100         None
      192.0.2.2                                           None        10
      No As-Path                                                      524286
u*>i  192.0.1.22/32                                       100         None
      192.0.2.2                                           None        10
      No As-Path                                                      524286
u*>i  192.0.1.23/32                                       100         None
      192.0.2.2                                           None        10
      No As-Path                                                      524286
i     192.0.1.24/32                                       100         None
      192.0.2.2                                           None        10
      No As-Path                                                      524286
-------------------------------------------------------------------------------
Routes : 4
===============================================================================
```

Only the BGP tunnel between the "lo4" interfaces is not used by any service, so the last BGP label-IPv4 route is marked invalid in the RIB-IN when **selective-label-ipv4-install** is configured on PE-1, as follows:

```
[/]
A:admin@PE-1# show router bgp routes label-ipv4 hunt | match "Invalid" pre-lines 16

Network        : 192.0.1.24/32
Nexthop        : 192.0.2.2
Path Id        : None
From           : 192.0.2.2
Res. Nexthop   : 192.0.2.2 (ISIS Tunnel)
Local Pref.    : 100                    Interface Name : NotAvailable
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : 10
Connector      : None
Community      : No Community Members
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.2
Fwd Class      : None                   Priority       : None
IPv4 Label     : 524286
Flags          : Invalid IGP Label-Unicast-No-Svc
```

Conclusion

The **selective-label-ipv4-install** command allows BGP-LU IPv4 routes to be marked as invalid in the RIB-IN when these routes are received from a base router BGP peer and not needed by any eligible service. This is a technique to reduce the number of routes in the FIB/LFIB, which is mainly useful for small access routers having small FIB/LFIB sizes.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

G.8032 Ethernet Ring Protection Multiple Ring Topology

This chapter provides information about G.8032 Ethernet ring protection multiple ring topologies.

Topics in this chapter include:

Applicability

Initially, this chapter was written for SR OS Release 12.0.R5, but the MD-CLI in this edition is based on Release 23.3.R2.

Overview

G.8032 Ethernet ring protection is supported for data service SAPs within a regular VPLS service, a PBB VPLS (I/B-component), or a routed VPLS (R-VPLS). G.8032 is one of the fastest protection schemes for Ethernet networks. This chapter describes the advanced topic of multiple ring control, sometimes referred to as multi-chassis protection, with access rings being the most common form of multiple ring topologies. Single rings are covered in the G.8032 Ethernet Ring Protection Single Ring Topology chapter. This chapter will use a VPLS service to illustrate the configuration of G.8032. For very large ring topologies, provider backbone bridging (PBB) can also be used, but that is not configured in this chapter.

ITU-T G.8032v2 specifies protection switching mechanisms and a protocol for Ethernet layer network (ETH) Ethernet rings. Ethernet rings can provide wide-area multipoint connectivity more economically due to their reduced number of links. The mechanisms and protocol defined in ITU-T G.8032v2 are highly reliable with stable protection and never form loops, which would negatively affect network operation and service availability. Each ring node is connected to adjacent nodes participating in the same ring using two independent paths, which use ring links (configured on ports or link aggregation groups (LAGs)). A ring link is bounded by two adjacent nodes and a port for a ring link is called a ring port. The minimum number of nodes on a ring is two.

The fundamentals of this ring protection switching architecture are:

- the principle of loop avoidance and
- the utilization of learning, forwarding, and address table mechanisms defined in the ITU-T G.8032v2 Ethernet flow forwarding function (ETH_FF) (control plane).

Loop avoidance in the ring is achieved by guaranteeing that, at any time, traffic may flow on all but one of the ring links. This particular link is called the ring protection link (RPL) and under normal conditions this link is blocked, so it is not used for traffic. One designated node, the RPL owner, is responsible to block traffic over the one designated RPL. Under a ring failure condition, the RPL owner is responsible for unblocking the RPL, allowing the RPL to be used for traffic. The protocol ensures that even without an RPL owner defined, one link will be blocked and it operates as a *break before make* protocol, specifically the protocol guarantees that no link is restored until a different link in the ring is blocked. The other side of the RPL is configured as an RPL neighbor. An RPL neighbor blocks traffic on the RPL.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The event of a ring link or ring node failure results in protection switching of the traffic. This is achieved under the control of the ETH_FF functions on all ring nodes. A ring automatic protection switching (R-APS) protocol is used to coordinate the protection actions over the ring. The protection switching mechanisms and protocol supports a multi-ring/ladder network that consists of connected Ethernet rings.

Ring protection mechanism

The ring protection protocol is based on the following building blocks:

- ring status change on failure
  - idle → link failure → protection → recovery → idle
- ring control state changes
  - idle → protection → manual switch → forced switch → pending
- re-use existing ETH OAM
  - monitoring: ETH continuity check messages (CCM)
  - failure notification: Y.1731 signal failure
- forwarding database MAC flush on ring status change
- ring protection link (RPL)
  - defines blocked link in idle status

When subrings are used, they can either connect to a major ring (which is configured in the exact same way as a single ring) or another subring, or to a VPLS service. When connected to a major ring or to a subring, there is the option to extend the subring control service through the major ring or not. This gives the following three options for subring connectivity:

1. **subring to a major ring or to a subring with a virtual channel** — In this case, a data service on the major ring or subring is created which is used to forward the R-APS messages for the subring over the major ring or subring, between the interconnection points of the subring to the major ring or subring. This allows the subring to operate as a fully connected ring and is mandatory if the subring connects two major rings or subrings because the virtual channel is the only mechanism that the subrings can use to exchange control messages. It also could improve failover times if the subring was large as it provides two paths on the subring interconnection nodes to propagate the fault indication around the subring, whereas without a virtual channel the fault indication may need to traverse the entire subring. Each subring requires its own data service on the major ring or subring for the virtual channel.
2. **subring to a major ring or to a subring without a virtual channel** — In this case the subring is not fully connected and does not require any resources on the major ring or subring. This option requires that the R-APS messages are not blocked on the subring over its RPL.
3. **subring to a VPLS service** — This is similar to the preceding option, but it uses a VPLS service instead of a major ring or subring. In this option, subring failures can initiate the sending of an LDP MAC flush message into the VPLS service when spoke or MPLS mesh SDPs are used in the VPLS service.

Ethernet ring terminology

The implementation of Ethernet ring on SR OS uses a VPLS as the construct for a ring flow function (one for ETH_FF (solely for control) and one for each service_FF) and SAPs (on ports or LAGs) as ring links. The control VPLS must be a regular VPLS, but the data VPLS can be a regular VPLS, a PBB (B/I-) VPLS or a routed VPLS. The state of the data service SAPs is inherited from the state of the control

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

service SAPs. Table 2: Terminology comparison displays a comparison between the ITU-T and SR OS terminologies.

Table 2: Terminology comparison

ITU-T G.8032v2 terminology	SR OS terminology
ETH_FF	control vpls
service_FF	data vpls
east ring link	path a
west ring link	path b
RPL owner	rpl-node owner
RPL link	path {a|b} rpl-end
MEP	control-mep
ERP control process	eth-ring instance or ring-id
major ring	eth-ring
sub-ring	eth-ring sub-ring
ring node	ring node PE
ring-ID	not used; fixed at 1 per G.8032v2

There are various ways that multiple rings can be interconnected and the possible topologies may be large. Customers typically have two forms of networks: access ring edge networks or larger multiple ring networks. Both topologies require ring interconnection.

Figure 5: G.8032 major ring and subring shows a ring of six nodes, with a major ring (regular Ethernet ring) on the top four nodes and a subring on the bottom.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 5: G.8032 major ring and subring

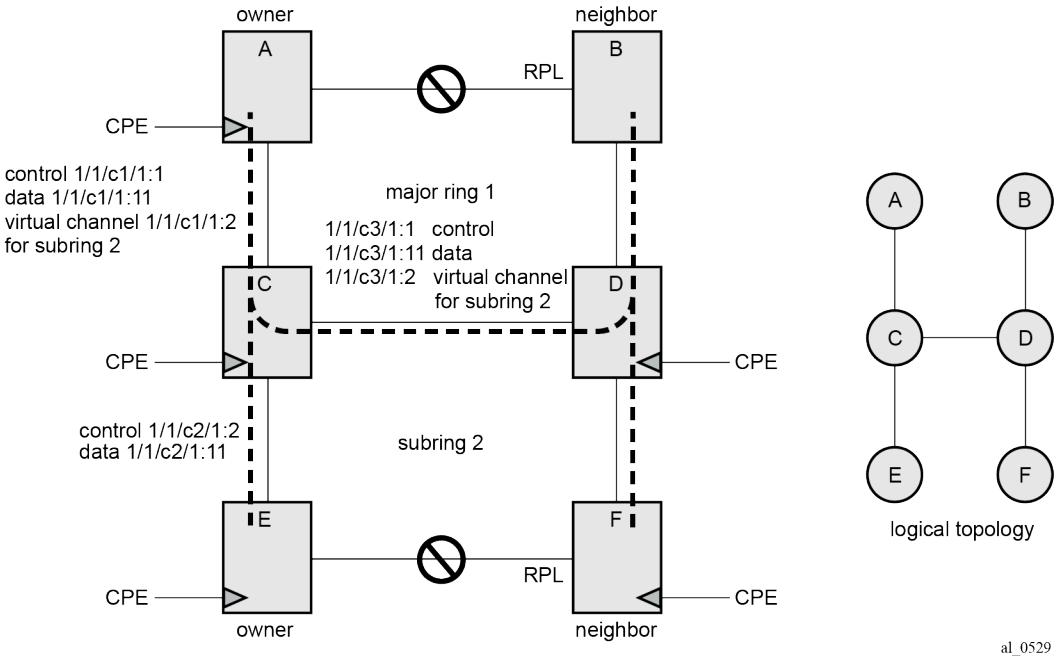

A major ring is a fully connected ring. A subring is a partial ring that depends on a major ring or a VPLS topology for part of the ring interconnect. Two major rings can be connected by a single subring. A subring can support other subrings.

In the major ring (on nodes A, B, C, and D), one path of the RPL owner is designated to be the RPL and the respective SAPs will be blocked in order to prevent a loop. The choice of where to put the RPL is up to the network administrator and can be different for different control instances of the ring allowing an RPL to be used for some other ring's traffic. In the subring, one path is designated as the RPL and will be blocked. Both the major ring and the subring have their own RPL. The subring interconnects to the major ring on nodes C and D and has a virtual channel on the major ring. SR OS supports both virtual channel and non-virtual channel rings. Schematics of the physical and logical topologies are also shown in Figure 5: G.8032 major ring and subring.

The G.8032 protocol defines a ring ID (1-255). The SR OS implementation only uses ring ID 1, which complies with G.8032v2. The configuration on a node uses a ring instance with a number but all rings use ring ID 1. This ring instance number is purely local and does not have to match on other ring nodes. Only the VLAN ID must match between SR OS ring nodes. For consistency in this example, VPLS instances and Ethernet ring instances are shown as matching for the same ring.

An RPL owner and RPL neighbor are configured for both the major ring and subring. The path and associated link will be the RPL when the ring is fully operational and will be blocked by the RPL owner whenever there is no fault on other ring links. Each ring RPL is independent. If a different ring link fails, then the RPL will be unblocked by the RPL owner. The link shared between a subring and the major ring is completely controlled by the major ring as if the subring were not there. Each ring can completely protect one fault within its ring. When the failed link recovers, it will initially be blocked by one of its adjacent nodes. The adjacent node sends an R-APS message across the ring to indicate the error is cleared and after a configurable time, if reversion is enabled, the RPL will revert to being blocked with all other links unblocked. This ensures that the ring topology when fully operational is predictable.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

If a specific RPL owner is not configured (not recommended by G.8032 specification), then the last link to become active will be blocked and the ring will remain in this state until another link fails. This operation makes the selection of the blocked link non-deterministic.

The protection protocol uses a specific control VLAN, with the associated data VLANs taking their forwarding state from the control VLAN. The control VLAN cannot carry data.

Load balancing with multiple ring instances

Each control ring is independent of the other control rings on the same topology. Therefore, because the RPL is used by one control ring, it is often desirable to set up a second control ring that uses a different link as RPL. This spreads out traffic in the topology, but if there is a link failure in the ring, all traffic will be on the remaining links. In the following examples, only a single control ring instance is configured. Other control and data rings could be configured if desired.

Provider backbone bridging (PBB)

PBB services also support G.8032 as data services (the services used for the control VPLS must be a regular VPLS). B/I-VPLS rings support both major rings and subrings. B-VPLS rings support multi-chassis link aggregation group (MC-LAG) as a dual homing option when aggregating I-VPLS traffic onto a B-VPLS ring. In other words, I-VPLS rings should not be dual-homed into two backbone edge bridge (BEB) nodes where the B-VPLS uses G.8032 to get connected to the rest of the B-VPLS network because the only mechanism that can propagate MAC flushes between an I-VPLS and B-VPLS is an LDP MAC flush.

SR OS implementation

G.8032 is built from VPLS components and each ring consists of the configuration components illustrated in Figure 6: G.8032 ring components .

Figure 6: G.8032 ring components

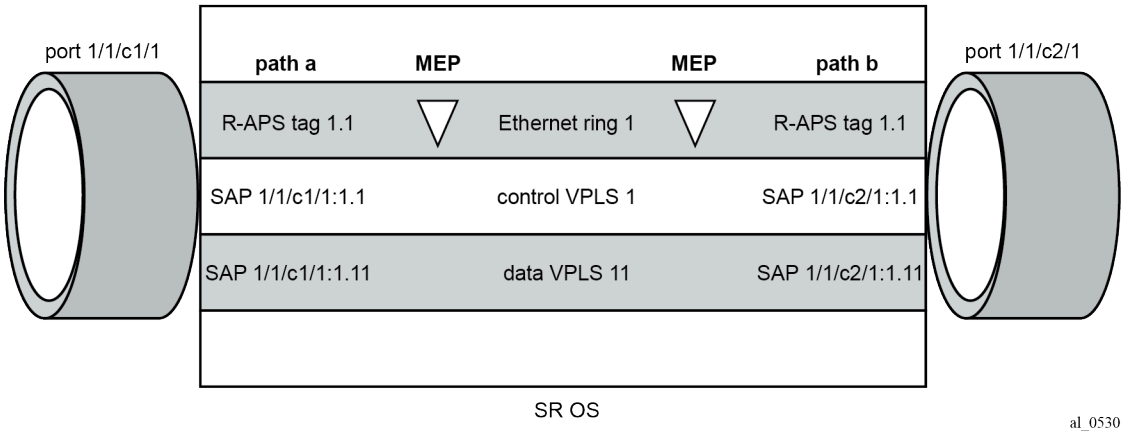

These components consist of:

- the Ethernet ring instance which defines the R-APS tags, the MEPs, and the ring behavior.
- the control VPLS which has SAPs with an encapsulation that matches the R-APS tags.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- the data VPLS which is linked to the ring. All of the data VPLS SAPs follow the operational state of the control VPLS SAPs in that each blocked SAP controlled by the ring is blocked for all control and data instances.

Figure 7: G.8032 subring interconnection components shows the major ring and subring interconnection components:

Figure 7: G.8032 subring interconnection components

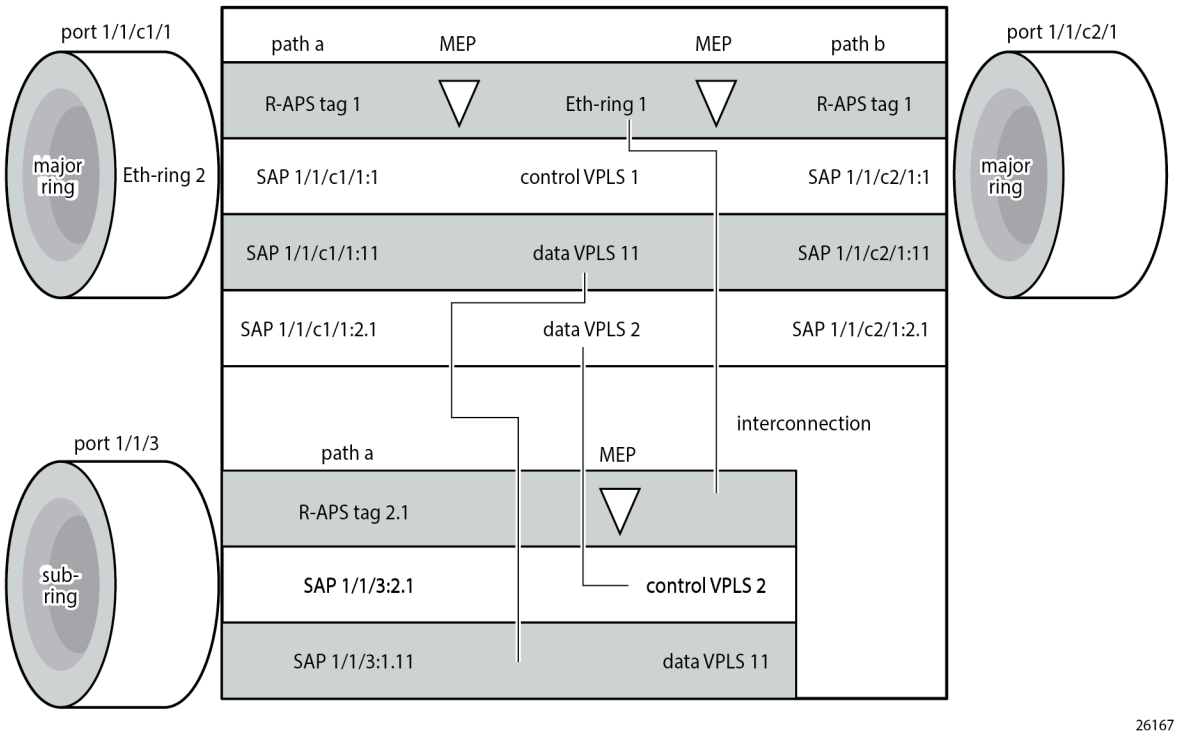

For a subring, the configuration is the same as a single ring except at the junction of the major ring and the subring. The interconnection of a subring and a major ring links the control VPLS of the subring to a data VPLS of the major ring when a virtual link is used. Similarly, the data VPLS of the subring is linked to a data VPLS of the major ring. Figure 7: G.8032 subring interconnection components illustrates the relationship of a subring and a major ring. Because this subring has a virtual channel, the data VPLS 2 has both data SAPs from the subring and data SAPs from the major ring. The virtual channel is also optional and in non-virtual-link cases, no VPLS instance is required (see non-virtual-link in the section Configuration of a subring to a VPLS service).

In Figure 7: G.8032 subring interconnection components, the inner tag values are kept the same for clarity, but in fact any encapsulation that is consistent with the next ring link will work. In other words, ring SAPs can perform VLAN ID translation and even when connecting a subring to a major ring. This also means that other ports may reuse the same tags when connecting independent services.

The R-APS tags and SAPs on the rings can either be dot1Q or QinQ encapsulated. It is also possible to have the control VPLS using single tagged frames with the data VPLSs using double tagged frames; this requires the system to be configured with the **extended-default-qinq-sap-lookup** parameter (**configure service system extended-default-qinq-sap-lookup true**), with the ring path R-APS tags and control VPLS SAPs configured as qtag.0, and the data VPLSs configured as QinQ SAP: qtag1.qtag2. Spanning tree protocol (STP) cannot be enabled on SAPs connected to Ethernet rings.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

R-APS messages received from other nodes are normally blocked on the RPL interface but the subring case with non-virtual channel recommends that R-APS messages be propagated over the RPL. Configuring **sub-ring type non-virtual-link** on all nodes on the subring is required to ensure propagation of R-APS messages around the subring.

R-APS messages are forwarded out of the egress using forwarding class network control (NC) and should be prioritized accordingly in the SAP egress QoS policy to ensure that congestion does not cause R-APS messages to be dropped which could cause the ring to switch to another path.

Configuration

This section describes the configuration of multiple rings. The Ethernet ring configuration commands are as follows.

```
configure {
    eth-ring <ring-index [1..128]> {
        ccm-hold-time {
            down <number>   # Hold timer for down event dampening in centiseconds
            up <number>     # Hold timer for recovery reporting in deciseconds
        }
        compatible-version <number>            # [1..2] - Default: 2
        description <string>
        guard-time <number>       # [1..20] in deciseconds - Default: 5
        node-id <mac-address>     # MAC address of the RPL <xx:xx:xx:xx:xx:xx>
        path <string>             # path ID: string of 1 character
            admin-state <keyword>   # default: disable
            description <string>
            port-and-raps-tag <port-and-encap>     # Port ID and ring APS tag ID
            eth-cfm {
                mep md-admin-name <reference> ma-admin-name <reference> mep-id <number> {
                    admin-state <keyword>   # default: disable
                    ccm <boolean>           # default: false
                    control-mep <boolean>   # default: false
                }
            }
            rpl-end <boolean>               # default: false
        }
        revert-time <number>      # <0,60..720> in seconds - Default: 300
        rpl-node <keyword>        # owner | neighbor
        sub-ring
            interconnect {
                propagate-topology-change <boolean>     # default: false
                ring-id      # Ring instance of the connection ring for the subring
                vpls         # Connect subring to VPLS ID that contains subring SAP
            }
            type             # Subring type (virtual-link|non-virtual-link)
        }
```

Parameters:

- **<ring-index>** — The ring index is the number by which the ring is referenced; values: 1 to128.
- **ccm-hold-time { [down <down-timeout>] [up <up-timeout>] }**
  - **down** — This command specifies the timer which controls the delay between detecting that ring path is down and reporting it to the G.8032 protection module. If a non-zero value is configured, the system will wait for the time specified in the value parameter before reporting it to the G.8032 protection module. This parameter applies only to ring path CCM. It does not apply to the ring port link state. To dampen ring port link state transitions, use the hold-time parameter from the physical

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

member port. This is useful if the underlying path between two nodes is going across an optical system which implements its own protection.

– **up** — This command specifies the timer which controls the delay between detecting that ring path is up and reporting it to the G.8032 protection module. If a non-zero value is configured, the system will wait for the time specified in the value parameter before reporting it to the G.8032 protection module. This parameter applies only to ring path CCM. It does not apply to the member port link state. To dampen member port link state transitions, use the hold-time parameter from the physical member port.

  Values:

```
*[ex:/configure eth-ring 1 ccm-hold-time]
A:admin@PE-1# down ?

 down <number>
 <number> - <1..5000> - centiseconds

    Hold timer for down event dampening

*[ex:/configure eth-ring 1 ccm-hold-time]
A:admin@PE-1# up ?

 up <number>
 <number> - <0..5000> - deciseconds
 Default  - 20

    Hold timer for recovery reporting
```

- The **admin-state** command allows to enable or disable the Ethernet ring.
- The **compatible-version** command configures the Ethernet ring compatibility version for the G.8032 state machine and messages. The default is version 2 (ITU G.8032v2) and all SR OS systems use version 2. If there is a need to interwork with third party devices that only support version 1, this can be set to version 1 allowing the reception of version 1 PDUs. Version 2 is encoded as 1 in the R-APS messages. Compatibility allows the reception of version 1 (encoded as 0) R-APS PDUs but, as per the G.8032 specification, higher versions are ignored on reception. For SR OS, messages are always originated with version 2. Therefore, if a third party switch supports version 3 (encoded as 2) or higher, interworking is also supported provided the other switch is compatible with version 2.
- The **description** includes a text string of maximum 80 characters to describe the use of the Ethernet ring.
- **guard-time** *<time>* — The forwarding method, in which R-APS messages are copied and forwarded at every Ethernet ring node, can result in a message corresponding to an old request, that is no longer relevant, being received by Ethernet ring nodes. Reception of an old R-APS message may result in erroneous ring state interpretation by some Ethernet ring nodes. The guard timer is used to prevent Ethernet ring nodes from acting upon outdated R-APS messages and prevents the possibility of forming a closed loop. Messages are not forwarded when the guard-timer is running.

  Values:

```
*[ex:/configure eth-ring 1]
A:admin@PE-1# guard-time ?

 guard-time <number>
 <number> - <1..20> - deciseconds
 Default  - 5
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- The **node-id <xx-xx-xx-xx-xx-xx>** allows the node identifier to be explicitly configured. By default, the chassis MAC is used. The node ID is not required in typical configurations.
- The **path** parameter defines the paths around the ring, of which there are two in different directions on the ring: an "a" path and a "b" path, except on the interconnection node where a subring connects to another major ring or subring in which case there is one path (either a or b) configured together with the **sub-ring** command. The paths are configured on a dot1Q or QinQ encapsulated access or hybrid port or a LAG with the encapsulation used for the R-APS messages on the ring. These can be either single tagged or double tagged.
  - The **admin-state** command allows to enable or disable the path.
  - The **port-and-raps-tag** specifies the port ID and the R-APS tag.
  - The **description** includes a text string of maximum 80 characters to describe the use of the path.
  - The **eth-cfm** context includes the associated Ethernet CFM parameters.
    - **mep md-domain-name** *<reference>* **ma-domain-name** *<reference>* **mep-id** *<number>* — The MEP defined under the path is used for the G.8032 protocol messages, which are based on IEEE 802.1ag/Y.1731 CFM frames.
  - **rpl-end** — When configured, this path is expected to be one end of the RPL. This parameter must be configured in conjunction with the **rpl-node** parameter.
- The **revert-time** command configures the revert time for an Ethernet ring. The revert time is the time that the RPL will wait before returning to the blocked state, after a failure condition has been fixed. Values: [0, 60..720] in seconds - Default: 300.
- When the **rpl-node** parameter is configured, a node can be designated as either the **owner** of the RPL, in which case this node is responsible for the RPL, or the **neighbor**, in which case the node is expected to be the neighbor to the RPL owner across the RPL. The **neighbor** parameter is optional and is included to be compliant with the specification. The **rpl-node** parameter must be configured in conjunction with the **rpl-end** command. On a subring without virtual channel it is mandatory to configure **sub-ring type non-virtual-link** on all nodes on the subring to ensure propagation of the R-APS messages around the subring.
- The **sub-ring** command is configured on the interconnection node between the subring and its major ring or subring to indicate that this ring is a subring. A ring configured as a subring can only be configured with a single path.
  - The **type {virtual-link|non-virtual-link}** parameter specifies whether it uses a virtual link through the major ring or subring for the R-APS messages or not.
  - **interconnect [ring-id** *<ring-index>* **| vpls]** — A subring connects to either another ring or to a VPLS service. If it connects to another ring (either a major ring or another subring), the ring identifier must be specified and the ring to which it connects must be configured with both a path "a" and a path "b", meaning that it is not possible to connect a subring to another subring on an interconnection node. Alternatively, the **vpls** parameter is used to indicate the subring connects to a VPLS service. Interconnection using a VPLS service requires the subring to be configured with **type non-virtual-link**.
    - **propagate-topology-change** — If a topology change event happens in the subring, it can be optionally propagated with the use of this parameter to either the major ring or subring it is connected to, using R-APS messages, or to the LDP VPLS SDP peers using an LDP "flush-all-from-me" message if the subring is connected to a VPLS service.

The example topology is shown in Figure 8: Ethernet example topology.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 8: Ethernet example topology

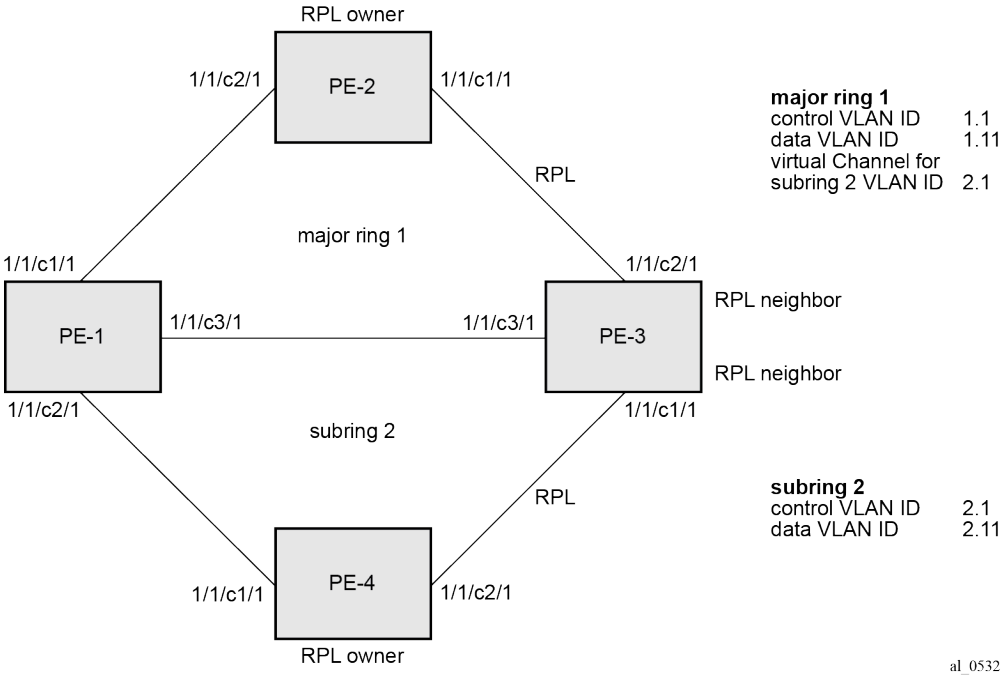

The configuration is divided into the following sections:

- a subring connected to a major ring using a virtual link through the major ring
- a subring connected to a major ring without a virtual link
- a subring connected to a VPLS service (without a virtual link)

Configure a subring to a major ring with a virtual link

To configure an Ethernet ring using R-APS, there will be at least two VPLS services required for one Ethernet ring instance, one for the control channel and the others for data channels. The control channel is used for R-APS signaling while the data channel is for user data traffic. The state of the data channels is inherited from the state of the control channel.

The following needs to be configured:

- encapsulation type for each ring port
- Ethernet CFM
- Ethernet ring for major ring 1
- Ethernet ring for subring 2
- control channel service and Ethernet ring SAPs
- user data channel services

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Configure the encapsulation for the ring ports.

An Ethernet ring needs an R-APS tag to send and receive G.8032 signaling messages. To configure a control channel, an access SAP configuration is required on each path (a or b) port. The SAP configuration follows that of the port and must be either dot1Q or QinQ, consequently the control and data packets are either single tagged or double tagged.

Note:
Single tagged control frames are supported on a QinQ port by configuring the system with the **extended-default-qinq-sap-lookup** parameter (**configure service system extended-default-qinq-sap-lookup true**), and the ring path R -APS tags and control VPLS SAPs configured as qtag.0.

In this example, QinQ tags are used. For example, the port configuration on PE-1 is as follows:

```
# on PE-1:
configure {
    port 1/1/c1/1 {
        admin-state enable
        ethernet {
            mode access
            encap-type qinq
        }
    }
    port 1/1/c2/1 {
        admin-state enable
        ethernet {
            mode access
            encap-type qinq
        }
    }
    port 1/1/c3/1 {
        admin-state enable
        ethernet {
            mode access
            encap-type qinq
        }
    }
```

Configure Ethernet CFM

Configuring the Ethernet CFM domain, association, and MEP is required before configuring an Ethernet ring. The standard domain format is **none** and the association name must be ITU carrier code-based (ICC-based - Y.1731); however, the SR OS implementation is flexible in that it supports both IEEE and ICC formats. The Ethernet ring MEP requires a CCM interval with values such as 1s, 100ms, or 10ms to be configured.

The MEPs used for R-APS control normally will have CCM configured on the control channel path MEPs for failure detection. Alternatively, detecting a failure of the ring may be achieved by running Ethernet in the first mile (EFM) at the port level if CCM is not possible at 1s, 100ms, or 10ms. Also rings can be run without CFM although the Ethernet CFM association must be configured for R-APS messages to be exchanged. To omit the failure detecting CCMs, it is necessary to remove the **ccm true** from under the path MEPs and to remove the **remote-mep** on the corresponding ETH CFM configuration.

Loss-of-signal, in conjunction with other OAM mechanisms, is applicable only when the nodes are directly connected.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 9: ETH-CFM MEP associations shows the details of the MEPs and their associations configured when both the major rings and subrings are used. The associations only need to be pairwise unique but for clarity five unique associations are used. Any name format can be used, but it must be consistent on both adjacent nodes.

Figure 9: ETH-CFM MEP associations

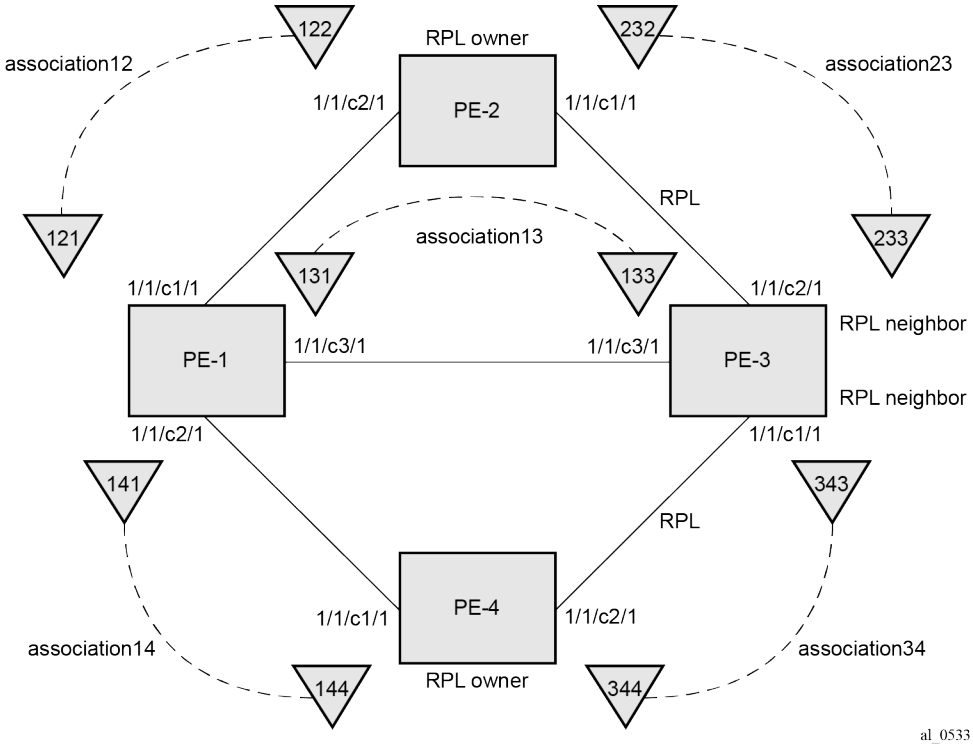

The configuration of Ethernet CFM for the major and subrings on each node is as follows. The CCMs for failure detection are configured for 1 second intervals.

On ring node PE-1, the associations "association-12" and "association-13" are used for the major ring and association "association-14" is used for the subring.

```
# on PE-1:
configure {
    eth-cfm {
        domain "domain-1" {
            level 2
            format none
            md-index 1
            association "association-12" {
                icc-based "Association12"
                ma-index 12
                ccm-interval 1s
                remote-mep 122 {
                }
            }
            association "association-13" {
                icc-based "Association13"
                ma-index 13
                ccm-interval 1s
                remote-mep 133 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
            }
            association "association-14" {
                icc-based "Association14"
                ma-index 14
                ccm-interval 1s
                remote-mep 144 {
                }
            }
        }
```

On ring node PE-2, the associations "association-12" and "association-23" are used for the major ring.

```
# on PE-2:
configure {
    eth-cfm {
        domain "domain-1" {
            level 2
            format none
            md-index 1
            association "association-12" {
                icc-based "Association12"
                ma-index 12
                ccm-interval 1s
                remote-mep 121 {
                }
            }
            association "association-23" {
                icc-based "Association23"
                ma-index 23
                ccm-interval 1s
                remote-mep 233 {
                }
            }
        }
```

On ring node PE-3, the associations "association-13" and "association-23" are used for the major ring and association "association-34" is used for the subring.

```
# on PE-3:
configure {
    eth-cfm {
        domain "domain-1" {
            level 2
            format none
            md-index 1
            association "association-13" {
                icc-based "Association13"
                ma-index 13
                ccm-interval 1s
                remote-mep 131 {
                }
            }
            association "association-23" {
                icc-based "Association23"
                ma-index 23
                ccm-interval 1s
                remote-mep 232 {
                }
            }
            association "association-34" {
                icc-based "Association34"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                ma-index 34
                ccm-interval 1s
                remote-mep 344 {
                }
            }
        }
```

On ring node PE-4, the associations 14 and 34 are used for the subring.

```
# on PE-4
configure {
    eth-cfm {
        domain "domain-1" {
            level 2
            format none
            md-index 1
            association "association-14" {
                icc-based "Association14"
                ma-index 14
                ccm-interval 1s
                remote-mep 141 {
                }
            }
            association "association-34" {
                icc-based "Association34"
                ma-index 34
                ccm-interval 1s
                remote-mep 343 {
                }
            }
        }
```

Configuring Ethernet ring – major ring 1

Two paths must be configured to form a ring. In this example, VLAN tag 1.1 is used as control channel for R-APS signaling for the major ring (Ethernet ring 1) on the ports shown in Figure 8: Ethernet example topology using the Ethernet CFM information shown in Figure 9: ETH-CFM MEP associations. The revert time is set to the value of 60 seconds and CCM messages are enabled on the MEP. The **control-mep true** command indicates that this MEP is used for ring R-APS messages.

The configuration of Ethernet ring 1 on ring node PE-1 is as follows:

```
# on PE-1:
configure {
    eth-ring 1 {
        admin-state enable
        description "Ethernet ring 1_major ring"
        revert-time 60
        path "a" {
            admin-state enable
            description "Ethernet ring 1_path a"
            port-and-raps-tag 1/1/c1/1:1.1
            eth-cfm {
                mep md-admin-name "domain-1" ma-admin-name "association-12" mep-id 121 {
                    admin-state enable
                    ccm true
                    control-mep true
                }
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        path "b" {
            admin-state enable
            description "Ethernet ring 1_path b"
            port-and-raps-tag 1/1/c3/1:1.1
            eth-cfm {
                mep md-admin-name "domain-1" ma-admin-name "association-13" mep-id 131 {
                    admin-state enable
                    ccm true
                    control-mep true
                }
            }
        }
```

It is mandatory to configure a MEP in the path context, otherwise the following error is displayed:

```
*[ex:/configure eth-ring 1 path "a"]
A:admin@PE-1# commit
MINOR: MGMT_CORE #4001: configure eth-ring 1 path "a" admin-state - Cannot enable path
without eth-cfm mep configured - configure eth-ring 1 path "a" admin-state
```

While MEPs are mandatory, enabling CCMs on the MEPs under the paths as a failure detection mechanism is optional as explained earlier.

Ring node PE-2 is configured as the RPL owner with the RPL being on path "a" as indicated by the **rpl-end** parameter. The revert time is 60 seconds.

```
# on PE-2:
configure {
    eth-ring 1 {
        admin-state enable
        description "Ethernet ring 1_major ring"
        revert-time 60
        rpl-node owner
        path "a" {
            admin-state enable
            description "Ethernet ring 1_path a"
            port-and-raps-tag 1/1/c1/1:1.1
            rpl-end true
            eth-cfm {
                mep md-admin-name "domain-1" ma-admin-name "association-23" mep-id 232 {
                    admin-state enable
                    ccm true
                    control-mep true
                }
            }
        }
        path "b" {
            admin-state enable
            description "Ethernet ring 1_path b"
            port-and-raps-tag 1/1/c2/1:1.1
            eth-cfm {
                mep md-admin-name "domain-1" ma-admin-name "association-12" mep-id 122 {
                    admin-state enable
                    ccm true
                    control-mep true
                }
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

It is not permitted to configure a path as an RPL end without having configured the node on this ring to be either the RPL owner or RPL neighbor, otherwise the following error message is reported.

```
*[ex:/configure eth-ring 1 path "a"]
A:admin@PE-2# commit
MINOR: MGMT_CORE #4001: configure eth-ring 1 path "a" rpl-end - rpl-node must be set -
 configure eth-ring 1 rpl-node
```

Ring node PE-3 is configured as the RPL neighbor with the RPL being on path "b" as indicated by the **rpl-end** parameter. The revert time is 60 seconds.

```
# on PE-3:
configure {
    eth-ring 1 {
        admin-state enable
        description "Ethernet ring 1_major ring"
        revert-time 60
        rpl-node neighbor
        path "a" {
            admin-state enable
            description "Ethernet ring 1_path a"
            port-and-raps-tag 1/1/c3/1:1.1
            eth-cfm {
                mep md-admin-name "domain-1" ma-admin-name "association-13" mep-id 133 {
                    admin-state enable
                    ccm true
                    control-mep true
                }
            }
        }
        path "b" {
            admin-state enable
            description "Ethernet ring 1_path b"
            port-and-raps-tag 1/1/c2/1:1.1
            rpl-end true
            eth-cfm {
                mep md-admin-name "domain-1" ma-admin-name "association-23" mep-id 233 {
                    admin-state enable
                    ccm true
                    control-mep true
                }
            }
        }
```

The link between PE-2 and PE-3 will be the RPL with PE-2 and PE-3 blocking that link when the ring is fully operational. In this example, the RPL is using path "a" on PE-2 and path "b" on PE-3.

Configuring Ethernet ring – subring 2

Ring nodes PE-1, PE-3, and PE-4 form a subring. The subring attaches to the major ring (Ethernet ring 1). The subring in this case uses a virtual link. The interconnection ring instance identifier (**ring-id 1**) is specified and **propagate-topology-change true** indicates that subring flushing will be propagated to the major ring. Only one path (path a) is specified because the other path (path b) is not required at an interconnection node. Subrings are almost identical to major rings in operation except that subrings send MAC flushes towards their connected ring (either a major ring or a subring). Major rings or subrings never send MAC flushes to their subrings. Therefore a couple of subrings connected to a major ring can

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

cause MACs to flush on the major ring but the major ring will not propagate a subring MAC flush to other subrings.

Ring node PE-1 provides an interconnection between the major ring (ring 1) and the subring (ring 2). Ring 2 is configured to be a subring which interconnects to ring 1. It will use a virtual link on ring 1 to send R-APS messages to the other interconnection node and topology changes will be propagated from subring 2 to the major ring 1.

```
# on PE-1:
configure {
    eth-ring 2 {
        admin-state enable
        description "Ethernet subring 2 on major ring 1"
        revert-time 60
        sub-ring {
            type virtual-link
            interconnect {
                ring-id 1
                propagate-topology-change true
            }
        }
        path "a" {
            admin-state enable
            description "Ethernet subring 2_path a"
            port-and-raps-tag 1/1/c2/1:2.1
            eth-cfm {
                mep md-admin-name "domain-1" ma-admin-name "association-14" mep-id 141 {
                    admin-state enable
                    ccm true
                    control-mep true
                }
            }
        }
```

Subring 2 is not configured on PE-2.

The configuration of subring 2 on PE-3 is similar to PE-1, but PE-3 is the RPL neighbor, with the RPL end on path "a", for the RPL between PE-3 and PE-4.

```
# on PE-3:
configure {
    eth-ring 2
        admin-state enable
        description "Ethernet subring 2 on major ring 1"
        revert-time 60
        rpl-node neighbor
        sub-ring {
            type virtual-link
            interconnect {
                ring-id 1
                propagate-topology-change true
            }
        }
        path "a" {
            admin-state enable
            description "Ethernet subring 2_path a"
            port-and-raps-tag 1/1/c1/1:2.1
            rpl-end true
            eth-cfm {
                mep md-admin-name "domain-1" ma-admin-name "association-34" mep-id 343 {
                    admin-state enable
                    ccm true
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    control-mep true
                }
            }
        }
```

Ring node PE-4 only has configuration for the subring 2, not for major ring 1. PE-4 is the RPL owner, with path "b" being the RPL end, for the RPL between PE-3 and PE-4.

```
# on PE-4:
configure {
    eth-ring 2
        admin-state enable
        description "Ethernet subring 2"
        revert-time 60
        rpl-node owner
        path "a" {
            admin-state enable
            description "Ethernet subring 2_path a"
            port-and-raps-tag 1/1/c1/1:2.1
            eth-cfm {
                mep md-admin-name "domain-1" ma-admin-name "association-14" mep-id 144 {
                    admin-state enable
                    ccm true
                    control-mep true
                }
            }
        }
        path "b" {
            admin-state enable
            description "Ethernet subring 2_path b"
            port-and-raps-tag 1/1/c2/1:2.1
            rpl-end true
            eth-cfm {
                mep md-admin-name "domain-1" ma-admin-name "association-34" mep-id 344 {
                    admin-state enable
                    ccm true
                    control-mep true
                }
            }
        }
```

Until the Ethernet ring instance is attached to a VPLS service, the ring operational status is down and the forwarding status of each port is blocked. This prevents the operator from creating a loop by misconfiguration. This state can be seen on ring node PE-1 as follows:

```
[/]
A:admin@PE-1# show eth-ring 1

===============================================================================
Ethernet Ring 1 Information
===============================================================================
Description        : Ethernet ring 1_major ring
Admin State        : Up                 Oper State       : Down
Node ID            : 02:09:ff:00:00:00
Guard Time         :    5 deciseconds   RPL Node         : rplNone
Max Revert Time    :   60 seconds       Time to Revert   : N/A
CCM Hold Down Time :    0 centiseconds  CCM Hold Up Time :   20 deciseconds
Compatible Version : 2
APS Tx PDU         : Request State: 0xB
                     Sub-Code     : 0x0
                     Status       : 0x20  ( BPR )
                     Node ID      : 02:09:ff:00:00:00
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Defect Status      :

Sub-Ring Type      : none

-------------------------------------------------------------------------------
Ethernet Ring Path Summary
-------------------------------------------------------------------------------
Path Port             Raps-Tag     Admin/Oper      Type            Fwd State
-------------------------------------------------------------------------------
  a  1/1/c1/1         1.1             Up/Down      normal          blocked
  b  1/1/c3/1         1.1             Up/Down      normal          blocked
===============================================================================
```

Configure the control channel VPLS service

Path "a" and "b" configured in the Ethernet ring must be added as SAPs into a VPLS service (standard VPLS) using the **eth-ring** parameter. The SAP encapsulation values must match the values of the R-APS tag configured for the associated path.

G.8032 uses the same R-APS tag value on all nodes on the ring, as configured in this example. However, the SR OS implementation relaxes this constraint by requiring the tag to match only on adjacent nodes.

In this example VPLS "control-VPLS-1" is configured on PE-1, PE-2, and PE-3 for the control channel for the major ring (ring 1), and VPLS "control-VPLS-2" is used on PE-1, PE-3, and PE-4 for the subring (ring 2).

VPLS "control-VPLS-1" is the control service for the major ring and is defined for PE-1, PE-2, and PE-3, as follows:

```
# on PE-1:
configure {
    service {
         vpls "control-VPLS-1" {
            admin-state enable
            description "Control channel VID 1.1 for ring 1 - major ring"
            service-id 1
            customer "1"
            sap 1/1/c1/1:1.1 {
                eth-ring 1
            }
            sap 1/1/c3/1:1.1 {
                eth-ring 1
            }
        }
```

```
# on PE-2:
configure {
    service {
         vpls "control-VPLS-1" {
            admin-state enable
            description "Control channel VID 1.1 for ring 1 - major ring"
            service-id 1
            customer "1"
            sap 1/1/c1/1:1.1 {
                eth-ring 1
            }
            sap 1/1/c2/1:1.1 {
                eth-ring 1
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
```

```
# on PE-3:
configure {
    service {
         vpls "control-VPLS-1" {
            admin-state enable
            description "Control channel VID 1.1 for ring 1 - major ring"
            service-id 1
            customer "1"
            sap 1/1/c2/1:1.1 {
                eth-ring 1
            }
            sap 1/1/c3/1:1.1 {
                eth-ring 1
            }
        }
```

SAPs or SDPs can be added to a control channel VPLS on condition the **eth-ring** parameter is present. Any attempt to add a SAP without the **eth-ring** parameter to a control channel VPLS results in the following messages being displayed.

```
*[ex:/configure service vpls "control-VPLS-1" sap 1/1/c4/1:1]
A:admin@PE-1# commit
MINOR: SVCMGR #2590: configure service vpls "control-VPLS-1" sap 1/1/c3/1:1.1 - Invalid
Ethernet ring configuration - Ring control SAP cannot be added to service that contains non-
ring SAPs or SDP bindings - configure service vpls "control-VPLS-1" sap 1/1/c3/1:1.1 eth-ring
MINOR: SVCMGR #2590: configure service vpls "control-VPLS-1" sap 1/1/c3/1:1.1 - Invalid
Ethernet ring configuration - Ethernet Ring Control service should only have controls saps
from same ring - configure service vpls "control-VPLS-1" sap 1/1/c3/1:1.1 eth-ring
MINOR: SVCMGR #2590: configure service vpls "control-VPLS-1" sap 1/1/c1/1:1.1 - Invalid
Ethernet ring configuration - Ring control SAP cannot be added to service that contains non-
ring SAPs or SDP bindings - configure service vpls "control-VPLS-1" sap 1/1/c1/1:1.1 eth-ring
MINOR: SVCMGR #2590: configure service vpls "control-VPLS-1" sap 1/1/c1/1:1.1 - Invalid
Ethernet ring configuration - Ethernet Ring Control service should only have controls saps
from same ring - configure service vpls "control-VPLS-1" sap 1/1/c1/1:1.1 eth-ring
```

For the subring, the configuration of a split horizon group for the virtual channel on the major ring on the interconnection nodes is recommended. This avoids the looping of control R-APS messages in the case there is a misconfiguration in the major ring.

On ring node PE-1, the control service for the subring "control-VPLS-2" is configured as follows. SAP 1/1/c1/1:2.1 and SAP 1/1/c3/1:2.1 connect to the major ring (ring 1) for the virtual channel, whereas SAP 1/1/c2/1:2.1 connects to the subring (ring 2).

```
# on PE-1:
configure {
    service {
        vpls "control-VPLS-2" {
            admin-state enable
            description "Control/virtual channel VID 2.1 for ring 2"
            service-id 2
            customer "1"
            split-horizon-group "shg-ring2" {
            }
            sap 1/1/c1/1:2.1 {
                description "ring 2 interconnection using ring 1"
                eth-ring 1
                split-horizon-group "shg-ring2"
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            sap 1/1/c2/1:2.1 {
                eth-ring 2
            }
            sap 1/1/c3/1:2.1 {
                description "ring 2 interconnection using ring 1"
                eth-ring 1
                split-horizon-group "shg-ring2"
            }
        }
```

On ring node PE-2, subring 2 is not present. However, the control service "control-VPLS-2" for the subring must be configured on PE-2, because the virtual channel for subring 2 needs to exist throughout major ring 1.

```
# on PE-2:
configure {
    service {
        vpls "control-VPLS-2" {
            admin-state enable
            description "virtual channel VID 2.1 for ring 2"
            service-id 2
            customer "1"
            sap 1/1/c1/1:2.1 {
                eth-ring 1
            }
            sap 1/1/c2/1:2.1 {
                eth-ring 1
            }
        }
```

If multiple virtual channels are used (due to the aggregation of multiple subrings into the same major ring), their configuration could be simplified on non-interconnection nodes on the major ring. To achieve this on a ring node such as PE-2, a default SAP could be used rather than configuring a VPLS per virtual channel. If QinQ SAPs are used then default SAPs 1/1/c1/1:qtag.* and 1/1/c2/1:qtag.* could be used but this requires all control channels for subrings to be using qtag as the outer VLAN ID, or 1/1/c1/1:* and 1/1/c2/1:* if dot1Q SAPs were used. This is because the SAPs match explicit SAP definitions first and the default SAP will handle any other traffic.

The following configuration for control service "control-VPLS-2" for the subring on ring node PE-3 is similar to the configuration of PE-1.

```
# on PE-3:
configure {
    service {
        vpls "control-VPLS-2" {
            admin-state enable
            description "control/virtual channel VID 2.1 for ring 2"
            service-id 2
            customer "1"
            split-horizon-group "shg-ring2" {
            }
            sap 1/1/c1/1:2.1 {
                eth-ring 2
            }
            sap 1/1/c2/1:2.1 {
                description "ring 2 interconnection using ring 1"
                eth-ring 1
                split-horizon-group "shg-ring2"
            }
            sap 1/1/c3/1:2.1 {
                description "ring 2 interconnection using ring 1"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                eth-ring 1
                split-horizon-group "shg-ring2"
            }
        }
```

On ring node PE-4, control service "control-VPLS-2" for the subring is configured as follows. Both SAPs are configured on the subring (ring 2).

```
# on PE-4:
configure {
    service {
        vpls "control-VPLS-2" {
            admin-state enable
            description "control VID 2.1 for subring 2"
            service-id 2
            customer "1"
            sap 1/1/c1/1:2.1 {
                eth-ring 2
            }
            sap 1/1/c2/1:2.1 {
                eth-ring 2
            }
        }
```

At this point, the Ethernet ring 1 is operationally up and the RPL is blocking successfully RPL end port 1/1/c1/1 on RPL owner PE-2 and RPL end port 1/1/c2/1 on RPL neighbor PE-3.

Show output

An overview of all of the rings can be shown using the following commands, in this case on PE-1.

The following command shows the Ethernet ring status on PE-1.

```
[/]
A:admin@PE-1# show eth-ring status

===============================================================================
Ethernet Ring (Status information)
===============================================================================
Ring   Admin  Oper       Path Information                MEP Information
ID     State  State  Path         Tag       State     Ctrl-MEP CC-Intvl Defects
-------------------------------------------------------------------------------
1      Up     Up     a - 1/1/c1/1     1.1   Up        Yes      1         -----
                     b - 1/1/c3/1     1.1   Up        Yes      1         -----
2      Up     Up     a - 1/1/c2/1     2.1   Up        Yes      1         -----
                     b - N/A                -         -        -         -----
===============================================================================
Ethernet Tunnel MEP Defect Legend:
R = Rdi, M = MacStatus, C = RemoteCCM, E = ErrorCCM, X = XconCCM
```

It is expected that the state is "up", even on ring paths which are blocked. The "Defects" column refers to the CFM defects of the MEPs. If there is a problem, these will be flagged.

The following command shows the ring and path forwarding states on PE-1.

```
[/]
A:admin@PE-1# show eth-ring

===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Ethernet Rings (summary)
===============================================================================
Ring Int  Admin Oper           Paths Summary                      Path States
ID   ID   State State                                             a     b
-------------------------------------------------------------------------------
1    -    Up    Up    a - 1/1/c1/1     1.1   b - 1/1/c3/1     1.1   U     U
2    1    Up    Up    a - 1/1/c2/1     2.1   b - Not configured     U     -
===============================================================================
Ethernet Ring Summary Legend:   B - Blocked    U - Unblocked
```

The following command shows specific information for major ring 1 on ring node PE-1:

```
[/]
A:admin@PE-1# show eth-ring 1

===============================================================================
Ethernet Ring 1 Information
===============================================================================
Description        : Ethernet ring 1_major ring
Admin State        : Up                 Oper State       : Up
Node ID            : 02:09:ff:00:00:00
Guard Time         :    5 deciseconds   RPL Node         : rplNone
Max Revert Time    :   60 seconds       Time to Revert   : N/A
CCM Hold Down Time :    0 centiseconds  CCM Hold Up Time :   20 deciseconds
Compatible Version : 2
APS Tx PDU         : N/A
Defect Status      :

Sub-Ring Type      : none

-------------------------------------------------------------------------------
Ethernet Ring Path Summary
-------------------------------------------------------------------------------
Path Port             Raps-Tag     Admin/Oper      Type            Fwd State
-------------------------------------------------------------------------------
  a  1/1/c1/1         1.1            Up/Up        normal          unblocked
  b  1/1/c3/1         1.1            Up/Up        normal          unblocked
===============================================================================
```

The status around the major ring can also be checked.

The following command shows specific information for major ring 1 on RPL owner PE-2:

```
[/]
A:admin@PE-2# show eth-ring 1

===============================================================================
Ethernet Ring 1 Information
===============================================================================
Description        : Ethernet ring 1_major ring
Admin State        : Up                 Oper State       : Up
Node ID            : 02:0b:ff:00:00:00
Guard Time         :    5 deciseconds   RPL Node         : rplOwner
Max Revert Time    :   60 seconds       Time to Revert   : N/A
CCM Hold Down Time :    0 centiseconds  CCM Hold Up Time :   20 deciseconds
Compatible Version : 2
APS Tx PDU         : Request State: 0x0
                     Sub-Code     : 0x0
                     Status       : 0x80  ( RB )
                     Node ID      : 02:0b:ff:00:00:00
Defect Status      :

Sub-Ring Type      : none
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
Ethernet Ring Path Summary
-------------------------------------------------------------------------------
Path Port             Raps-Tag     Admin/Oper      Type            Fwd State
-------------------------------------------------------------------------------
  a  1/1/c1/1         1.1             Up/Up        rplEnd          blocked
  b  1/1/c2/1         1.1             Up/Up        normal          unblocked
===============================================================================
```

PE-2 is the RPL owner with port 1/1/c1/1 as an RPL end, which is blocked as expected. The revert time is also shown to be the configured value of 60 seconds. Detailed information is shown relating to the R-APS PDUs being transmitted on this ring because PE-2 is the RPL owner.

When a revert is pending after a link failure has been removed, the "Time to Revert" will show the number of seconds remaining before the revert occurs.

The following command shows specific information for major ring 1 on RPL neighbor PE-3:

```
[/]
A:admin@PE-3# show eth-ring 1

===============================================================================
Ethernet Ring 1 Information
===============================================================================
Description        : Ethernet ring 1_major ring
Admin State        : Up                 Oper State       : Up
Node ID            : 02:0d:ff:00:00:00
Guard Time         :    5 deciseconds   RPL Node         : rplNeighbor
Max Revert Time    :   60 seconds       Time to Revert   : N/A
CCM Hold Down Time :    0 centiseconds  CCM Hold Up Time :   20 deciseconds
Compatible Version : 2
APS Tx PDU         : N/A
Defect Status      :

Sub-Ring Type      : none

-------------------------------------------------------------------------------
Ethernet Ring Path Summary
-------------------------------------------------------------------------------
Path Port             Raps-Tag     Admin/Oper      Type            Fwd State
-------------------------------------------------------------------------------
  a  1/1/c3/1         1.1             Up/Up        normal          unblocked
  b  1/1/c2/1         1.1             Up/Up        rplEnd          blocked
===============================================================================
```

PE-3 is the RPL neighbor with port 1/1/c2/1 as an RPL end which is blocked as expected.

The information for the subring can also be shown using a similar command. The following command shows specific information for subring 2 on ring node PE-1:

```
[/]
A:admin@PE-1# show eth-ring 2

===============================================================================
Ethernet Ring 2 Information
===============================================================================
Description        : Ethernet subring 2 on major ring 1
Admin State        : Up                 Oper State       : Up
Node ID            : 02:09:ff:00:00:00
Guard Time         :    5 deciseconds   RPL Node         : rplNone
Max Revert Time    :   60 seconds       Time to Revert   : N/A
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
CCM Hold Down Time :    0 centiseconds  CCM Hold Up Time :   20 deciseconds
Compatible Version : 2
APS Tx PDU         : N/A
Defect Status      :

Sub-Ring Type      : virtualLink        Interconnect-ID  : 1
Topology Change    : Propagate

-------------------------------------------------------------------------------
Ethernet Ring Path Summary
-------------------------------------------------------------------------------
Path Port             Raps-Tag     Admin/Oper      Type            Fwd State
-------------------------------------------------------------------------------
  a  1/1/c2/1         2.1             Up/Up        normal          unblocked
  b  -                -                 -/-         -               -
===============================================================================
```

Only path "a" is active and unblocked. Path "b" is not configured because only one path is required on an interconnection node. The "Sub-Ring Type" is shown to be a virtual link interconnecting to ring 1, with topology propagation enabled.

The following command shows specific information for subring 2 on ring node PE-3:

```
[/]
A:admin@PE-3# show eth-ring 2

===============================================================================
Ethernet Ring 2 Information
===============================================================================
Description        : Ethernet subring 2 on major ring 1
Admin State        : Up                 Oper State       : Up
Node ID            : 02:0d:ff:00:00:00
Guard Time         :    5 deciseconds   RPL Node         : rplNeighbor
Max Revert Time    :   60 seconds       Time to Revert   : N/A
CCM Hold Down Time :    0 centiseconds  CCM Hold Up Time :   20 deciseconds
Compatible Version : 2
APS Tx PDU         : N/A
Defect Status      :

Sub-Ring Type      : virtualLink        Interconnect-ID  : 1
Topology Change    : Propagate

-------------------------------------------------------------------------------
Ethernet Ring Path Summary
-------------------------------------------------------------------------------
Path Port             Raps-Tag     Admin/Oper      Type            Fwd State
-------------------------------------------------------------------------------
  a  1/1/c1/1         2.1             Up/Up        rplEnd          blocked
  b  -                -                 -/-         -               -
===============================================================================
```

PE-3 is the RPL neighbor with port 1/1/c1/1 as an RPL end, which is blocked as expected.

The following command shows specific information for subring 2 on ring node PE-4:

```
[/]
A:admin@PE-4# show eth-ring 2

===============================================================================
Ethernet Ring 2 Information
===============================================================================
Description        : Ethernet subring 2
Admin State        : Up                 Oper State       : Up
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Node ID            : 02:0f:ff:00:00:00
Guard Time         :    5 deciseconds   RPL Node         : rplOwner
Max Revert Time    :   60 seconds       Time to Revert   : N/A
CCM Hold Down Time :    0 centiseconds  CCM Hold Up Time :   20 deciseconds
Compatible Version : 2
APS Tx PDU         : Request State: 0x0
                     Sub-Code     : 0x0
                     Status       : 0xE0  ( RB DNF BPR )
                     Node ID      : 02:0f:ff:00:00:00
Defect Status      :

Sub-Ring Type      : none

-------------------------------------------------------------------------------
Ethernet Ring Path Summary
-------------------------------------------------------------------------------
Path Port             Raps-Tag     Admin/Oper      Type            Fwd State
-------------------------------------------------------------------------------
  a  1/1/c1/1         2.1             Up/Up        normal          unblocked
  b  1/1/c2/1         2.1             Up/Up        rplEnd          blocked
===============================================================================
```

PE-4 is the RPL owner with port 1/1/c2/1 as an RPL end, which is blocked as expected.

The following command shows the details of an individual path.

```
[/]
A:admin@PE-1# show eth-ring 1 path a

===============================================================================
Ethernet Ring 1 Path Information
===============================================================================
Description        : Ethernet ring 1_path a
Port               : 1/1/c1/1           Raps-Tag         : 1.1
Admin State        : Up                 Oper State       : Up
Path Type          : normal             Fwd State        : unblocked
                                        Fwd State Change : 05/16/2023 08:14:44
Last Switch Command: noCmd
APS Rx PDU         : Request State: 0x0
                     Sub-Code     : 0x0
                     Status       : 0x80  ( RB )
                     Node ID      : 02:0b:ff:00:00:00

===============================================================================
```

The ring hierarchy created can be shown, either for all rings, or as follows for a specific ring.

```
[/]
A:admin@PE-1# show eth-ring 1 hierarchy

===============================================================================
Ethernet Ring 1 (hierarchy)
===============================================================================
Ring Int  Admin Oper           Paths Summary                      Path States
ID   ID   State State                                             a     b
-------------------------------------------------------------------------------
1    -    Up    Up    a - 1/1/c1/1     1.1   b - 1/1/c3/1     1.1   U     U
2    1    Up    Up    a - 1/1/c2/1     2.1   b - Not configured     U     -
===============================================================================
Ethernet Ring Summary Legend:   B - Blocked    U - Unblocked
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Configure the user data channel VPLS service

The user data channels are created on a separate VPLS, "VPLS-11" in this example, using VLAN tag 1.11. The ring data channels must be on the same ports as the corresponding control channels configured above. The access into the data services can use normal SAPs or SDPs, for example the SAP on port 1/1/c4/1 in the following output. Customer data traverses the ring on a data SAP. Multiple parallel data SAPs in different data services can be controlled by one control ring instance, Ethernet ring 1 in the example.

Data VPLS "VPLS-11" on ring node PE-1 has data SAPs 1/1/c1/1:1.11 and 1/1/c3/1:1.11 on major ring 1, while SAP 1/1/c2/1:1.11 is the data SAP on subring 2.

```
# on PE-1:
configure {
    service {
        vpls "VPLS-11" {
            admin-state enable
            description "data VPLS"
            service-id 11
            customer "1"
            sap 1/1/c1/1:1.11 {
                eth-ring 1
            }
            sap 1/1/c2/1:1.11 {
                eth-ring 2
            }
            sap 1/1/c3/1:1.11 {
                eth-ring 1
            }
            sap 1/1/c4/1:11 {
                description "sample customer service SAP"
            }
        }
```

The configuration of data VPLS "VPLS-11" on ring node PE-3 (not shown) is similar to ring node PE-1.

The configuration of data VPLS "VPLS-11" on ring node PE-2 has data SAPs 1/1/c1/1:1.11 and 1/1/c3/1:1.11 on major ring 1.

```
# on PE-2:
configure {
    service {
        vpls "VPLS-11" {
            admin-state enable
            description "data VPLS"
            service-id 11
            customer "1"
            sap 1/1/c1/1:1.11 {
                eth-ring 1
            }
            sap 1/1/c2/1:1.11 {
                eth-ring 1
            }
            sap 1/1/c4/1:11 {
                description "sample customer service SAP"
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The configuration of data VPLS "VPLS-11" on ring node PE-4 has data SAPs 1/1/c1/1:1.11 and 1/1/c3/1:1.11 on subring 2.

```
# on PE-4:
configure {
    service {
        vpls "VPLS-11" {
            admin-state enable
            description "data VPLS"
            service-id 11
            customer "1"
            sap 1/1/c1/1:1.11 {
                eth-ring 2
            }
            sap 1/1/c2/1:1.11 {
                eth-ring 2
            }
            sap 1/1/c4/1:11 {
                description "sample customer service SAP"
            }
        }
```

All the SAPs which are configured to use Ethernet rings can be displayed. The following output is taken from PE-1, where there are:

- two SAPs in VPLS 1 for the control channel of ring 1 (VLAN ID 1.1)
- two SAPs in VPLS 2 on ring 1 for the virtual channel for ring 2 (VLAN ID 2.1)
- one SAP in VPLS 2 on ring 2 for the control channel for ring 2 (VLAN ID 2.1)
- three SAPs in VPLS 11, two on ring 1 and one on ring 2, for the data service (VLAN ID 1.11). This matches the information in Figure 7: G.8032 subring interconnection components.

```
[/]
A:admin@PE-1# show service sap-using eth-ring

===============================================================================
Service Access Points (Ethernet Ring)
===============================================================================
SapId                   SvcId       Eth-Ring Path Admin Oper  Blocked Control/
                                                  State State         Data
-------------------------------------------------------------------------------
1/1/c1/1:1.1            1           1        a    Up    Up    No      Ctrl
1/1/c3/1:1.1            1           1        b    Up    Up    No      Ctrl
1/1/c1/1:2.1            2           1        a    Up    Up    No      Ctrl
1/1/c2/1:2.1            2           2        a    Up    Up    No      Ctrl
1/1/c3/1:2.1            2           1        b    Up    Up    No      Ctrl
1/1/c1/1:1.11           11          1        a    Up    Up    No      Data
1/1/c2/1:1.11           11          2        a    Up    Up    No      Data
1/1/c3/1:1.11           11          1        b    Up    Up    No      Data
-------------------------------------------------------------------------------
Number of SAPs : 8
===============================================================================
```

Statistics are available showing both the CCM and R-APS messages sent and received on a node. An associated **clear** command is available.

```
[/]
A:admin@PE-1# show eth-cfm statistics

===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
ETH-CFM System Statistics
===============================================================================
Rx Count           : 5973              Tx Count           : 6820
Dropped Congestion : 0                 Discarded Error    : 0
AIS Currently Act  : 0                 AIS Currently Fail : 0
===============================================================================

================================
ETH-CFM System Op-code Statistics
================================
Op-code      Rx Count   Tx Count
--------------------------------
ccm              4936       6002
lbr                 0          0
lbm                 0          0
ltr                 0          0
ltm                 0          0
ais                 0          0
lck                 0          0
tst                 0          0
laps                0          0
raps             1037        818
mcc                 0          0
lmr                 0          0
lmm                 0          0
1dm                 0          0
dmr                 0          0
dmm                 0          0
exr                 0          0
exm                 0          0
csf                 0          0
vsr                 0          0
vsm                 0          0
1sl                 0          0
slr                 0          0
slm                 0          0
gnm                 0          0
other               0          0
--------------------------------
Total            5973       6820
================================
```

To see an example of the messages in log "99" on a ring failure, when the unblocked port 1/1/c2/1 on PE-2 is disabled, the following messages are displayed. When logging is enabled from main to console, the same messages can be seen on the console.

```
# on PE-2:
configure {
    port 1/1/c2/1
        admin-state disable
```

```
74 2023/05/16 08:43:40.190 CEST MINOR: ETH_CFM #2001 Base
"MEP 1/12/122 highest defect is now defRemoteCCM"

73 2023/05/16 08:43:36.641 CEST MAJOR: SVCMGR #2210 Base
"Processing of an access port state change event is finished and the status of all
affected SAPs on port 1/1/c2/1 has been updated."

72 2023/05/16 08:43:36.640 CEST MINOR: ERING #2001 Base eth-ring-1
"Eth-Ring 1 path a changed fwd state to unblocked"

71 2023/05/16 08:43:36.640 CEST MINOR: ERING #2001 Base eth-ring-1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
"Eth-Ring 1 path b changed fwd state to blocked"

70 2023/05/16 08:43:36.640 CEST WARNING: SNMP #2004 Base 1/1/c2/1
"Interface 1/1/c2/1 is not operational"
```

For troubleshooting, the **tools dump eth-ring** *<ring-index>* command displays path information, the internal state of the control protocol, related statistics information, and up to the last 16 protocol events (including messages sent and received, and the expiration of timers). An associated **clear** parameter exists, which clears the event information in this output when the command is entered. The following is an example of the output on PE-2 after port 1/1/c2/1 has been enabled.

```
[/]
A:admin@PE-2# tools dump eth-ring 1

ringId 1 (Up/Up): numPaths 2 nodeId 02:0b:ff:00:00:00
 SubRing: none (interconnect ring 0, propagateTc  No), Cnt 0
  path-a, port 1/1/c1/1 (Up), tag 1.1(Up) status (Up/Up/Blk)
      cc (Dn/Up): Cnt 1/1 tm 000 00:25:53.470/000 00:36:36.970
      state: Cnt 7 B/F 000 01:10:19.470/000 01:05:31.920, flag: 0x0
  path-b, port 1/1/c2/1 (Up), tag 1.1(Up) status (Up/Up/Fwd)
      cc (Dn/Up): Cnt 2/2 tm 000 01:05:35.480/000 01:09:02.180
      state: Cnt 4 B/F 000 01:05:31.920/000 01:10:19.470, flag: 0x0
  FsmState=  IDLE, Rpl = Owner, revert = 60 s, guard = 5 ds
    Defects =
    Running Timers = PduReTx
    lastTxPdu = 0x0080 Nr(RB )
    path-a Rpl, RxId(I)= 02:09:ff:00:00:00, rx(F)= v1-0x0000 Nr, cmd= None
    path-b Normal, RxId(I)= 02:09:ff:00:00:00, rx(F)= v1-0x0000 Nr, cmd= None
  DebugInfo:  aPathSts 3, bPathSts 3, pm (set/clr) 0/0, txFlush 0
    RxRaps: ok 9 nok 0 self 36, TmrExp - wtr 2(1), grd 2, wtb 0
    Flush: cnt 8 (6/1/1) tm 000 01:10:19.470-000 01:10:19.470 Out/Ack 0/1
    RxRawRaps: aPath 85 bPath 45 vPath 0
    Now: 000 01:20:37.210 , softReset: No - noTx 0

  Seq Event  RxInfo(Path: NodeId-Bytes)
             state:TxInfo (Bytes)            Dir  pA  pB        Time
  === ===== ============================  ===== === === ================
  001  bAdd
              PROT  : 0xb020  Sf            TxF-> Blk Blk 000 00:25:49.470
  002   aUp
              PROT  : 0xb060  Sf(DNF)       Tx--> Fwd Blk 000 00:25:49.470
  003   aDn
              PROT  : 0xb000  Sf            TxF-> Blk Blk 000 00:25:53.470
  004   pdu B: 02:09:ff:00:00:00-0xb040 Sf(DNF)
              PROT  : 0xb000  Sf            Rx<-- Blk Blk 000 00:36:38.180
  005   pdu B: 02:09:ff:00:00:00-0x0000 Nr
              PROT  : 0xb000  Sf            Rx<-- Blk Blk 000 00:36:38.480
  006   pdu A: 02:0d:ff:00:00:00-0x0020 Nr
              PROT  : 0xb000  Sf            Rx<-- Blk Blk 000 00:36:38.960
  007   pdu B: 02:0d:ff:00:00:00-0x0020 Nr
              PROT  : 0xb000  Sf            Rx<-- Blk Blk 000 00:36:39.160
  008   bUp
              PROT  : 0xb040  Sf(DNF)       Tx--> Blk Fwd 000 00:36:39.170
  009   aUp
              PEND-G: 0x0000  Nr            Tx--> Blk Fwd 000 00:36:39.470
  010   pdu B: 02:09:ff:00:00:00-0xe000 Ev
              PEND  : 0x0000  Nr            Frx<- Blk Fwd 000 00:36:40.430
  011   pdu A: 02:0d:ff:00:00:00-0x0020 Nr
              PEND  : 0x0000  Nr            Rx<-- Blk Fwd 000 00:36:40.440
  012   pdu
              PEND  :                       ----- Fwd Fwd 000 00:36:40.440
  013   pdu B: 02:0d:ff:00:00:00-0x0020 Nr
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
             PEND  :                       Rx<-- Fwd Fwd 000 00:36:40.440
  014  xWtr
             IDLE  : 0x0080  Nr(RB )       TxF-> Blk Fwd 000 00:37:40.470
  015   bDn
             PROT  : 0xb020  Sf            TxF-> Fwd Blk 000 01:05:31.920
  016   pdu A: 02:09:ff:00:00:00-0xb000 Sf
             PROT  : 0xb020  Sf            RxF<- Fwd Blk 000 01:05:35.190
  017   pdu B: 02:09:ff:00:00:00-0x0000 Nr
             PROT  : 0xb020  Sf            Rx<-- Fwd Blk 000 01:09:03.480
  018   pdu A: 02:09:ff:00:00:00-0x0000 Nr
             PROT  : 0xb020  Sf            Rx<-- Fwd Blk 000 01:09:03.490
  019   bUp
             PEND-G: 0x0020  Nr            Tx--> Fwd Blk 000 01:09:04.170
  000  xWtr
             IDLE  : 0x0080  Nr(RB )       TxF-> Blk Fwd 000 01:10:19.470
```

Configuration of a subring to a major ring with a non-virtual link

The differences from the preceding virtual link configuration with a non-virtual link for the subring are:

- The subring configuration on the interconnection nodes, PE-1 and PE-3, is modified to indicate that the subring is not using a virtual link, otherwise it remains the same.
- The subring configuration on the subring node PE-4 is also modified to indicate that this is part of a subring that is not using a virtual link. This is mandatory on all non-interconnection nodes on the subring in order to ensure the propagation of R-APS messages around the subring.
- The virtual link services and SAPs must be removed from PE-1, PE-2, and PE3, that is:
  - On PE-1 and PE-3, the SAPs in VPLS 2 around the major ring (configured with the parameter **eth-ring 1**) are removed.
  - The service "control-VPLS-2" is removed completely from PE-2.

The new configuration of subring 2 on PE-1 is as follows, the configuration on PE-3 is similar.

```
# on PE-1:
configure {
    eth-ring 2 {
        admin-state enable
        description "Ethernet subring 2 on major ring 1"
        revert-time 60
        sub-ring {
            type non-virtual-link
            interconnect {
                ring-id 1
                propagate-topology-change true
            }
        }
        path "a" {
            admin-state enable
            description "Ethernet subring 2_path a"
            port-and-raps-tag 1/1/c2/1:2.1
            eth-cfm {
                mep md-admin-name "domain-1" ma-admin-name "association-14" mep-id 141 {
                    admin-state enable
                    ccm true
                    control-mep true
                }
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The configuration of subring 2 on non-interconnection node PE-4 must include the **type non-virtual-link** parameter, as follows:

```
# on PE-4:
configure {
    eth-ring 2 {
        admin-state enable
        description "Ethernet subring 2"
        revert-time 60
        rpl-node owner
        sub-ring {
            type non-virtual-link
        }
        path "a" {
            admin-state enable
            description "Ethernet subring 2_path a"
            port-and-raps-tag 1/1/c1/1:2.1
            eth-cfm {
                mep md-admin-name "domain-1" ma-admin-name "association-14" mep-id 144 {
                    admin-state enable
                    ccm true
                    control-mep true
                }
            }
        }
        path "b" {
            admin-state enable
            description "Ethernet subring 2_path b"
            port-and-raps-tag 1/1/c2/1:2.1
            rpl-end true
            eth-cfm {
                mep md-admin-name "domain-1" ma-admin-name "association-34" mep-id 344 {
                    admin-state enable
                    ccm true
                    control-mep true
                }
            }
        }
```

The SAP usage on PE-1 is as follows with only the control and data SAPs to PE-4 now using subring 2.

```
[/]
A:admin@PE-1# show service sap-using eth-ring

===============================================================================
Service Access Points (Ethernet Ring)
===============================================================================
SapId                    SvcId       Eth-Ring Path Admin Oper  Blocked Control/
                                                   State State         Data
-------------------------------------------------------------------------------
1/1/c1/1:1.1             1           1        a    Up    Up    No      Ctrl
1/1/c3/1:1.1             1           1        b    Up    Up    No      Ctrl
1/1/c2/1:2.1             2           2        a    Up    Up    No      Ctrl
1/1/c1/1:1.11            11          1        a    Up    Up    No      Data
1/1/c2/1:1.11            11          2        a    Up    Up    No      Data
1/1/c3/1:1.11            11          1        b    Up    Up    No      Data
-------------------------------------------------------------------------------
Number of SAPs : 6
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The information relating to subring 2 is as follows and it can be seen that this is now not using a virtual link, but subring 2 is still connected to major ring 1 and propagation is still enabled from the subring to the major ring. The single ring path "a" is unblocked because the RPL is configured between PE-3 and PE-4.

```
[/]
A:admin@PE-1# show eth-ring 2

===============================================================================
Ethernet Ring 2 Information
===============================================================================
Description        : Ethernet subring 2 on major ring 1
Admin State        : Up                 Oper State       : Up
Node ID            : 02:09:ff:00:00:00
Guard Time         :    5 deciseconds   RPL Node         : rplNone
Max Revert Time    :   60 seconds       Time to Revert   : N/A
CCM Hold Down Time :    0 centiseconds  CCM Hold Up Time :   20 deciseconds
Compatible Version : 2
APS Tx PDU         : N/A
Defect Status      :

Sub-Ring Type      : nonVirtualLink     Interconnect-ID  : 1
Topology Change    : Propagate

-------------------------------------------------------------------------------
Ethernet Ring Path Summary
-------------------------------------------------------------------------------
Path Port             Raps-Tag     Admin/Oper      Type            Fwd State
-------------------------------------------------------------------------------
  a  1/1/c2/1         2.1             Up/Up        normal          unblocked
  b  -                -                 -/-         -               -
===============================================================================
```

Configuration of a subring to a VPLS service

Subrings can be connected to VPLS services, in which case a virtual link is not used and is not configurable. While similar to the ring interconnect, there are a few differences.

Flush propagation is from the subring to the VPLS, in the same way as it was for the subring to the major ring. The same configuration parameter is used to propagate topology changes. In this case, LDP "flush-all-from-me" messages are sent into the LDP portion of the network to account for ring changes without the need to configure anything in the VPLS service.

As with other rings, until an Ethernet ring instance is attached to the VPLS service, the ring operational status is down and the forwarding status of each port is blocked. This prevents operators from creating a loop by misconfiguration.

The topology for this case is shown in Figure 10: Subring to VPLS topology. The configuration is very similar to the subring with a non-virtual link described earlier, but ring 1 is replaced by a VPLS service using LDP-signaled mesh SDPs between PE-1, PE-2, and PE-3 to create a fully meshed VPLS service. Both spoke and mesh SDPs using LDP can be used for the VPLS; however, only mesh SDPs have been used in this example.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 10: Subring to VPLS topology

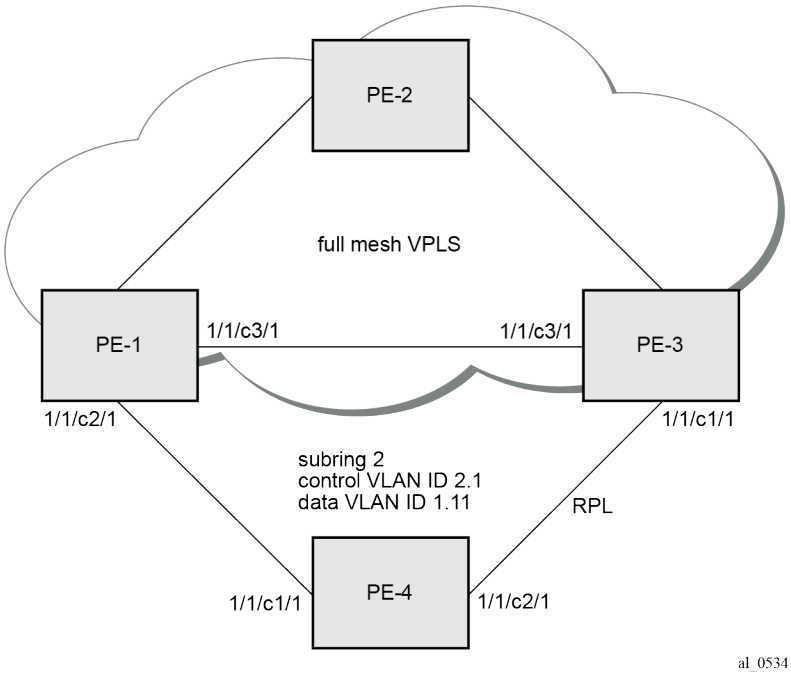

The differences for the VPLS service connection to the configuration when the subring is connected to a major ring without a virtual link are:

- The subring configuration on the interconnection nodes, PE-1 and PE-3, is modified to indicate that the subring is connected to a VPLS service.
- The subring configuration on the non-interconnection node PE-4 indicates that this is part of a subring that is not using a virtual link (same configuration as in the scenario when a subring is connected to a major ring without a virtual link). This is mandatory on all non-interconnection nodes on the subring in order to ensure the propagation of R-APS messages around the subring.
- The control VPLS "control-VPLS-1" and SAPs relating to the major ring 1 on PE-1, PE-2, and PE-3 are removed. These are replaced by routed IP interfaces configured with a routing protocol and LDP in order to signal the required MPLS labels, together with the necessary SDPs to provide interconnection at a service level.
- The data service "VPLS-11" is configured with mesh SDPs between PE-1, PE-2, and PE-3.

The configuration on PE-1 of the subring 2 is as follows with the interconnect indicating a VPLS service. The configuration on PE-3 is similar.

```
# on PE-1:
configure {
    eth-ring 2 {
        admin-state enable
        description "Ethernet subring 2 on VPLS"
        revert-time 60
        sub-ring {
            type non-virtual-link
            interconnect {
                vpls
                propagate-topology-change true
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
        path "a" {
            admin-state enable
            description "Ethernet subring 2_path a"
            port-and-raps-tag 1/1/c2/1:2.1
            eth-cfm {
                mep md-admin-name "domain-1" ma-admin-name "association-14" mep-id 141 {
                    admin-state enable
                    ccm true
                    control-mep true
                }
            }
        }
```

The following configuration of subring 2 on non-interconnection node PE-4 includes the **type non-virtual-link** parameter:

```
# on PE-4:
configure {
    eth-ring 2 {
        admin-state enable
        description "Ethernet subring 2"
        revert-time 60
        rpl-node owner
        sub-ring {
            type non-virtual-link
        }
        path "a" {
            admin-state enable
            description "Ethernet subring 2_path a"
            port-and-raps-tag 1/1/c1/1:2.1
            eth-cfm {
                mep md-admin-name "domain-1" ma-admin-name "association-14" mep-id 144 {
                    admin-state enable
                    ccm true
                    control-mep true
                }
            }
        }
        path "b" {
            admin-state enable
            description "Ethernet subring 2_path b"
            port-and-raps-tag 1/1/c2/1:2.1
            rpl-end true
            eth-cfm {
                mep md-admin-name "domain-1" ma-admin-name "association-34" mep-id 344 {
                    admin-state enable
                    ccm true
                    control-mep true
                }
            }
        }
```

The data service on PE-1 is as follows. The configuration on PE-3 is similar.

```
# on PE-1:
configure {
    service {
        vpls "VPLS-11" {
            admin-state enable
            description "data VPLS"
            service-id 11
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            customer "1"
            mesh-sdp 12:11 {
            }
            mesh-sdp 13:11 {
            }
            sap 1/1/c2/1:1.11 {
                eth-ring 2
            }
            sap 1/1/c4/1:11 {
                description "sample customer service SAP"
            }
```

The state of the subring is as follows and shows the subring is not using a virtual link, is connected to a VPLS service, and has propagation of topology change events enabled. As earlier, the single ring path "a" is unblocked because the RPL is configured between PE-3 and PE-4.

```
[/]
A:admin@PE-1# show eth-ring 2

===============================================================================
Ethernet Ring 2 Information
===============================================================================
Description        : Ethernet subring 2 on VPLS
Admin State        : Up                 Oper State       : Up
Node ID            : 02:09:ff:00:00:00
Guard Time         :    5 deciseconds   RPL Node         : rplNone
Max Revert Time    :   60 seconds       Time to Revert   : N/A
CCM Hold Down Time :    0 centiseconds  CCM Hold Up Time :   20 deciseconds
Compatible Version : 2
APS Tx PDU         : N/A
Defect Status      :

Sub-Ring Type      : nonVirtualLink     Interconnect-ID  : VPLS
Topology Change    : Propagate

-------------------------------------------------------------------------------
Ethernet Ring Path Summary
-------------------------------------------------------------------------------
Path Port             Raps-Tag     Admin/Oper      Type            Fwd State
-------------------------------------------------------------------------------
  a  1/1/c2/1         2.1             Up/Up        normal          unblocked
  b  -                -                 -/-         -               -
===============================================================================
```

In this case, if a topology change event occurs in the subring, an LDP "flush-all-from-me" message is sent by PE-1 and PE-3 to their LDP peers. This can be seen by enabling the following debugging for PE-1:

```
# on PE-1:
debug {
    router "Base" {
        ldp {
            peer 192.0.2.2 {
                packet {
                    init {
                    }
                }
            }
            peer 192.0.2.3 {
                packet {
                    init {
                    }
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
```

The topology change is forced by disabling port 1/1/c2/1 on PE-1.

```
# on PE-1:
configure {
    port 1/1/c2/1
        admin-state disable
```

The log shows the following messages on the console (combination of log 1 for debug-trace and log 2 for main), where packets 1 and 2 are the LDP flush messages.

```
88 2023/05/16 09:39:27.293 CEST WARNING: SNMP #2004 Base 1/1/c2/1
"Interface 1/1/c2/1 is not operational"

89 2023/05/16 09:39:27.293 CEST MINOR: ERING #2001 Base eth-ring-2
"Eth-Ring 2 path a changed fwd state to blocked"

1 2023/05/16 09:39:27.294 CEST MINOR: DEBUG #2001 Base LDP
"LDP: LDP
Send Address Withdraw packet (msgId 191) to 192.0.2.2:0
 MAC Flush (All MACs learned from me)
Service FEC PWE3: ENET(5)/11 Group ID = 0 cBit = 0
"

2 2023/05/16 09:39:27.294 CEST MINOR: DEBUG #2001 Base LDP
"LDP: LDP
Send Address Withdraw packet (msgId 190) to 192.0.2.3:0
 MAC Flush (All MACs learned from me)
Service FEC PWE3: ENET(5)/11 Group ID = 0 cBit = 0
"

90 2023/05/16 09:39:27.299 CEST MAJOR: SVCMGR #2210 Base
"Processing of an access port state change event is finished and the status of a
ll affected SAPs on port 1/1/c2/1 has been updated."

91 2023/05/16 09:39:30.909 CEST MINOR: ETH_CFM #2001 Base
"MEP 1/14/141 highest defect is now defRemoteCCM"
```

Operational procedures

Operators may wish to configure rings with or without control over reversion. Reversion can be controlled by timers or the ring can be run without reversion allowing the operator to choose when the ring reverts. To change a ring topology, the **manual** or **force** switch command may be used to block a specified ring path. A ring will still address failures when run without reversion but will not automatically revert to the RPL when resources are restored. A **clear** command can be used to clear the manual or force state of a ring.

The following **tools** commands are available to control the state of paths on a ring.

```
tools perform eth-ring clear <ring-index>
tools perform eth-ring force <ring-index> path {a|b}
tools perform eth-ring manual <ring-index> path {a|b}
```

In the following output, both ports of Ethernet ring 1 are unblocked.

```
[/]
A:admin@PE-1# show eth-ring 1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Ethernet Ring 1 Information
===============================================================================
Description        : Ethernet ring 1_major ring
Admin State        : Up                 Oper State       : Up
Node ID            : 02:09:ff:00:00:00
Guard Time         :    5 deciseconds   RPL Node         : rplNone
Max Revert Time    :   60 seconds       Time to Revert   : N/A
CCM Hold Down Time :    0 centiseconds  CCM Hold Up Time :   20 deciseconds
Compatible Version : 2
APS Tx PDU         : N/A
Defect Status      :

Sub-Ring Type      : none

-------------------------------------------------------------------------------
Ethernet Ring Path Summary
-------------------------------------------------------------------------------
Path Port             Raps-Tag     Admin/Oper      Type            Fwd State
-------------------------------------------------------------------------------
  a  1/1/c1/1         1.1             Up/Up        normal          unblocked
  b  1/1/c3/1         1.1             Up/Up        normal          unblocked
===============================================================================
```

The following command blocks path "b" of Ethernet ring 1 manually:

```
*A:PE-1# tools perform eth-ring manual 1 path b
```

In the following output, path "b" of Ethernet ring 1 is blocked:

```
[/]
A:admin@PE-1# show eth-ring 1

===============================================================================
Ethernet Ring 1 Information
===============================================================================
Description        : Ethernet ring 1_major ring
Admin State        : Up                 Oper State       : Up
Node ID            : 02:09:ff:00:00:00
Guard Time         :    5 deciseconds   RPL Node         : rplNone
Max Revert Time    :   60 seconds       Time to Revert   : N/A
CCM Hold Down Time :    0 centiseconds  CCM Hold Up Time :   20 deciseconds
Compatible Version : 2
APS Tx PDU         : Request State: 0x7
                     Sub-Code     : 0x0
                     Status       : 0x20  ( BPR )
                     Node ID      : 02:09:ff:00:00:00
Defect Status      :

Sub-Ring Type      : none

-------------------------------------------------------------------------------
Ethernet Ring Path Summary
-------------------------------------------------------------------------------
Path Port             Raps-Tag     Admin/Oper      Type            Fwd State
-------------------------------------------------------------------------------
  a  1/1/c1/1         1.1             Up/Up        normal          unblocked
  b  1/1/c3/1         1.1             Up/Up        normal          blocked
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The following command on PE-1 clears Ethernet ring 1:

```
[/]
A:admin@PE-1# tools perform eth-ring clear 1
```

After Ethernet ring 1 is cleared on PE-1, both paths are unblocked again.

```
[/]
A:admin@PE-1# show eth-ring 1

===============================================================================
Ethernet Ring 1 Information
===============================================================================
Description        : Ethernet ring 1_major ring
Admin State        : Up                 Oper State       : Up
Node ID            : 02:09:ff:00:00:00
Guard Time         :    5 deciseconds   RPL Node         : rplNone
Max Revert Time    :   60 seconds       Time to Revert   : N/A
CCM Hold Down Time :    0 centiseconds  CCM Hold Up Time :   20 deciseconds
Compatible Version : 2
APS Tx PDU         : N/A
Defect Status      :

Sub-Ring Type      : none

-------------------------------------------------------------------------------
Ethernet Ring Path Summary
-------------------------------------------------------------------------------
Path Port             Raps-Tag     Admin/Oper      Type            Fwd State
-------------------------------------------------------------------------------
  a  1/1/c1/1         1.1             Up/Up        normal          unblocked
  b  1/1/c3/1         1.1             Up/Up        normal          unblocked
===============================================================================
```

Both the **manual** and **force** command block the path specified, however, the **manual** command fails if there is an existing forced switch or signal fail event in the ring, as seen in the following output. The **force** command will block the port regardless of any existing ring state and there can be multiple force states simultaneously on a ring on different nodes.

```
[/]
A:admin@PE-1# tools perform eth-ring force 1 path b

[/]
A:admin@PE-1# tools perform eth-ring manual 1 path b
INFO: ERMGR #1001: Not permitted - The switch command is not compatible to the current state
 (FS), effective priority (FS) or rpl-node type (None)
```

Conclusion

Ethernet ring APS provides an optimal solution for designing native Ethernet services with ring topology. With subrings, both multiple rings and access rings increase the versatility of G.8032. G.8032 has been expanded to more of the SR platforms by allowing R-APS with slower MEPs (including CCMs intervals of 1 second). This protocol provides simple configuration, operation, and guaranteed fast protection time. The implementation also has a flexible encapsulation that allows dot1Q, QinQ, or PBB for the ring traffic. It can be utilized on various services such as mobile backhaul, business VPN access, aggregation, and core.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

G.8032 Ethernet Ring Protection Single Ring Topology

This chapter provides information about G.8032 Ethernet ring protection single ring topology.

Topics in this chapter include:

Applicability

The chapter was initially written for SR OS Release 8.0.R7, but the MD-CLI in the current edition corresponds to SR OS Release 23.3.R2. This chapter describes ring protection for a single ring topology. Protection for multiple ring topologies is covered in G.8032 Ethernet Ring Protection Multiple Ring Topology.

Overview

G.8032 Ethernet ring protection is supported for data service SAPs within a regular VPLS service, a provider backbone bridging (PBB) VPLS (I/B-component), or a routed VPLS (R-VPLS). G.8032 is one of the fastest protection schemes for Ethernet networks.

ITU-T G.8032v2 specifies protection switching mechanisms and a protocol for Ethernet layer network (ETH) Ethernet rings. Ethernet rings can provide wide-area multi-point connectivity more economically due to their reduced number of links. The mechanisms and protocol defined in ITU-T G.8032v2 achieve highly reliable and stable protection and never form loops, which would negatively affect network operation and service availability. Each ring node is connected to adjacent nodes participating in the same ring using two independent paths, which use ring links that are configured on ports or link aggregation groups (LAGs). A ring link is bounded by two adjacent nodes and a port for a ring link is called a ring port. The minimum number of nodes on a ring is two.

The fundamentals of this ring protection switching architecture are:

- the principle of loop avoidance and
- the utilization of learning, forwarding, and address table mechanisms defined in the ITU-T G.8032v2 Ethernet flow forwarding function (ETH_FF) (control plane).

Loop avoidance in the ring is achieved by guaranteeing that, at any time, traffic may flow on all but one of the ring links. This particular link is called the ring protection link (RPL) and under normal conditions this link is blocked, so it is not used for traffic. One designated node, the RPL owner, is responsible to block traffic over the one designated RPL. Under a ring failure condition, the RPL owner is responsible for unblocking the RPL, allowing the RPL to be used for traffic. The protocol ensures that even without an RPL owner defined, one link will be blocked and it operates as a *break before make protocol*, specifically the protocol guarantees that no link is restored until a different link in the ring is blocked. The other side of the RPL is configured as an RPL neighbor. An RPL neighbor blocks traffic on the link.

The event of a ring link or ring node failure results in protection switching of the traffic. This is achieved under the control of the ETH_FF functions on all ring nodes. A ring automatic protection switching (R-APS)

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

protocol is used to coordinate the protection actions over the ring. The protection switching mechanisms and protocol supports a multi-ring/ladder network that consists of connected Ethernet rings, however, that is not covered in this chapter.

Ring protection mechanism

The ring protection protocol is based on the following building blocks:

- ring status change on failure
  - idle → link failure → protection → recovery → idle
- ring control state changes
  - idle → protection → manual switch → forced switch → pending
- re-use existing ETH OAM
  - monitoring: Ethernet continuity check messages
  - failure notification: Y.1731 signal failure
- forwarding database MAC flush on ring status change
- ring protection link (RPL) defines blocked link in idle status

Figure 11: G.8032 operation and topologies shows a ring of six nodes, with the RPL owner on the top right. One link of the RPL owner is designated to be the RPL and will be blocked in order to prevent a loop. Schematics of the physical and logical topologies are also shown.

When an RPL owner and RPL end are configured, the associated link will be the RPL when the ring is fully operational and so be blocked by the RPL owner. If a different ring link fails, then the RPL will be unblocked by the RPL owner. When the failed link recovers, it will initially be blocked by one of its adjacent nodes. The adjacent node sends an R-APS message across the ring to indicate the error is cleared and after a configurable time, if reversion is enabled, the RPL will revert to being blocked with all other links unblocked. This ensures that the ring topology is predictable when fully operational.

If a specific RPL owner is not configured, then the last link to become active will be blocked and the ring will remain in this state until another link fails. However, this operation makes the selection of the blocked link non-deterministic.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 11: G.8032 operation and topologies

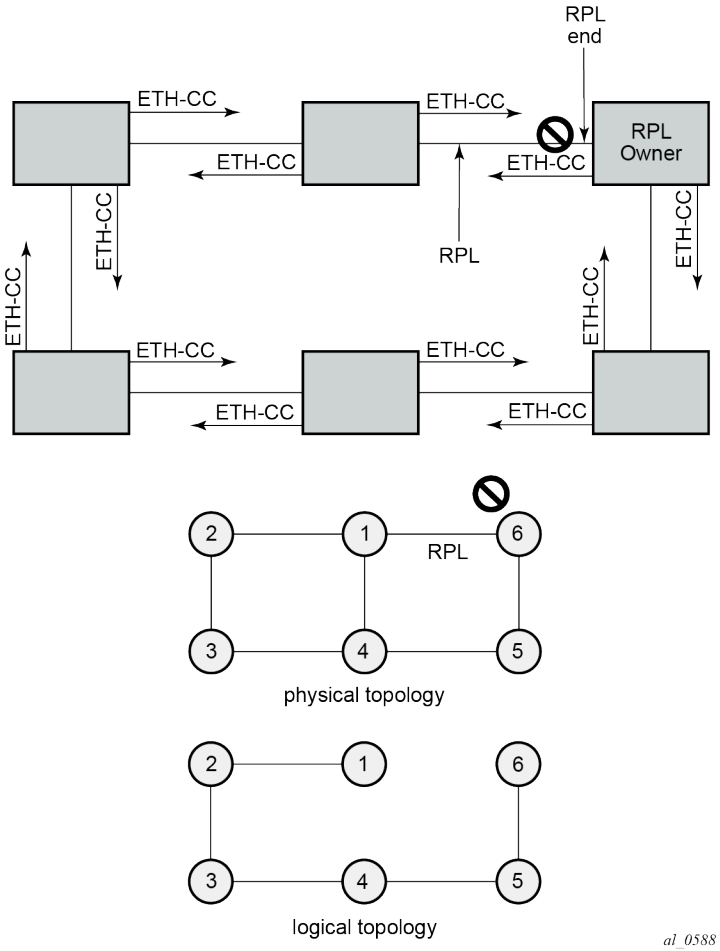

The protection protocol uses a specific control VLAN, with the associated data VLANs taking their forwarding state from the control VLAN.

Configuration

The example topology is shown in Figure 12: Example topology.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 12: Example topology

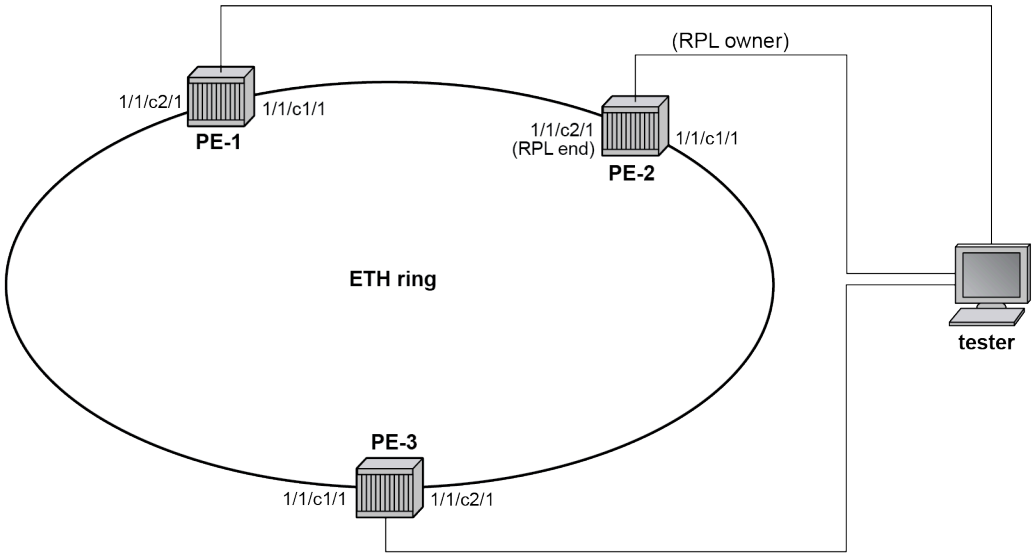

** control channel: VPLS 1, tag 1
** data channel: VPLS 100, tag 100

al_0589

The Ethernet ring configuration commands are as follows:

```
configure {
    eth-ring <ring-index [1..128]> {
        ccm-hold-time {
            down <number>   # Hold timer for down event dampening
            up <number>     # Hold timer for recovery reporting
        }
        compatible-version <number>            # [1..2] - Default: 2
        description <string>
        guard-time <number>       # [1..20] in deciseconds - Default: 5
        node-id <mac-address>     # MAC address of the RPL <xx:xx:xx:xx:xx:xx>
        path <string>   # path ID: string of 1 character
            admin-state <keyword>   # default: disable
            description <string>
            port-and-raps-tag <port-and-encap>     # Port ID and Ring APS tag ID
            eth-cfm {
                mep md-admin-name <reference> ma-admin-name <reference> mep-id <number> {
                    admin-state <keyword>   # default: disable
                    ccm <boolean>           # default: false
                    control-mep <boolean>   # default: false
                }
            }
            rpl-end <boolean>               # default: false
        }
        revert-time <number>      # <0,60..720> in seconds - Default: 300
        rpl-node <keyword>        # owner | neighbor
        sub-ring        # beyond the scope
```

Parameters:

- *ring-index* — This is the number by which the ring is referenced, values: 1 to 128.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- **ccm-hold-time {[down** <*hold timer for down event dampening*>**] [up** <*hold timer for recovery reporting*>**]}**
  - **down** — This command specifies the timer that controls the delay between detecting that ring path is down and reporting it to the G.8032 protection module. If a non-zero value is configured, the system will wait for the time specified in the value parameter before reporting it to the G.8032 protection module. This parameter applies only to the ring path continuity check message (CCM); it does not apply to the ring port link state. To dampen ring port link state transitions, use the **hold-time** parameter from the physical member port. This is useful if the underlying path between two nodes is going across an optical system which implements its own protection.
  - **up** — This command specifies the timer which controls the delay between detecting that the ring path is up and reporting it to the G.8032 protection module. If a non-zero value is configured, the system will wait for the time specified in the value parameter before reporting it to the G.8032 protection module. This parameter applies only to ring path CCM; it does not apply to the member port link state. To dampen member port link state transitions, use the **hold-time** parameter from the physical member port.
  - hold timer values:

```
*[ex:/configure eth-ring 1 ccm-hold-time]
A:admin@PE-1# down ?

 down <number>
 <number> - <1..5000> - centiseconds

    Hold timer for down event dampening

*[ex:/configure eth-ring 1 ccm-hold-time]
A:admin@PE-1# up ?

 up <number>
 <number> - <0..5000> - deciseconds
 Default  - 20

    Hold timer for recovery reporting
```

- **compatible version** — This command configures the Ethernet ring compatibility version for the G.8032 state machine and messages. The default is version 2 (ITU G.8032v2) and all SR OS nodes use version 2. If there is a need to interwork with third party devices that only support version 1, this can be set to version 1 allowing the reception of version 1 PDUs. Version 2 is encoded as 1 in the R-APS messages. Compatibility allows the reception of version 1 (encoded as 0) R-APS PDUs but, as per the G.8032 specification, higher versions are ignored on reception. For SR OS nodes, messages are always originated with version 2. Therefore, if a third party switch supported version 3 (encoded as 2) or higher, interworking is also supported provided the other switch is compatible with version 2 (encoded as 1).
- **description <string>** — This configures a text string, up to 80 characters, which can be used to describe the use of the Ethernet ring.
- **guard-time <number>** — The forwarding method, in which R-APS messages are copied and forwarded at every Ethernet ring node, can result in a message corresponding to an old request, that is no longer relevant, being received by Ethernet ring nodes. Reception of an old R-APS message may result in erroneous ring state interpretation by some Ethernet ring nodes. The guard timer is used to prevent Ethernet ring nodes from acting upon outdated R-APS messages and prevents the possibility of forming a closed loop. Messages are not forwarded when the guard-timer is running.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The guard time is configured in 10ths of seconds and the default guard time is 0.5 s:

```
*[ex:/configure eth-ring 1]
A:admin@PE-1# guard-time ?

 guard-time <number>
 <number> - <1..20> - deciseconds
 Default  - 5

    Ethernet ring guard time
```

- **node-id <mac-address>** — The node identifier can be explicitly configured. In typical configurations, the node ID is not configured; by default, the chassis MAC address is used as node ID.
- **path [path-index] <1-character string> [port-and-raps-tag]** — The **path** parameter defines the paths around the ring, of which there are two in different directions on the ring: an "a" path and a "b" path. In addition, the path command configures the encapsulation used for the R-APS messages on the ring. These can be either single or double tagged.
  - **description <string>** — The description is a text string with up to 80 characters, that can be used to describe the use of the path.
  - **eth-cfm** — Configures the associated Ethernet connectivity fault management (CFM) parameters.
    - **mep md-admin-name <reference> ma-admin-name <reference> <mep-id>** — The maintenance endpoint (MEP) defined under the path is used for the G.8032 protocol messages, which are based on IEEE 802.1ag/Y.1731 CFM frames.
  - **rpl-end** — When configured, this path is expected to be one end of the RPL. This parameter must be configured in conjunction with the **rpl-node**.
  - **admin-state <keyword>** — This command enables or disables the path.
- **revert-time <number>** — This command configures the revert time for an Ethernet ring. The revert time is the time that the RPL will wait before returning to the blocked state.

  Values:

```
*[ex:/configure eth-ring 1]
A:admin@PE-1# revert-time ?

 revert-time <number>
 <number> - <0,60..720> - seconds
 Default  - 300
```

- **rpl-node <keyword{owner|neighbor}>** — A node can be designated as either the **owner** of the RPL, in which case this node is responsible for the RPL, or the **neighbor**, in which case this node is expected to be the neighbor to the RPL owner across the RPL. The neighbor is optional and is included to be compliant with the specification. This parameter must be configured in conjunction with the **rpl-end** parameter.
- **admin-state <keyword>** — This command enables or disables the Ethernet ring.
- **sub-ring** — The **sub-ring** command is beyond the scope of this chapter because it is only required for multiple ring topologies.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Logging

Create following log-id on PE-2 to see major events logged to the console on PE-2. This is an optional step; alternatively, log 99 can be consulted.

```
# on PE-2:
configure {
    log {
        log-id "log1" {
            source {
                main true
            }
            destination {
                console
            }
        }
```

Configure encapsulation for ring ports

To configure R-APS, there should be at least two VPLS services for one Ethernet ring instance, one VPLS for the control channel and the other VPLSs for data channels. The control channel is used for R-APS signaling while the data channel is for user data traffic. The state of the data channels is inherited from the state of the control channel.

- An Ethernet ring needs R-APS tags to send and receive G.8032 signaling messages. To configure a control channel, an access SAP configuration is required on each path a port and path b port. The SAP configuration follows that of the port and must be either dot1Q or QinQ, so the control and data packets are either single tagged or double tagged. It is also possible to have the control VPLS using single tagged frames with the data VPLSs using double tagged frames; this requires the system to be configured with the **extended-default-qinq-sap-lookup** parameter (**configure service system extended-default-qinq-sap-lookup true**), with the ring path R-APS tags and control VPLS SAPs configured as qtag.0, and the data VPLS SAPs configured as qtag1.qtag2.

  In the example in this chapter, single tags are used so the ports on the ring nodes are configured as follows:

```
# on PE-1, PE-2, PE-3:
configure {
    port 1/1/c1/1 {
        admin-state enable
        ethernet {
            mode access
            encap-type dot1q
        }
    }
    port 1/1/c2/1 {
        admin-state enable
        ethernet {
            mode access
            encap-type dot1q
        }
    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Configure Ethernet CFM

Ethernet ring requires Ethernet CFM domains, associations, and MEPs being configured. The domain format must be none and association name must be ITU-T carrier code-based (ICC-based - Y.1731). The minimum CCM interval for the SR OS nodes is 10ms. The Ethernet ring MEP requires a CCM interval, such as 10ms, 100ms, or 1s, to be configured.

The MEPs used for R-APS control normally have CCM configured on the control channel path MEPs for failure detection. Alternatively, detecting a failure of the ring may be achieved by running Ethernet in the first mile (EFM) at the port level if CCM is not possible at 10ms, 100ms, or 1s. Loss-of-signal, in conjunction with other OAM, is applicable only when the nodes are directly connected.

To omit the failure detecting CCMs, remove the **ccm true** from under the path MEPs and remove the **remote-mep** from under the **eth-cfm>domain>association** on all nodes.

Figure 13: Ethernet CFM configuration shows the Ethernet CFM configuration used here.

Figure 13: Ethernet CFM configuration

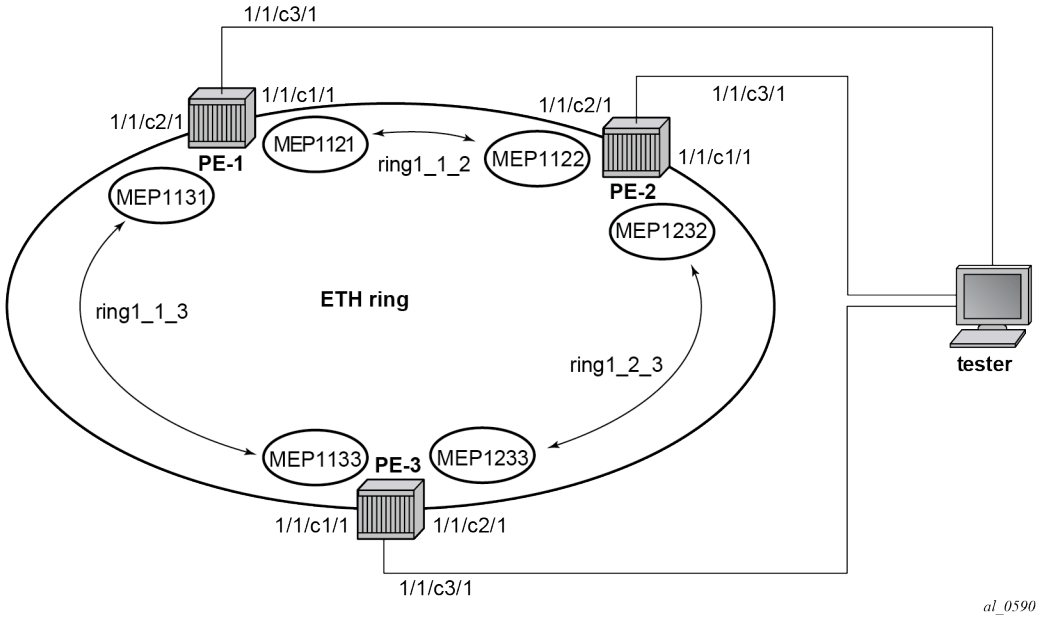

The Ethernet CFM configuration of the nodes is as follows.

```
# on PE-1:
configure {
    eth-cfm {
        domain "domain-1" {
            level 3
            format none
            md-index 1
            association "association-1" {
                icc-based "ring1_1_2"
                ma-index 1
                ccm-interval 1s
                remote-mep 1122 {
                }
            }
            association "association-2" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                icc-based "ring1_1_3"
                ma-index 2
                ccm-interval 1s
                remote-mep 1133 {
                }
            }
        }
```

```
# on PE-2:
configure {
    eth-cfm {
        domain "domain-1" {
            level 3
            format none
            md-index 1
            association "association-1" {
                icc-based "ring1_2_3"
                ma-index 1
                ccm-interval 1s
                remote-mep 1233 {
                }
            }
            association "association-2" {
                icc-based "ring1_1_2"
                ma-index 2
                ccm-interval 1s
                remote-mep 1121 {
                }
            }
        }
```

```
# on PE-3:
configure {
    eth-cfm {
        domain "domain-1" {
            level 3
            format none
            md-index 1
            association "association-1" {
                icc-based "ring1_1_3"
                ma-index 1
                ccm-interval 1s
                remote-mep 1131 {
                }
            }
            association "association-2" {
                icc-based "ring1_2_3"
                ma-index 2
                ccm-interval 1s
                remote-mep 1232 {
                }
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Configure Ethernet ring

Two paths need to be configured to form a ring: path a and path b. In this example, VLAN tag 1 is used as control channel for R-APS signaling in the ring.

```
# on PE-1:
configure {
    eth-ring 1 {
        admin-state enable
        path "a" {
            admin-state enable
            port-and-raps-tag 1/1/c1/1:1
            eth-cfm {
                mep md-admin-name "domain-1" ma-admin-name "association-1" mep-id 1121 {
                    admin-state enable
                    ccm true
                    control-mep true
                }
            }
        }
        path "b" {
            admin-state enable
            port-and-raps-tag 1/1/c2/1:1
            eth-cfm {
                mep md-admin-name "domain-1" ma-admin-name "association-2" mep-id 1131 {
                    admin-state enable
                    ccm true
                    control-mep true
                }
            }
        }
```

It is mandatory to configure a MEP in the path context, otherwise the following error is displayed:

```
*[ex:/configure eth-ring 1]
A:admin@PE-1# commit
MINOR: MGMT_CORE #4001: configure eth-ring 1 path "a" admin-state - Cannot enable path
without eth-cfm mep configured - configure eth-ring 1 path "a" admin-state
MINOR: MGMT_CORE #4001: configure eth-ring 1 path "b" admin-state - Cannot enable path
without eth-cfm mep configured - configure eth-ring 1 path "b" admin-state
```

While MEPs are mandatory, enabling CCM on the MEP in the path context as a failure detection mechanism is optional.

In order to define the RPL, node PE-2 is configured as the RPL owner and path b as the RPL end. The link between nodes PE-1 and PE-2 will be the RPL with node PE-2 blocking that link when the ring is fully operational.

```
# on PE-2:
configure {
    eth-ring 1 {
        admin-state enable
        revert-time 60
        rpl-node owner
        path "a" {
            admin-state enable
            port-and-raps-tag 1/1/c1/1:1
            eth-cfm {
                mep md-admin-name "domain-1" ma-admin-name "association-1" mep-id 1232 {
                    admin-state enable
                    ccm true
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    control-mep true
                }
            }
        }
        path "b" {
            admin-state enable
            port-and-raps-tag 1/1/c2/1:1
            rpl-end true
            eth-cfm {
                mep md-admin-name "domain-1" ma-admin-name "association-2" mep-id 1122 {
                    admin-state enable
                    ccm true
                    control-mep true
                }
            }
        }
```

It is not allowed to configure a path as an RPL end without having configured the node on this ring to be either the RPL **owner** or **neighbor** otherwise the following error message is reported.

```
*[ex:/configure eth-ring 1]
A:admin@PE-2# commit
MINOR: MGMT_CORE #4001: configure eth-ring 1 path "b" rpl-end - rpl-node must be set -
configure eth-ring 1 rpl-node
```

```
# on PE-3:
configure {
    eth-ring 1 {
        admin-state enable
        path "a" {
            admin-state enable
            port-and-raps-tag 1/1/c1/1:1
            eth-cfm {
                mep md-admin-name "domain-1" ma-admin-name "association-1" mep-id 1133 {
                    admin-state enable
                    ccm true
                    control-mep true
                }
            }
        }
        path "b" {
            admin-state enable
            port-and-raps-tag 1/1/c2/1:1
            eth-cfm {
                mep md-admin-name "domain-1" ma-admin-name "association-2" mep-id 1233 {
                    admin-state enable
                    ccm true
                    control-mep true
                }
            }
        }
```

Until the Ethernet ring instance is attached to the service (VPLS in this case), the ring operational status is down and the forwarding status of each port is blocked. This prevents operators from creating a loop by misconfiguration. This state can be seen on ring node PE-1 as follows:

```
[/]
A:admin@PE-1# show eth-ring 1

===============================================================================
Ethernet Ring 1 Information
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Description        : (Not Specified)
Admin State        : Up                 Oper State       : Down
Node ID            : 02:09:ff:00:00:00
Guard Time         :    5 deciseconds   RPL Node         : rplNone
Max Revert Time    :  300 seconds       Time to Revert   : N/A
CCM Hold Down Time :    0 centiseconds  CCM Hold Up Time :   20 deciseconds
Compatible Version : 2
APS Tx PDU         : Request State: 0xB
                     Sub-Code     : 0x0
                     Status       : 0x20  ( BPR )
                     Node ID      : 02:09:ff:00:00:00
Defect Status      :

Sub-Ring Type      : none

-------------------------------------------------------------------------------
Ethernet Ring Path Summary
-------------------------------------------------------------------------------
Path Port             Raps-Tag     Admin/Oper      Type            Fwd State
-------------------------------------------------------------------------------
  a  1/1/c1/1         1               Up/Down      normal          blocked
  b  1/1/c2/1         1               Up/Down      normal          blocked
===============================================================================
```

Configure control channel VPLS service

Paths a and b defined in the Ethernet ring must be added as SAPs into a VPLS service (standard VPLS in this example) using the **eth-ring** parameter. The SAP encapsulation values must match the values of the **port-and-raps-tag** configured for the associated path.

G.8032 uses the same R-APS tag value on all nodes on the ring, as configured in this example. However, the SR OS implementation relaxes this constraint by requiring the tag to match only on adjacent nodes.

```
# on PE-1:
configure {
    service {
        vpls "VPLS-1" {
            admin-state enable
            description "control channel VPLS 1 tag 1"
            service-id 1
            customer "1"
            sap 1/1/c1/1:1 {
                eth-ring 1
            }
            sap 1/1/c2/1:1 {
                eth-ring 1
            }
```

```
# on PE-2:
configure {
    service {
        vpls "VPLS-1" {
            admin-state enable
            description "control channel VPLS 1 tag 1"
            service-id 1
            customer "1"
            sap 1/1/c1/1:1 {
                eth-ring 1
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            sap 1/1/c2/1:1 {
                eth-ring 1
            }
```

```
# on PE-3:
configure {
    service {
        vpls "VPLS-1" {
            admin-state enable
            description "control channel VPLS 1 tag 1"
            service-id 1
            customer "1"
            sap 1/1/c1/1:1 {
                eth-ring 1
            }
            sap 1/1/c2/1:1 {
                eth-ring 1
            }
```

A normal SAP or SDP can be added in a control channel VPLS on condition the **eth-ring** parameter is present. Any attempt to add a SAP or SDP without this parameter into a control channel VPLS results in error messages being displayed. To trigger the following error messages, SAP 1/1/c3/1:1 is added without the eth-ring parameter.

```
*[ex:/configure service vpls "VPLS-1" sap 1/1/c3/1:1]
A:admin@PE-1# commit
MINOR: SVCMGR #2590: configure service vpls "VPLS-1" sap 1/1/c1/1:1 - Invalid Ethernet ring
configuration - Ring control SAP cannot be added to service that contains non-ring SAPs or SDP
bindings - configure service vpls "VPLS-1" sap 1/1/c1/1:1 eth-ring
MINOR: SVCMGR #2590: configure service vpls "VPLS-1" sap 1/1/c1/1:1 - Invalid Ethernet ring
configuration - Ethernet Ring Control service should only have controls saps from same ring -
configure service vpls "VPLS-1" sap 1/1/c1/1:1 eth-ring
MINOR: SVCMGR #2590: configure service vpls "VPLS-1" sap 1/1/c2/1:1 - Invalid Ethernet ring
configuration - Ring control SAP cannot be added to service that contains non-ring SAPs or SDP
bindings - configure service vpls "VPLS-1" sap 1/1/c2/1:1 eth-ring
MINOR: SVCMGR #2590: configure service vpls "VPLS-1" sap 1/1/c2/1:1 - Invalid Ethernet ring
configuration - Ethernet Ring Control service should only have controls saps from same ring -
configure service vpls "VPLS-1" sap 1/1/c2/1:1 eth-ring
```

In non-failure conditions, the Ethernet ring is operationally up and the RPL is blocking successfully on ring node PE-2 port 1/1/c2/1, as expected from the RPL owner and RPL end configuration.

An overview of all of the rings can be shown using the following commands, in this case on node PE-2.

The following command on PE-2 shows the Ethernet ring status.

```
[/]
A:admin@PE-2# show eth-ring status

===============================================================================
Ethernet Ring (Status information)
===============================================================================
Ring   Admin  Oper       Path Information                MEP Information
ID     State  State  Path         Tag       State     Ctrl-MEP CC-Intvl Defects
-------------------------------------------------------------------------------
1      Up     Up     a - 1/1/c1/1     1     Up        Yes      1         -----
                     b - 1/1/c2/1     1     Up        Yes      1         -----
===============================================================================
Ethernet Tunnel MEP Defect Legend:
R = Rdi, M = MacStatus, C = RemoteCCM, E = ErrorCCM, X = XconCCM
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The following command on PE-2 shows the ring and path forwarding states.

```
[/]
A:admin@PE-2# show eth-ring

===============================================================================
Ethernet Rings (summary)
===============================================================================
Ring Int  Admin Oper           Paths Summary                      Path States
ID   ID   State State                                             a     b
-------------------------------------------------------------------------------
1    -    Up    Up    a - 1/1/c1/1     1     b - 1/1/c2/1     1     U     B
===============================================================================
Ethernet Ring Summary Legend:   B - Blocked    U - Unblocked
```

The **show eth-ring 1** command on the different nodes shows specific information for Ethernet ring 1:

```
[/]
A:admin@PE-1# show eth-ring 1

===============================================================================
Ethernet Ring 1 Information
===============================================================================
Description        : (Not Specified)
Admin State        : Up                 Oper State       : Up
Node ID            : 02:09:ff:00:00:00
Guard Time         :    5 deciseconds   RPL Node         : rplNone
Max Revert Time    :  300 seconds       Time to Revert   : N/A
CCM Hold Down Time :    0 centiseconds  CCM Hold Up Time :   20 deciseconds
Compatible Version : 2
APS Tx PDU         : N/A
Defect Status      :

Sub-Ring Type      : none

-------------------------------------------------------------------------------
Ethernet Ring Path Summary
-------------------------------------------------------------------------------
Path Port             Raps-Tag     Admin/Oper      Type            Fwd State
-------------------------------------------------------------------------------
  a  1/1/c1/1         1               Up/Up        normal          unblocked
  b  1/1/c2/1         1               Up/Up        normal          unblocked
===============================================================================
```

```
[/]
A:admin@PE-2# show eth-ring 1

===============================================================================
Ethernet Ring 1 Information
===============================================================================
Description        : (Not Specified)
Admin State        : Up                 Oper State       : Up
Node ID            : 02:0b:ff:00:00:00
Guard Time         :    5 deciseconds   RPL Node         : rplOwner
Max Revert Time    :   60 seconds       Time to Revert   : N/A
CCM Hold Down Time :    0 centiseconds  CCM Hold Up Time :   20 deciseconds
Compatible Version : 2
APS Tx PDU         : Request State: 0x0
                     Sub-Code     : 0x0
                     Status       : 0xE0  ( RB DNF BPR )
                     Node ID      : 02:0b:ff:00:00:00
Defect Status      :
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Sub-Ring Type      : none

-------------------------------------------------------------------------------
Ethernet Ring Path Summary
-------------------------------------------------------------------------------
Path Port             Raps-Tag     Admin/Oper      Type            Fwd State
-------------------------------------------------------------------------------
  a  1/1/c1/1         1               Up/Up        normal          unblocked
  b  1/1/c2/1         1               Up/Up        rplEnd          blocked
===============================================================================
```

Node PE-2 is the RPL owner and port 1/1/c2/1 is the RPL end. The **Max Revert Time** shows the configured value.

When a revert is pending after a failure restoration, the **Time to Revert** shows the number of seconds remaining before the revert occurs, as follows:

```
[/]
A:admin@PE-2# show eth-ring 1

===============================================================================
Ethernet Ring 1 Information
===============================================================================
Description        : (Not Specified)
Admin State        : Up                 Oper State       : Up
Node ID            : 02:0b:ff:00:00:00
Guard Time         :    5 deciseconds   RPL Node         : rplOwner
Max Revert Time    :   60 seconds       Time to Revert   :   49 seconds
CCM Hold Down Time :    0 centiseconds  CCM Hold Up Time :   20 deciseconds
Compatible Version : 2
APS Tx PDU         : N/A
Defect Status      :

Sub-Ring Type      : none

-------------------------------------------------------------------------------
Ethernet Ring Path Summary
-------------------------------------------------------------------------------
Path Port             Raps-Tag     Admin/Oper      Type            Fwd State
-------------------------------------------------------------------------------
  a  1/1/c1/1         1               Up/Up        normal          unblocked
  b  1/1/c2/1         1               Up/Up        rplEnd          unblocked
===============================================================================
```

On reversion, the following message is logged in log 99.

```
78 2023/05/05 16:12:16.588 CEST MINOR: ERING #2001 Base eth-ring-1
"Eth-Ring 1 path b changed fwd state to blocked"
```

The status of Ethernet ring 1 on PE-3 is as follows:

```
[/]
A:admin@PE-3# show eth-ring 1

===============================================================================
Ethernet Ring 1 Information
===============================================================================
Description        : (Not Specified)
Admin State        : Up                 Oper State       : Up
Node ID            : 02:0d:ff:00:00:00
Guard Time         :    5 deciseconds   RPL Node         : rplNone
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Max Revert Time    :  300 seconds       Time to Revert   : N/A
CCM Hold Down Time :    0 centiseconds  CCM Hold Up Time :   20 deciseconds
Compatible Version : 2
APS Tx PDU         : N/A
Defect Status      :

Sub-Ring Type      : none

-------------------------------------------------------------------------------
Ethernet Ring Path Summary
-------------------------------------------------------------------------------
Path Port             Raps-Tag     Admin/Oper      Type            Fwd State
-------------------------------------------------------------------------------
  a  1/1/c1/1         1              Up/Up        normal          unblocked
  b  1/1/c2/1         1              Up/Up        normal          unblocked
===============================================================================
```

Finally, the following commands on PE-2 show the details of the individual paths:

```
[/]
A:admin@PE-2# show eth-ring 1 path a

===============================================================================
Ethernet Ring 1 Path Information
===============================================================================
Description        : (Not Specified)
Port               : 1/1/c1/1           Raps-Tag         : 1
Admin State        : Up                 Oper State       : Up
Path Type          : normal             Fwd State        : unblocked
                                        Fwd State Change : 05/05/2023 16:11:17
Last Switch Command: noCmd
APS Rx PDU         : Request State: 0x0
                     Sub-Code     : 0x0
                     Status       : 0x20  ( BPR )
                     Node ID      : 02:0d:ff:00:00:00

===============================================================================

[/]
A:admin@PE-2# show eth-ring 1 path b

===============================================================================
Ethernet Ring 1 Path Information
===============================================================================
Description        : (Not Specified)
Port               : 1/1/c2/1           Raps-Tag         : 1
Admin State        : Up                 Oper State       : Up
Path Type          : rplEnd             Fwd State        : blocked
                                        Fwd State Change : 05/05/2023 16:12:17
Last Switch Command: noCmd
APS Rx PDU         : Request State: 0x0
                     Sub-Code     : 0x0
                     Status       : 0x20  ( BPR )
                     Node ID      : 02:0d:ff:00:00:00

===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Configure user data channel VPLS service

The user data channels are created on a separate VPLS, "VPLS-100" in the example. The ring data channels must be on the same ports as the corresponding control channels configured above. The access into the data services can use SAPs or SDPs.

```
# on PE-1:
configure {
    service {
        vpls "VPLS-100" {
            admin-state enable
            description "data channel VPLS 100"
            service-id 100
            customer "1"
            sap 1/1/c1/1:100 {
                eth-ring 1
            }
            sap 1/1/c2/1:100 {
                eth-ring 1
            }
            sap 1/1/c3/1:100 {
            }
        }
```

```
# on PE-2:
configure {
    service {
        vpls "VPLS-100" {
            admin-state enable
            description "data channel VPLS 100"
            service-id 100
            customer "1"
            sap 1/1/c1/1:100 {
                eth-ring 1
            }
            sap 1/1/c2/1:100 {
                eth-ring 1
            }
            sap 1/1/c3/1:100 {
            }
        }
```

```
# on PE-3:
configure {
    service {
        vpls "VPLS-100" {
            admin-state enable
            description "data channel VPLS 100"
            service-id 100
            customer "1"
            sap 1/1/c1/1:100 {
                eth-ring 1
            }
            sap 1/1/c2/1:100 {
                eth-ring 1
            }
            sap 1/1/c3/1:100 {
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The following command on PE-1 shows all the SAPs that are configured to use Ethernet rings.

```
[/]
A:admin@PE-1# show service sap-using eth-ring

===============================================================================
Service Access Points (Ethernet Ring)
===============================================================================
SapId                   SvcId       Eth-Ring Path Admin Oper  Blocked Control/
                                                  State State         Data
-------------------------------------------------------------------------------
1/1/c1/1:1              1           1        a    Up    Up    No      Ctrl
1/1/c2/1:1              1           1        b    Up    Up    No      Ctrl
1/1/c1/1:100            100         1        a    Up    Up    No      Data
1/1/c2/1:100            100         1        b    Up    Up    No      Data
-------------------------------------------------------------------------------
Number of SAPs : 4
===============================================================================
```

Debug

To emulate a failure on Ethernet ring 1, the unblocked port 1/1/c1/1 on node PE-2 is disabled, as follows.

```
# on PE-2:
configure {
    port 1/1/c1/1 {
        admin-state disable
```

The following messages are logged in log 99 when the failure occurs:

```
88 2023/05/05 16:15:44.598 CEST MINOR: ETH_CFM #2001 Base
"MEP 1/1/1232 highest defect is now defRemoteCCM"

87 2023/05/05 16:15:40.798 CEST MAJOR: SVCMGR #2210 Base
"Processing of an access port state change event is finished and the status of all affected
SAPs on port 1/1/c1/1 has been updated."

86 2023/05/05 16:15:40.783 CEST MINOR: ERING #2001 Base eth-ring-1
"Eth-Ring 1 path b changed fwd state to unblocked"

85 2023/05/05 16:15:40.783 CEST MINOR: ERING #2001 Base eth-ring-1
"Eth-Ring 1 path a changed fwd state to blocked"

84 2023/05/05 16:15:40.782 CEST WARNING: SNMP #2004 Base 1/1/c1/1
"Interface 1/1/c1/1 is not operational"
```

For troubleshooting, the **tools dump eth-ring** *<ring-index>* command displays path information, the internal state of the control protocol, related statistics information and up to the last 20 protocol events (including messages sent and received, and the expiration of timers). An associated parameter **clear** exists, clearing the event information in this output when the command is entered. The following is an example of the output on node PE-2 with port 1/1/c1/1 disabled.

```
[/]
A:admin@PE-2# tools dump eth-ring 1

ringId 1 (Up/Up): numPaths 2 nodeId 02:0b:ff:00:00:00
 SubRing: none (interconnect ring 0, propagateTc  No), Cnt 0
  path-a, port 1/1/c1/1 (Down), tag 1.0(Dn) status (Up/Dn/Blk)
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    cc (Dn/Up): Cnt 3/2 tm 000 00:18:38.370/000 00:14:03.350
    state: Cnt 7 B/F 000 00:18:34.550/000 00:14:10.340, flag: 0x0
path-b, port 1/1/c2/1 (Up), tag 1.0(Up) status (Up/Up/Fwd)
    cc (Dn/Up): Cnt 2/2 tm 000 00:08:53.380/000 00:08:56.620
    state: Cnt 10 B/F 000 00:15:10.360/000 00:18:34.550, flag: 0x0
FsmState=  PROT, Rpl = Owner, revert = 60 s, guard = 5 ds
  Defects =
  Running Timers = PduReTx
  lastTxPdu = 0xb000 Sf
  path-a Normal, RxId(I)= 02:0d:ff:00:00:00, rx= v1-0x0020 Nr, cmd= None
  path-b Rpl, RxId= 02:0d:ff:00:00:00, rx= v1-0xb020 Sf, cmd= None
DebugInfo:  aPathSts 6, bPathSts 5, pm (set/clr) 0/0, txFlush 0
  RxRaps: ok 28 nok 0 self 218, TmrExp - wtr 3(1), grd 4, wtb 0
  Flush: cnt 15 (9/6/0) tm 000 00:18:38.340-000 00:18:38.340 Out/Ack 0/1
  RxRawRaps: aPath 156 bPath 169 vPath 0
  Now: 000 00:19:09.390 , softReset: No - noTx 0

Seq Event  RxInfo(Path: NodeId-Bytes)
           state:TxInfo (Bytes)            Dir  pA  pB        Time
=== ===== ============================  ===== === === ================
007   pdu B: 02:09:ff:00:00:00-0xb040 Sf(DNF)
            PEND-G: 0x0020  Nr            Rx<-- Fwd Blk 000 00:08:59.410
008   pdu B: 02:09:ff:00:00:00-0x0000 Nr
            PEND-G: 0x0020  Nr            Rx<-- Fwd Blk 000 00:08:59.410
009   pdu A: 02:09:ff:00:00:00-0xb040 Sf(DNF)
            PEND-G: 0x0020  Nr            Rx<-- Fwd Blk 000 00:08:59.420
010   pdu A: 02:09:ff:00:00:00-0x0000 Nr
            PEND-G: 0x0020  Nr            Rx<-- Fwd Blk 000 00:08:59.420
011   pdu B: 02:09:ff:00:00:00-0x0000 Nr
            PEND-G: 0x0020  Nr            Rx<-- Fwd Blk 000 00:08:59.510
012   pdu A: 02:09:ff:00:00:00-0x0000 Nr
            PEND-G: 0x0020  Nr            Rx<-- Fwd Blk 000 00:08:59.510
013   pdu B: 02:09:ff:00:00:00-0x0000 Nr
            PEND  : 0x0020  Nr            Rx<-- Fwd Blk 000 00:08:59.610
014   pdu A: 02:09:ff:00:00:00-0x0000 Nr
            PEND  : 0x0020  Nr            Rx<-- Fwd Blk 000 00:08:59.610
015  xWtr
            IDLE  : 0x00e0  Nr(RB DNF)    Tx--> Fwd Blk 000 00:10:14.360
016   aDn
            PROT  : 0xb000  Sf            TxF-> Blk Fwd 000 00:13:17.200
017   pdu B: 02:0d:ff:00:00:00-0xb020 Sf
            PROT  : 0xb000  Sf            RxF<- Blk Fwd 000 00:13:20.350
018   pdu A: 02:0d:ff:00:00:00-0x0020 Nr
            PROT  : 0xb000  Sf            Rx<-- Blk Fwd 000 00:14:04.440
019   pdu B: 02:0d:ff:00:00:00-0x0020 Nr
            PROT  : 0xb000  Sf            Rx<-- Blk Fwd 000 00:14:04.440
000   aUp
            PEND-G: 0x0000  Nr            Tx--> Blk Fwd 000 00:14:05.360
001   pdu A: 02:0d:ff:00:00:00-0x0020 Nr
            PEND  : 0x0000  Nr            Rx<-- Blk Fwd 000 00:14:10.340
002   pdu
            PEND  :                       ----- Fwd Fwd 000 00:14:10.340
003   pdu B: 02:0d:ff:00:00:00-0x0020 Nr
            PEND  :                       Rx<-- Fwd Fwd 000 00:14:10.340
004  xWtr
            IDLE  : 0x00a0  Nr(RB )       TxF-> Fwd Blk 000 00:15:10.360
005   aDn
            PROT  : 0xb000  Sf            TxF-> Blk Fwd 000 00:18:34.550
006   pdu B: 02:0d:ff:00:00:00-0xb020 Sf
            PROT  : 0xb000  Sf            RxF<- Blk Fwd 000 00:18:38.340
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Conclusion

Ethernet ring APS provides an optimal solution for designing native Ethernet services with ring topology. This protocol provides simple configuration, operation, and guaranteed fast protection time. SR OS also has a flexible encapsulation that allows dot1Q, QinQ, or PBB for the ring traffic. Ethernet ring APS can be utilized for various services such as mobile backhaul, business VPN access, aggregation, and core.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

GRE Tunnel Origination and Termination Using Non-system IP Addresses

This chapter provides information about GRE tunnel origination and termination using non-system IP addresses.

Topics in this chapter include:

Applicability

This chapter was initially written based on SR OS Release 16.0.R5, but the MD-CLI in the current edition corresponds to SR OS Release 23.3.R2. GRE SDPs and auto-bind GRE tunnels can originate and terminate on a non-system IP address in SR OS Release 16.0.R4 or later.

Overview

For scaling purposes, service providers typically deploy seamless MPLS or inter-AS scenarios. In many cases, the system IP address cannot be leaked between domains and a separate loopback address is used to terminate tunnels. GRE termination on a non-system IP address is supported in the following services:

- VPLS with manually configured GRE spoke-SDPs
- VPLS with BGP-AD using provisioned GRE SDPs (**provisioned-sdp use** or **provisioned-sdp prefer** commands)
- BGP-VPLS using provisioned GRE SDPs
- Epipe with manually configured GRE spoke-SDPs
- Epipe with BGP-VPWS using provisioned GRE SDPs
- VPRN with manually configured GRE spoke-SDPs
- VPRN with auto-bind GRE tunnel
- IES with manually configured GRE spoke-SDPs

This chapter focuses on MPLS-over-GRE termination, but IP-over-GRE termination is also supported.

MPLS-over-GRE termination

GRE termination applies to GRE SDPs and auto-bind GRE tunnels concurrently on a system interface and on non-system interfaces with a subnet that is up to and including /16. In the following example, the non-system loopback address 10.0.1.1 with a subnet of /24 is configured as GRE termination on PE-1:

```
# on PE-1:
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
configure {
    router "Base" {
        interface "lo1" {
            loopback
            gre-termination true
            ipv4 {
                primary {
                    address 10.0.1.1
                    prefix-length 24
                }
            }
        }
```

Only one interface can be configured as GRE termination. The following error is raised when attempting to configure a second loopback interface "lo2" as GRE termination on PE-1:

```
*[ex:/configure router "Base" interface "lo2"]
A:admin@Dut-A# commit
MINOR: COMMON #238: configure router "Base" interface "lo2" - Configuration change failed
validation - Multiple interfaces with gre-termination set in the router
```

Although the preceding examples are for loopback interfaces, GRE termination can also be configured on other router interfaces, but only one per node. The following shows an attempt to configure interface "int-PE-1-PE-2" on PE-1 as GRE termination. The same error message is raised. However, if it were the first interface on the node to be configured as GRE termination, the configuration would be accepted.

```
*[ex:/configure router "Base" interface "int-PE-1-PE-2"]
A:admin@Dut-A# commit
MINOR: COMMON #238: configure router "Base" interface "int-PE-1-PE-2" - Configuration change
failed validation - Multiple interfaces with gre-termination set in the router
```

The maximum size of the GRE termination subnet is /16.

GRE termination cannot be applied on the following interface types:

- Unnumbered network IP interfaces
- IES interfaces
- VPRN interfaces
- CSC VPRN interfaces

MPLS-over-GRE origination

GRE SDPs and auto-bind GRE tunnels can originate and terminate on a non-system IP address. Manually configured SDPs can be configured with a non-system IP address as the far-end address. Optionally, a non-system local-end address can be configured for generating GRE from an interface other than the system interface. In the following example on PE-1, GRE SDP 120 uses loopback address 10.0.1.1 as the local-end address and 10.0.2.1 on PE-2 as the far-end address.

```
# on PE-1:
configure {
    service {
        sdp 120 {
            admin-state enable
            local-end 10.0.1.1
            far-end {
                ip-address 10.0.2.1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
        }
```

The local-end IP address can only be configured for GRE SDPs; the following error message is raised when attempting to configure an MPLS SDP with a local-end address:

```
*[ex:/configure service sdp 122]
A:admin@PE-1# commit
MINOR: SVCMGR #7720: configure service sdp 122 local-end - Invalid SDP configuration -
local-end is not supported for this sdp delivery-type.
```

The **local-end** parameter value complies with the following rules:

- A maximum of 15 distinct address values can be configured for all GRE SDPs in the **configure service sdp local-end** context, and all L2oGRE SDPs under the **configure service system gre-eth-bridged tunnel-termination** context.
- The same source address cannot be used in both contexts because an address configured for an L2oGRE SDP matches an internally created interface that is not available to other applications.
- The local-end address of a GRE SDP, when different from the system address, need not match the primary address of an interface that has the MPLS-over-GRE termination subnet configured, unless a GRE SDP or tunnel from the far-end router terminates on this address.

The primary IPv4 address of any local network IP interface, loopback or not, may be used. The following shows that IP address 192.168.12.1, as the IP address of the previously mentioned interface "int-PE-1-PE-2" toward PE-2, can be used as the local-end address:

```
# on PE-1:
configure {
    service {
        sdp 123 {
            admin-state enable
            local-end 192.168.12.1
            far-end {
                ip-address 10.0.2.1
            }
        }
```

The following shows that an error message is raised when attempting to configure an invalid local-end IP address, that is, an IP address that is not primary on a local router interface. In this case, local-end IP address 10.99.1.1 does not exist on PE-1.

```
*[ex:/configure service sdp 120]
A:admin@PE-1# commit
MINOR: MGMT_CORE #4001: configure service sdp 120 - sdp 120: router interface with address
10.99.1.1 does not exist, or is not primary IPv4 address - configure router "Base" description
```

For services that support auto-binding to a GRE tunnel, the following command configures a single alternate source address (in this case, 10.0.1.1) per system:

```
# on PE-1:
configure {
    service {
        system {
            vpn-gre-source-ip 10.0.1.1
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The default value of the single source address is the primary IPv4 address of the system interface. The value of the **vpn-gre-source-ip** parameter can be changed at any time. After a new value is configured, the system address will not be used in services that bind to the GRE tunnel.

The **vpn-gre-source-ip** parameter value complies with the following rules:

- This single source address counts toward the maximum of 15 distinct address values per system used by all GRE SDPs under the **configure service sdp local-end** context and all L2oGRE SDPs under the **configure service system gre-eth-bridged tunnel-termination** context.
- The same source address can be used in both **vpn-gre-source-ip** and **configure service sdp local-end** contexts.
- The same source address cannot be used in both **vpn-gre-source-ip** and **configure service system gre-eth-bridged tunnel-termination** contexts because an address configured for an L2oGRE SDP matches an internally created interface that is not available to other applications.
- The **vpn-gre-source-ip** address, when different from the system IP address, need not match the primary address of an interface that has the MPLS-over-GRE termination subnet configured, unless a GRE SDP or tunnel from the far-end router terminates on this address.

Configuration

Figure 14: Example topology shows the example topology with three SR OS nodes in AS 64500. Services will be configured on PE-1 and PE-2, while RR-3 is a route reflector (RR).

Figure 14: Example topology

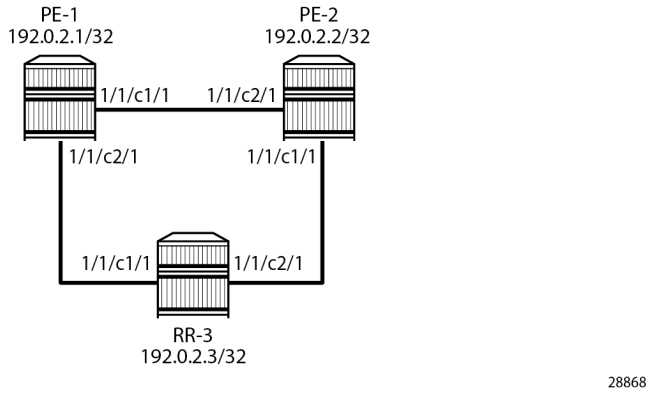

The initial configuration on the three PEs includes:

- cards, MDAs, ports
- router interfaces. The IP addresses shown on the figure are the system IP addresses 192.0.2.x/32.
- IS-IS as IGP (alternatively, OSPF can be used)

GRE SDP termination on non-system IP addresses will be configured in the following use cases:

- VPLS with manually configured T-LDP signaled SDP
- Epipe with manually configured T-LDP signaled SDP
- BGP-VPLS using a provisioned BGP-signaled SDP

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- BGP-AD in VPLS using a provisioned T-LDP signaled SDP
- BGP-VPWS using a provisioned BGP-signaled SDP
- VPRN with manually configured T-LDP signaled SDP
- VPRN with auto-bind to GRE tunnel
- IES with manually configured T-LDP signaled SDP

MPLS-over-GRE termination

On PE-1, PE-2, and RR-3, loopback interface "lo1" is configured as GRE termination with IPv4 address 10.0.x.1/24 for PE-x. The configuration on PE-1 is as follows:

```
# on PE-1:
configure {
    router "Base" {
        interface "lo1" {
            loopback
            gre-termination true
            ipv4 {
                primary {
                    address 10.0.1.1
                    prefix-length 24
                }
            }
        }
```

This loopback interface will be used in the SDP configuration. With a /24 subnet, the SDP origination can be any address in the subnet. This is useful for providing entropy in the outer IPv4 header for load-balancing over the IP network.

MPLS-over-GRE origination: SDP local end

The local-end address must be reachable from the far-end router that terminates the GRE SDP. Therefore, the interface for this address can be added to IGP or BGP. Alternatively, a static route can be configured on the far-end router. In this example, IS-IS is enabled on the loopback interface with GRE termination, as follows:

```
# on PE-1, PE-2, RR-3:
configure {
    router "Base" {
        isis 0 {
            interface "lo1" {
            }
```

On PE-1, the following SDPs are configured with far-end 10.0.2.1 on PE-2 and local-end 10.0.1.1: SDP 120 with T-LDP signaling (default) and SDP 121 with BGP signaling.

```
# on PE-1:
configure exclusive
    service {
        sdp 120 {
            admin-state enable
            local-end 10.0.1.1
            far-end {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                ip-address 10.0.2.1
            }
        }
        sdp 121 {
            admin-state enable
            local-end 10.0.1.1
            signaling bgp
            far-end {
                ip-address 10.0.2.1
            }
        }
```

T-LDP signaled GRE SDPs

When T-LDP signaled SDPs, such as SDP 120 in the preceding example, are configured, T-LDP sessions are auto-created toward the far end of the SDPs. By default, LDP uses the system IP address as source address. However, if the source address for the T-LDP session does not match the destination transport address set by the remote PE, the T-LDP session will not come up and the GRE SDP will remain down.Figure 15: Mismatched T-LDP transport addresses shows an example where SDP auto-created T-LDP sessions use the local system addresses 192.0.2.x and far-end addresses 10.0.0.x, so the GRE SDPs will not come up.

Figure 15: Mismatched T-LDP transport addresses

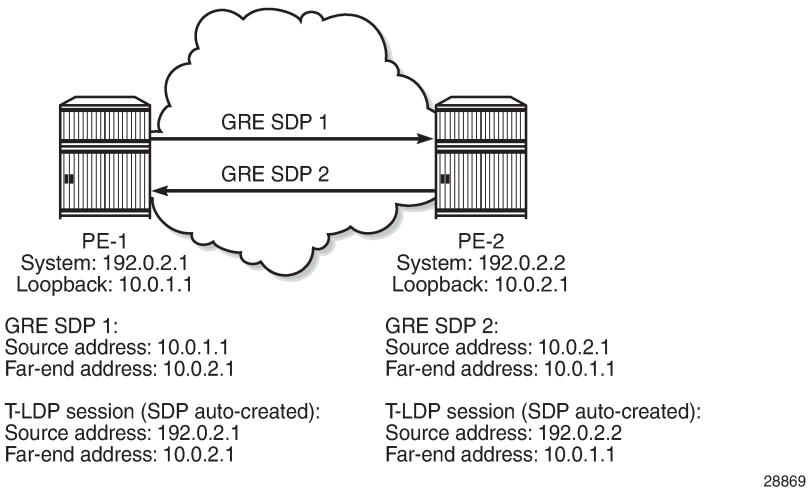

Therefore, the local transport address of the T-LDP session must match the local-end address of the GRE SDP in the PE. These T-LDP sessions can be manually provisioned or auto-created via peer templates. The following configures T-LDP sessions between the non-system IP addresses on PE-1 and PE-2.

```
# on PE-1:
configure {
    router "Base" {
        ldp {
            targeted-session {
                peer 10.0.2.1 {
                    local-lsr-id {
                        interface-name "lo1"
                    }
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }

# on PE-2:
configure {
    router "Base" {
        ldp {
            targeted-session {
                peer 10.0.1.1 {
                    local-lsr-id {
                        interface-name "lo1"
                    }
                }
            }
```

Figure 16: Matching T-LDP transport addresses shows the GRE T-LDP signaled SDPs with matching addresses for the T-LDP sessions.

Figure 16: Matching T-LDP transport addresses

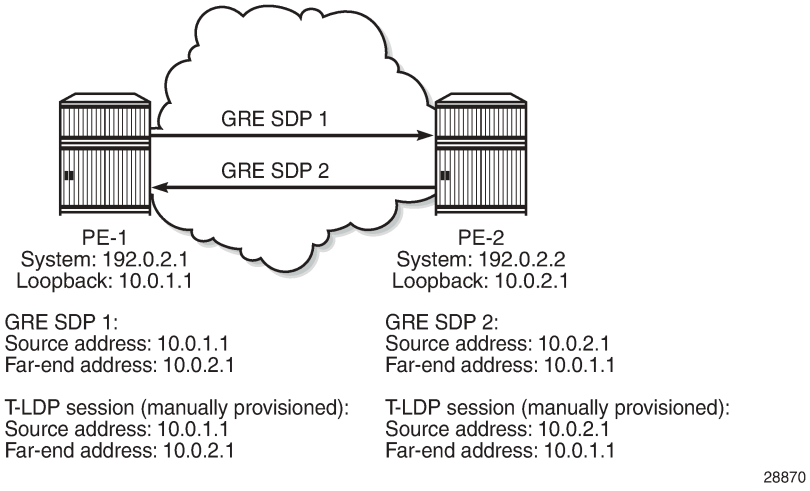

BGP configuration

In this example, the L2 and L3 services are configured on PE-1 and PE-2, while RR-3 acts as the RR. On PE-1, BGP is configured with neighbor 10.0.3.1 and local address 10.0.1.1, as follows. Address family L2-VPN is required for L2 services using BGP-VPLS, BGP-AD, and BGP-VPWS; address family VPN-IPv4 is used for VPRN services.

```
# on PE-1:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            rapid-withdrawal true
            split-horizon true
            group "internal" {
                type internal
                local-address 10.0.1.1
                family {
                    vpn-ipv4 true
                    l2-vpn true
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
            }
            neighbor "10.0.3.1" {
                group "internal"
            }
        }
```

The BGP configuration on PE-2 is similar with neighbor 10.0.3.1 and local address 10.0.2.1.

On RR-3, the BGP configuration is as follows.

```
# on RR-3:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            rapid-withdrawal true
            split-horizon true
            group "internal" {
                type internal
                local-address 10.0.3.1
                family {
                    vpn-ipv4 true
                    l2-vpn true
                }
                cluster {
                    cluster-id 10.0.3.1
                }
            }
            neighbor "10.0.1.1" {
                group "internal"
            }
            neighbor "10.0.2.1" {
                group "internal"
            }
        }
```

The loopback addresses 10.0.x.1 are configured for the local and neighbor addresses.

Note:
When the local address 10.0.x.1 is not configured, the system address 192.0.2.x will be used instead. However, in that case, no BGP sessions will be established and, therefore, no BGP routes will be exchanged between 192.0.2.x and 10.0.y.1, and no spoke-SDPs will be auto-created in L2 services using BGP-VPLS, BGP-AD, or BGP-VWPS. Likewise, no BGP-VPN routes will be exchanged between VPRNs on PE-1 and PE-2.

L2 services

Figure 17: L2 services on PE-1 and PE-2 shows the example topology with the following L2 services configured on PE-1 and PE-2:

- VPLS 1 with manually configured spoke-SDP 120:1
- Epipe 2 with manually configured spoke-SDP 120:2
- BGP-VPLS 3 using PW template 1 (BGP-signaled SDP 121 is used)
- LDP VPLS 4 with BGP-AD using PW template 1 (T-LDP signaled SDP 120 is used)
- BGP-VPWS Epipe 5 using PW template 1 (BGP-signaled SDP 121 is used)

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The CEs are VPRNs configured on the PEs and connected to the VPLSs via port cross-connect (PXC).

Figure 17: L2 services on PE-1 and PE-2

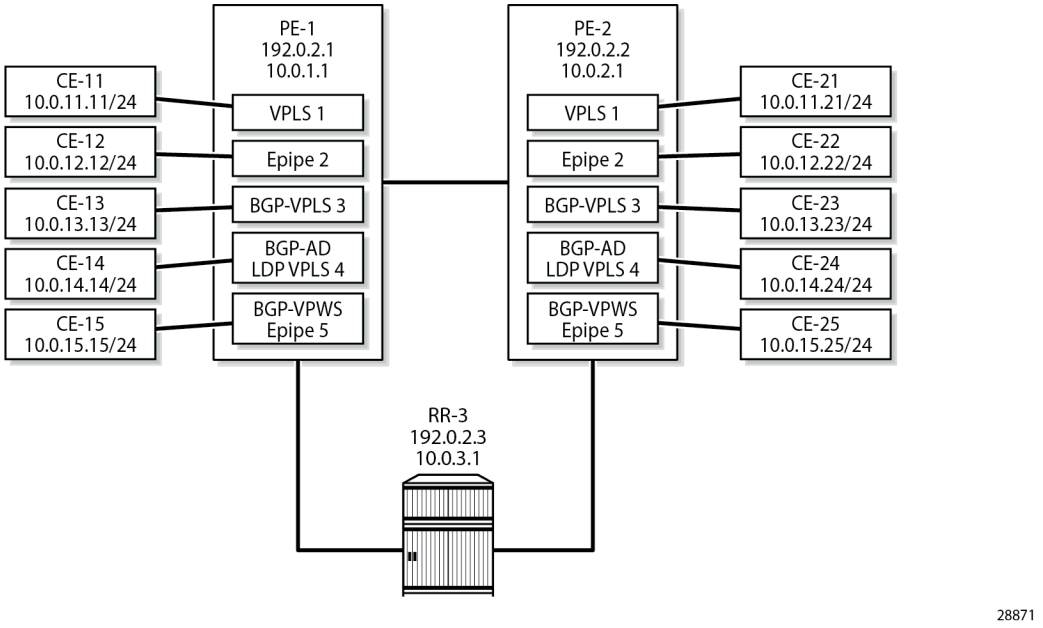

For a description of the BGP-VPLS parameters, see the BGP VPLS chapter; for BGP-AD, see the LDP VPLS Using BGP Auto-Discovery chapter; for BGP-VPWS, see the BGP Virtual Private Wire Services chapter. For BGP-VPLS, BGP-AD, and BGP-VPWS, PW template 1 is configured with the **provisioned-sdp use** command. The service configuration on PE-1 is as follows; the service configuration on PE-2 is similar.

```
# on PE-1:
configure {
    service {
        sdp 120 {
            admin-state enable
            local-end 10.0.1.1
            far-end {
                ip-address 10.0.2.1
            }
        }
        sdp 121 {
            admin-state enable
            local-end 10.0.1.1
            signaling bgp
            far-end {
                ip-address 10.0.2.1
            }
        }
        pw-template "PW1-use-prov-SDP" {
            pw-template-id 1
            provisioned-sdp use
        }
        vpls "VPLS-1" {
            admin-state enable
            description "VPLS 1 with manually configured spoke-SDP"
            service-id 1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            customer "1"
            spoke-sdp 120:1 {
            }
            sap pxc-10.a:1 {
            }
        }
        epipe "Epipe-2" {
            admin-state enable
            description "Epipe 2 with manually configured spoke-SDP"
            service-id 2
            customer "1"
            spoke-sdp 120:2 {
            }
            sap pxc-10.a:2 {
            }
        }
        vpls "BGP-VPLS-3" {
            admin-state enable
            description "BGP-VPLS with use provisioned SDP"
            service-id 3
            customer "1"
            bgp 1 {
                route-distinguisher "64500:3"
                route-target {
                    export "target:64500:3"
                    import "target:64500:3"
                }
                pw-template-binding "PW1-use-prov-SDP" {
                }
            }
            bgp-vpls {
                admin-state enable
                maximum-ve-id 100
                ve {
                    name "PE-1"
                    id 1
                }
            }
            sap pxc-10.a:3 {
            }
        }
        vpls "BGP-AD VPLS-4" {
            admin-state enable
            description "BGP-AD for LDP VPLS with use provisioned SDP"
            service-id 4
            customer "1"
            bgp 1 {
                route-distinguisher "64500:4"
                route-target {
                    export "target:64500:4"
                    import "target:64500:4"
                }
                pw-template-binding "PW1-use-prov-SDP" {
                }
            }
            bgp-ad {
                admin-state enable
                vpls-id "64500:4"
            }
            sap pxc-10.a:4 {
            }
        }
        epipe "BGP-VPWS-5" {
            admin-state enable
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            description "BGP-VPWS with use provisioned SDP"
            service-id 5
            customer "1"
            bgp 1 {
                route-distinguisher "64500:5"
                route-target {
                    export "target:64500:5"
                    import "target:64500:5"
                }
                pw-template-binding "PW1-use-prov-SDP" {
                }
            }
            bgp-vpws {
                admin-state enable
                local-ve {
                    name "PE-1"
                    id 1
                }
                remote-ve "PE-2" {
                    id 2
                }
            }
            sap pxc-10.a:5 {
            }
        }
```

The following BGP sessions are established between PE-1 and RR-3 for the VPN-IPv4 and L2VPN address families:

```
[/]
A:admin@PE-1# show router bgp summary all

===============================================================================
BGP Summary
===============================================================================
Legend : D - Dynamic Neighbor
===============================================================================
Neighbor
Description
ServiceId          AS PktRcvd InQ  Up/Down   State|Rcv/Act/Sent (Addr Family)
                      PktSent OutQ
-------------------------------------------------------------------------------
10.0.3.1
Def. Inst       64500      20    0 00h06m41s 0/0/0 (VpnIPv4)
                           23    0           3/3/3 (L2VPN)

-------------------------------------------------------------------------------
```

On PE-1, the following T-LDP session is established to 10.0.2.1 on PE-2:

```
[/]
A:admin@PE-1# show router ldp session ipv4

===============================================================================
LDP IPv4 Sessions
===============================================================================
Peer LDP Id         Adj Type  State         Msg Sent  Msg Recv  Up Time
-------------------------------------------------------------------------------
10.0.2.1:0          Targeted  Established   115       117       0d 00:09:26
-------------------------------------------------------------------------------
No. of IPv4 Sessions: 1
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

On PE-1, the following SDPs are created with far end 10.0.2.1 and GRE delivery. For SDP 120, T-LDP signaling is used; BGP signaling is used for SDP 121.

```
[/]
A:admin@PE-1# show service sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId  AdmMTU  OprMTU  Far End          Adm  Opr         Del     LSP   Sig
-------------------------------------------------------------------------------
120    0       8954    10.0.2.1         Up   Up          GRE     n/a   TLDP
121    0       8954    10.0.2.1         Up   Up          GRE     n/a   BGP
-------------------------------------------------------------------------------
Number of SDPs : 2
-------------------------------------------------------------------------------
Legend: R = RSVP, L = LDP, B = BGP, M = MPLS-TP, n/a = Not Applicable
        I = SR-ISIS, O = SR-OSPF, T = SR-TE, F = FPE
===============================================================================
```

On PE-1, the following SDP-bindings are used:

```
[/]
A:admin@PE-1# show service sdp-using

===============================================================================
SDP Using
===============================================================================
SvcId      SdpId              Type   Far End              Opr   I.Label E.Label
                                                         State
-------------------------------------------------------------------------------
1          120:1              Spok   10.0.2.1             Up    524278  524278
2          120:2              Spok   10.0.2.1             Up    524276  524276
3          121:4294967294     BgpVp* 10.0.2.1             Up    524280  524279
4          120:4294967295     BgpAd  10.0.2.1             Up    524275  524275
5          121:4294967293     BgpVp* 10.0.2.1             Up    524277  524277
-------------------------------------------------------------------------------
Number of SDPs : 5
-------------------------------------------------------------------------------
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

When the loopback interface "lo1" is configured as GRE termination on PE-1 and PE-2, the CEs can send traffic to each other. The following ping messages verify the connectivity between CE-11 and CE-21, CE-12 and CE-22, and so on:

```
[/]
A:admin@PE-1# ping 10.0.11.21 router-instance "CE-11" interval 0.1 output-format summary
PING 10.0.11.21 56 data bytes
!!!!!
---- 10.0.11.21 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 3.49ms, avg = 3.70ms, max = 4.11ms, stddev = 0.216ms

[/]
A:admin@PE-1# ping 10.0.12.22 router-instance "CE-12" interval 0.1 output-format summary
PING 10.0.12.22 56 data bytes
!!!!!
---- 10.0.12.22 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 3.54ms, avg = 4.58ms, max = 8.21ms, stddev = 1.82ms
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
[/]
A:admin@PE-1# ping 10.0.13.23 router-instance "CE-13" interval 0.1 output-format summary
PING 10.0.13.23 56 data bytes
!!!!!
---- 10.0.13.23 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 3.65ms, avg = 4.73ms, max = 8.67ms, stddev = 1.97ms

[/]
A:admin@PE-1# ping 10.0.14.24 router-instance "CE-14" interval 0.1 output-format summary
PING 10.0.14.24 56 data bytes
!!!!!
---- 10.0.14.24 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 3.76ms, avg = 6.67ms, max = 13.3ms, stddev = 3.82ms

[/]
A:admin@PE-1# ping 10.0.15.25 router-instance "CE-15" interval 0.1 output-format summary
PING 10.0.15.25 56 data bytes
!!!!!
---- 10.0.15.25 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 3.66ms, avg = 4.64ms, max = 7.89ms, stddev = 1.63ms
```

L3 services

Figure 18: L3 services on PE-1 and PE-2 shows the example topology with the following three L3 services configured on PE-1 and PE-2:

- VPRN 6 with manually configured spoke-SDP 120:6
- VPRN 7 with auto-bind to GRE tunnel
- IES 8 with manually configured spoke-SDP 120:8

Figure 18: L3 services on PE-1 and PE-2

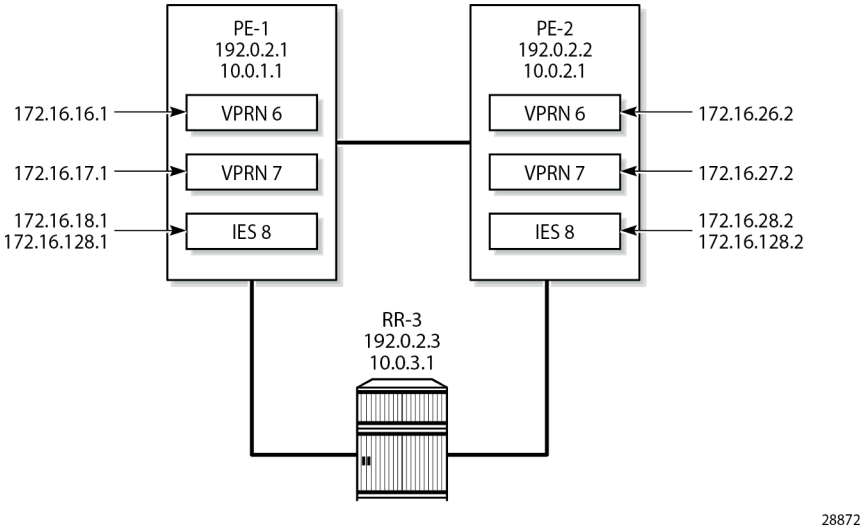

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

VPRN 6 is configured with a loopback interface and a GRE spoke-SDP, as follows:

```
# on PE-1:
configure {
    service {
        system {
            bgp-auto-rd-range {
                ip-address 10.0.1.1
                community-value {
                    start 60000
                    end 65000
                }
            }
        }
        vprn "VPRN-6 with GRE spoke-SDP" {
            admin-state enable
            service-id 6
            customer "1"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher auto-rd
                    vrf-target {
                        community "target:64500:6"
                    }
                }
            }
            interface "lo6" {
                loopback true
                ipv4 {
                    primary {
                        address 172.16.16.1
                        prefix-length 32
                    }
                }
            }
            spoke-sdp 120:6 {
            }
        }
```

The following forwarding information base (FIB) for VPRN 6 shows that the remote prefix is reachable via a transport tunnel using SDP 120:

```
[/]
A:admin@PE-1# show router 6 fib 1

===============================================================================
FIB Display
===============================================================================
Prefix [Flags]                                                  Protocol
  NextHop
-------------------------------------------------------------------------------
172.16.16.1/32                                                  LOCAL
  172.16.16.1 (lo6)
172.16.26.2/32                                                  BGP_VPN
  10.0.2.1 (VPRN Label:524273 Transport:SDP:120)
-------------------------------------------------------------------------------
Total Entries : 2
-------------------------------------------------------------------------------
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

VPRN 7 is configured with **auto-bind-tunnel** and the tunnel needs to be resolved using GRE. For services that support auto-binding to a GRE tunnel, the **vpn-gre-source-ip** parameter defines a single alternate source address for all VPRNs on the system. On PE-1, the configuration is as follows:

```
# on PE-1:
configure {
    service {
        system {
            vpn-gre-source-ip 10.0.1.1
        }
        vprn "VPRN-7 with auto-bind GRE" {
            admin-state enable
            service-id 7
            customer "1"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher auto-rd
                    vrf-target {
                        community "target:64500:7"
                    }
                    auto-bind-tunnel {
                        resolution filter
                        resolution-filter {
                            gre true
                        }
                    }
                }
            }
            interface "lo7" {
                loopback true
                ipv4 {
                    primary {
                        address 172.16.17.1
                        prefix-length 24
                    }
                }
            }
        }
```

The following FIB for VPRN 7 shows that the remote prefix is reachable via a GRE transport tunnel:

```
[/]
A:admin@PE-1# show router 7 fib 1

===============================================================================
FIB Display
===============================================================================
Prefix [Flags]                                                  Protocol
  NextHop
-------------------------------------------------------------------------------
172.16.17.0/24                                                  LOCAL
  172.16.17.0 (lo7)
172.16.27.0/24                                                  BGP_VPN
  10.0.2.1 (VPRN Label:524271 Transport:GRE)
-------------------------------------------------------------------------------
Total Entries : 2
-------------------------------------------------------------------------------
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

IES 8 has an interface with a manually configured GRE spoke-SDP, as follows:

```
# on PE-1:
configure {
    service {
        ies "IES-8" {
            admin-state enable
            service-id 8
            customer "1"
            interface "int-IES8-PE-1-PE-2" {
                spoke-sdp 120:8 {
                }
                ipv4 {
                    primary {
                        address 172.16.128.1
                        prefix-length 30
                    }
                }
            }
            interface "lo8" {
                loopback true
                ipv4 {
                    primary {
                        address 172.16.18.1
                        prefix-length 24
                    }
                }
            }
        }
```

On PE-1, the connectivity over the GRE spoke-SDP is verified as follows:

```
[/]
A:admin@PE-1# ping 172.16.128.2 interval 0.1 output-format summary
PING 172.16.128.2 56 data bytes
!!!!!
---- 172.16.128.2 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.57ms, avg = 2.73ms, max = 3.01ms, stddev = 0.168ms
```

Conclusion

By default, GRE SDPs and auto-bind GRE tunnels are originated and terminated on the system IP address, but it is possible to use non-system IP addresses. This is useful in cases where the system IP address cannot be leaked between domains and a separate loopback address must be used to terminate tunnels.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Network Group Encryption Helper

This chapter describes the network group encryption (NGE) helper.

Topics in this chapter include:

Applicability

The information and configuration in this chapter are based on SR OS Release 23.3.R1. Network group encryption (NGE) helpers require use of the VSR-a or the VSR-I and can be deployed with 7750 SR and 7950 XRS.

Overview

The NGE helper enables NGE security for services configured on the 7750 SR or 7950 XRS (hereafter referred to as the router) that require additional confidentiality and integrity.

Multiple NGE helpers can be deployed with a router depending on the encrypted services throughput requirements required by the operator. Figure 19: General architecture using an NGE helper shows the general architecture using an NGE helper.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 19: General architecture using an NGE helper

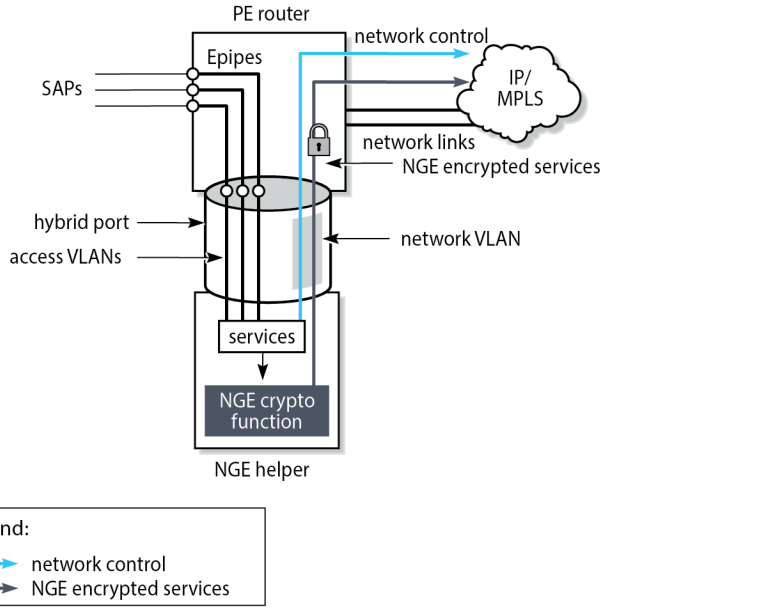

Each NGE helper is connected to the router using an access interface and a network interface, where both interfaces are configured on the NGE helper and on the router. A hybrid port can be used on the router and NGE helper to optimize the deployment, so one physical port is required on the router and NGE helper.

SAPs are configured on the router using an Epipe directed toward the NGE helper access interface. Unencrypted traffic that is received on the SAP interface is sent through the Epipe to the NGE helper which encrypts the traffic before sending it toward the network. The network interface on the NGE helper is enabled with minimal network control plane functions toward the router. The network control plane of the router performs the majority of network level processing and forwarding of NGE encrypted services.

The NGE helper supports services-based encryption, including:

- VPRN encryption
- SDP encryption
- PW-template encryption

Note: In SR OS Release 23.3.R1, all services-based encryption can be configured in classic CLI, whereas in MD-CLI, only PW-template encryption can be configured.

Router interface encryption and port-level encryption are not supported by the NGE helper.

Scenarios for encrypting services

The following main services scenarios are supported:

- **VPRN encryption using auto-bind services for both MPLS (LDP or RSVP-TE signaled tunnels) and GRE transport**

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

This scenario uses BGP to advertise the NGE helper IP address to remote NGE helpers. Remote NGE helpers can then send VPRN traffic to other NGE helpers to be processed for the associated destination SAP. This scenario uses VPRN-level NGE.

- **NG-MVPN with VPRN encryption using MLDP tunnels from the NGE helper to the router**

  This scenario uses a similar setup to VPRN encryption, with the difference that MLDP tunnels are also established between the NGE helper and the router where the point-to-multipoint tree branches from for the NG-MVPN service. This scenario uses VPRN-level NGE.

- **T-LDP signaled Epipe or VPLS services using LDP or RSVP-TE transport tunnels**

  T-LDP sessions are established from the NGE helper to the remote PEs to establish Epipe or VPLS services. The transport of these services focuses on LDP or LDP with RSVP-TE. Where GRE is possible, GRE support of VPLS or VPWS mainly uses BGP VPLS or BGP VPWS with auto-GRE SDP, because this use case is prevalent with SAR-Hm/Hmc deployments. This scenario uses SDP-level NGE.

- **L2 services using BGP VPLS or BGP VPWS auto-GRE SDP**

  This scenario is similar to the VPRN auto-bind scenario, except that a BGP session is used to advertise L2 routes to and from the NGE helper where remote PEs can send GRE L2 packets encrypted with the associated NGE configuration under the **pw-template** context.

Configuration

NGE configuration

NGE configuration is managed by the Network Services Platform Network Functions Manager - Packet (NSP NFM-P). Operators use the NSP NFM-P to configure:

- global encryption labels
- key groups
- VPRN-level encryption – setting the inbound and outbound key groups on VPRN-based services
- SDP-level encryption – setting the inbound and outbound key groups on selected SDPs
- PW-template level encryption – setting the inbound and outbound key groups on selected PW templates

Note: In this chapter, the NSP NFM-P is not used. Therefore, the remainder of the chapter focuses on PW-template level encryption, that can be configured in MD-CLI in SR OS Release 23.3.R1.

Group encryption configuration

In this example, encryption keygroup 1 is configured manually on NGE-1:

```
# on NGE-1:
configure {
    group-encryption {
        group-encryption-label 100
        encryption-keygroup 1 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            keygroup-name "KG1"
            active-outbound-security-association 1
            security-association 1 {
                authentication-key
0x11111111000000001111111100000000111111110000000011111111100000000
                encryption-key 0x11111111000000001111111100000000
            }
            security-association 2 {
                authentication-key
0x22222222000000002222222200000000222222220000000022222222200000000
                encryption-key 0x22222222000000002222222200000000
            }
            security-association 3 {
                authentication-key
0x33333333000000003333333300000000333333330000000033333333300000000
                encryption-key 0x33333333000000003333333300000000
            }
            security-association 4 {
                authentication-key
0x44444444000000004444444400000000444444440000000044444444400000000
                encryption-key 0x44444444000000004444444400000000
            }
        }
```

The authentication key and the encryption key are configured as cleartext. After configuration, they are never displayed in their cleartext form. The security parameter index (SPI) value in the security association is a node-wide unique value.

PW-template configuration

On NGE-1, PW template 2 is configured with encryption keygroup 1:

```
# on NGE-1:
    service {
        pw-template "2" {
            auto-gre-sdp true
            vc-type vlan
            split-horizon-group {
                name "SHG"
            }
            encryption-keygroup {
                inbound 1
                outbound 1
            }
        }
```

BGP configuration

BGP must be enabled on the router and the NGE helper for the L2-VPN address family for the following services:

- BGP VPWS with auto-GRE SDP (where NGE is configured under the **pw-template** context)
- BGP VPLS with auto-GRE SDP (where NGE is configured under the **pw-template** context)

Figure 20: BGP topology for learning BGP label routes shows the BGP topology for learning BGP label routes for those services.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 20: BGP topology for learning BGP label routes

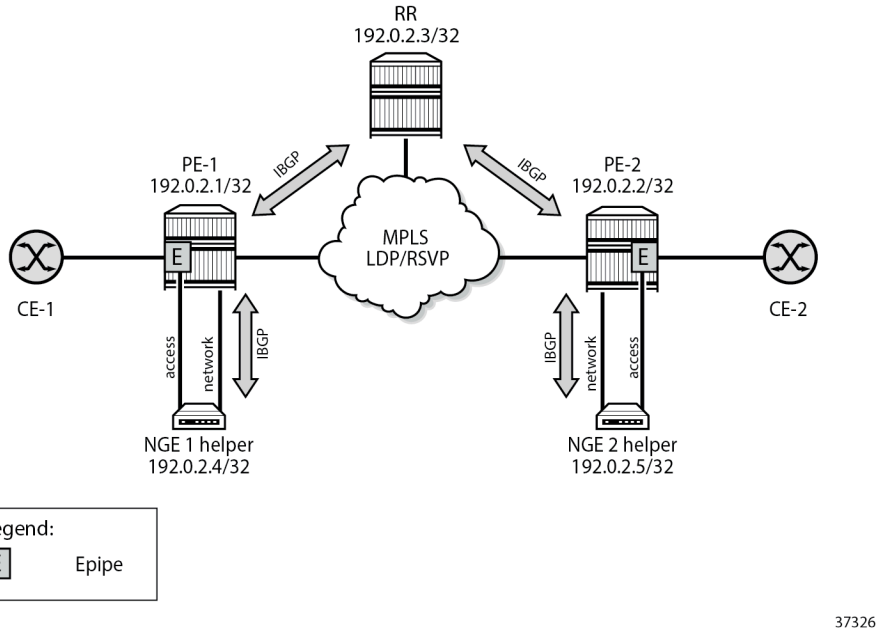

The following configures BGP on PE-1 to support the NGE 1 helper function:

```
# on PE-1:
configure {
    router "Base" {
        bgp {
            rapid-withdrawal true
            group "PE-1-NGE-1-RR" {
                peer-as 64496
                family {
                    l2-vpn true
                }
                cluster {
                    cluster-id 192.0.2.1
                }
            }
            group "core-RR" {
                peer-as 64496
                family {
                    l2-vpn true
                }
            }
            neighbor "192.0.2.3" {
                group "core-RR"
            }
            neighbor "192.0.2.4" {
                group "PE-1-NGE-1-RR"
            }
        }
```

The following configures BGP on PE-2 to support the NGE 2 helper function:

```
# on PE-2:
configure {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    router "Base" {
        bgp {
            rapid-withdrawal true
            group "PE-2-NGE-2-RR" {
                peer-as 64496
                family {
                    l2-vpn true
                }
                cluster {
                    cluster-id 192.0.2.2
                }
            }
            group "core-RR" {
                peer-as 64496
                family {
                    l2-vpn true
                }
            }
            neighbor "192.0.2.3" {
                group "core-RR"
            }
            neighbor "192.0.2.5" {
                group "PE-2-NGE-2-RR"
            }
        }
```

The BGP configuration on the NGE-1 helper is as follows:

```
# on NGE-1:
configure {
    router "Base" {
        bgp {
            rapid-withdrawal true
            group "RR-PE-1" {
                peer-as 64496
                family {
                    l2-vpn true
                }
            }
            neighbor "192.0.2.1" {
                group "RR-PE-1"
            }
        }
```

The BGP configuration on the NGE-2 helper is as follows:

```
# on NGE-2:
configure {
    router "Base" {
        bgp {
            rapid-withdrawal true
            group "RR-PE-2" {
                peer-as 64496
                family {
                    l2-vpn true
                }
            }
            neighbor "192.0.2.2" {
                group "RR-PE-2"
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Services configuration

BGP VPLS or BGP VPWS with auto-GRE SDP

Figure 21: NGE helper for BGP VPLS or BGP VPWS using GRE SDPs with auto-GRE SDP shows the operation of the NGE helper for BGP VPLS and BGP VPWS services that use GRE SDPs when auto-GRE SDP is configured on the associated PW template.

Figure 21: NGE helper for BGP VPLS or BGP VPWS using GRE SDPs with auto-GRE SDP

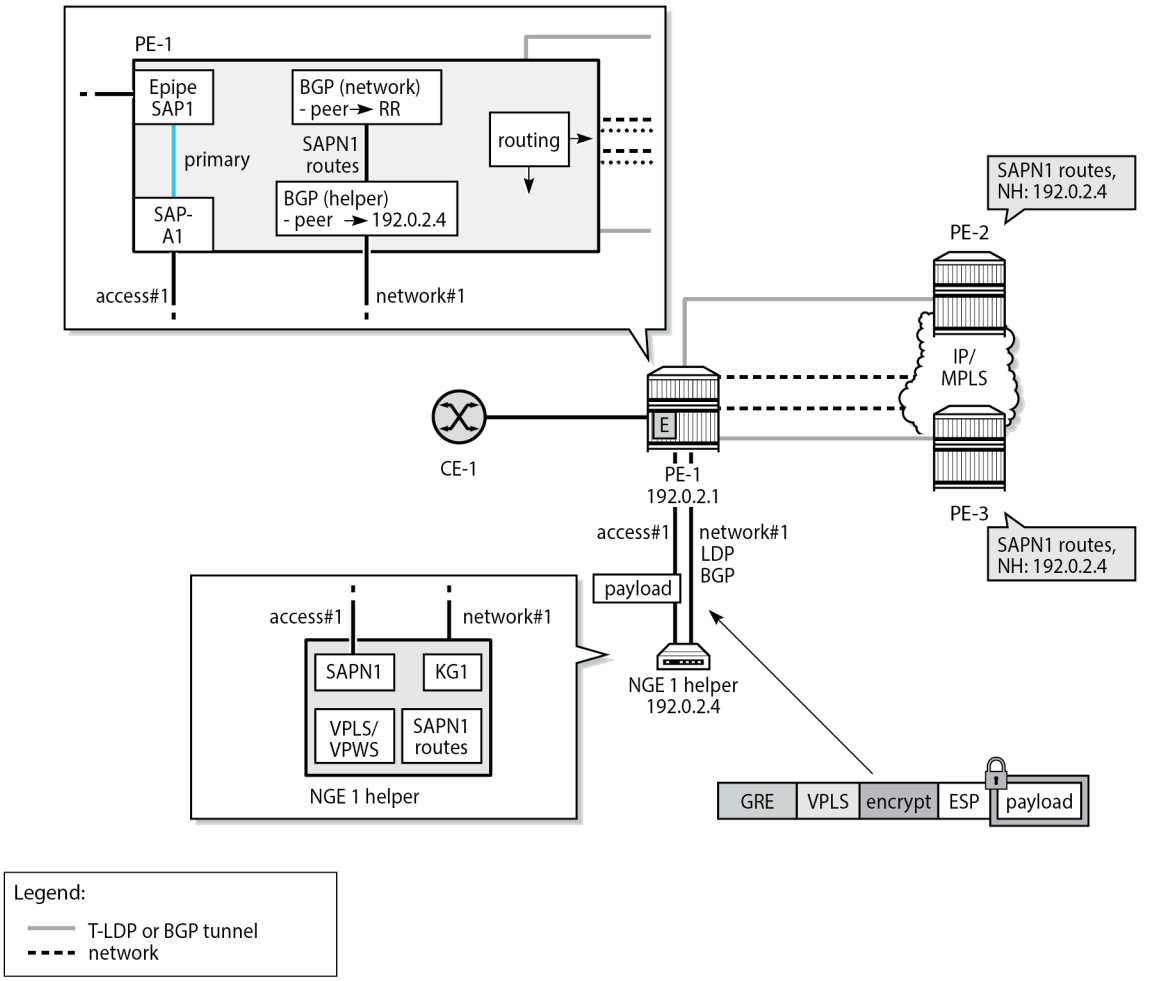

In this scenario, the VPLS or VPWS SAPN1 is configured on the NGE-1 helper. On PE-1, a local Epipe is configured that originates from the customer facing SAP1 and terminates on SAP-A1 connected to the NGE-1 helper. On PE-1, Epipes 100101 and 100201 are configured as follows:

```
# on PE-1:
configure {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    service {
        epipe "Epipe-100101" {
            admin-state enable
            service-id 100101
            customer "1"
            sap lag-1:101 {
                description "toward NGE-1 Epipe 101"
            }
            sap lag-11:101.1 {
                description "toward CE"
            }
        }
        epipe "Epipe-100201" {
            admin-state enable
            service-id 100201
            customer "1"
            sap lag-1:201 {
                description "toward NGE-1 VPLS 201"
            }
            sap lag-11:201.1 {
                description "toward CE"
            }
        }
```

On PE-1, the following network configurations are required to support encrypted services from the NGE-1 helper:

- any routing options that allow GRE packets received from the NGE helper to be routed to remote PEs
- BGP sessions for the L2-VPN address family, as described in the BGP configuration section

On the NGE-1 helper, the configuration includes:

- VPLS or VPWS SAPN1
- BGP session to PE-1 for the L2-VPN address family
- BGP VPLS or BGP VPWS using PW templates with auto-GRE SDP enabled
- NGE enabled on the PW templates for encrypting the VPLS or VPWS services using the PW templates

On NGE-1, Epipe 101 is a BGP VPWS with auto-GRE SDP. PW template 2 is configured with encryption keygroup 1. Epipe 101 is configured as follows:

```
# on NGE-1:
configure {
    service {
        pw-template "2" {
            auto-gre-sdp true
            vc-type vlan
            split-horizon-group {
                name "SHG"
            }
            encryption-keygroup {
                inbound 1
                outbound 1
            }
        }
        epipe "Epipe-101" {
            admin-state enable
            description "BGP VPWS auto-gre SDP_PW template 2"
            service-id 101
            customer "1"
            bgp 1 {
                route-distinguisher "101:1"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                route-target {
                    export "target:101:1"
                    import "target:101:1"
                }
                pw-template-binding "2" {
                }
            }
            bgp-vpws {
                admin-state enable
                local-ve {
                    name "pe-1"
                    id 1
                }
                remote-ve "pe-2" {
                    id 2
                }
            }
            sap lag-1:101 {
            }
        }
```

In a similar way, VPLS 201 is a BGP VPLS with auto-GRE SDP. PW template 2 is configured with encryption keygroup 1. VPLS 201 is configured as follows:

```
# on NGE-1:
configure {
    service {
        vpls "VPLS-201" {
            admin-state enable
            description "BGP VPLS auto-gre SDP_PW template 2"
            service-id 201
            customer "1"
            bgp 1 {
                route-distinguisher "201:1"
                route-target {
                    export "target:201:1"
                    import "target:201:1"
                }
                pw-template-binding "2" {
                }
            }
            bgp-vpls {
                admin-state enable
                maximum-ve-id 10
                ve {
                    name "pe-1"
                    id 1
                }
            }
            sap lag-1:201 {
            }
        }
```

Verification

The following base information for the services shows that the services are operationally up, as well as their SAPs and SDP bindings:

```
[/]
A:admin@NGE-1# show service id 101 base
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Service Basic Information
===============================================================================
Service Id        : 101                Vpn Id            : 0
Service Type      : Epipe
MACSec enabled    : no
Name              : Epipe-101
Description       : BGP VPWS auto-gre SDP_PW template 2
Customer Id       : 1                  Creation Origin   : manual
Last Status Change: 03/31/2023 15:44:58
Last Mgmt Change  : 03/31/2023 15:44:58
Test Service      : No
Admin State       : Up                 Oper State        : Up
---snip---

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                            Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:lag-1:101                          q-tag        8936    8936    Up   Up
sdp:32767:4294967295 SB(192.0.2.5)     BgpVpws      0       8890    Up   Up
===============================================================================

[/]
A:admin@NGE-1# show service id 201 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 201                Vpn Id            : 0
Service Type      : VPLS
MACSec enabled    : no
Name              : VPLS-201
Description       : BGP VPLS auto-gre SDP_PW template 2
Customer Id       : 1                  Creation Origin   : manual
Last Status Change: 03/31/2023 15:42:13
Last Mgmt Change  : 03/31/2023 15:44:58
Etree Mode        : Disabled
Admin State       : Up                 Oper State        : Up
MTU               : 1514
SAP Count         : 1                  SDP Bind Count    : 1
---snip---

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                            Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:lag-1:201                          q-tag        8936    8936    Up   Up
sdp:32766:4294967294 SB(192.0.2.5)     BgpVpls      0       8890    Up   Up
===============================================================================
```

The following command shows the encryption keygroup 1 with the associated SDPs: SDP 32767 is auto-provisioned by BGP-VPWS in Epipe 101, and SDP 32766 by BGP-VPLS in VPLS 201.

```
[/]
A:admin@NGE-1# show group-encryption encryption-keygroup 1

===============================================================================
Encryption Keygroup Configuration Detail
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Keygroup Id        : 1
Keygroup Name      : KG1
Description        : None
Authentication Algo: sha256
Encryption Algo    : aes128
Active Outbound SA : 1
Activation Time    : 03/31/2023 15:42:12

-------------------------------------------------------------------------------
Security Associations
-------------------------------------------------------------------------------
Spi                : 1
Install Time       : 03/31/2023 15:42:12
Key CRC            : 0xf57dcffc

Spi                : 2
Install Time       : 03/31/2023 15:42:12
Key CRC            : 0x26134d07

Spi                : 3
Install Time       : 03/31/2023 15:42:12
Key CRC            : 0xde19ce91

Spi                : 4
Install Time       : 03/31/2023 15:42:12
Key CRC            : 0x5bbf4eb0

-------------------------------------------------------------------------------
Encryption Keygroup Forwarded Statistics
-------------------------------------------------------------------------------
Encrypted Pkts          : 22            Encrypted Bytes         : 2124
Decrypted Pkts          : 22            Decrypted Bytes         : 2124
-------------------------------------------------------------------------------
Encryption Keygroup Outbound Discarded Statistics (Pkts)
-------------------------------------------------------------------------------
Total Discard           : 0             Other                   : 0
-------------------------------------------------------------------------------
Encryption Keygroup Inbound Discarded Statistics (Pkts)
-------------------------------------------------------------------------------
Total Discard           : 0             Invalid Spi             : 0
Authentication Failure *: 0             Padding Error           : 0
Other                   : 0


-----------------------------------------------------------------
SDP Keygroup Association Table
-----------------------------------------------------------------
SDP ID                          Direction
-----------------------------------------------------------------
32766                           Inbound   Outbound
32767                           Inbound   Outbound
-----------------------------------------------
Inbound Keygroup SDP Association Count:  2
Outbound Keygroup SDP Association Count: 2
-----------------------------------------------------------------
-------------------------------------------------------------------------------
VPRN Keygroup Association Table
-------------------------------------------------------------------------------
No entries found

-------------------------------------------------------------------------------
Network Interface Association Table
-------------------------------------------------------------------------------
No entries found
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
Wlan-GW Keygroup Association Table
-------------------------------------------------------------------------------
No entries found

===============================================================================
* indicates that the corresponding row element may have been truncated.
```

Conclusion

NGE is a security solution for encrypting traffic flows on a per-service basis. The NGE helper extends the NGE solution to 7750 SR and 7950 XRS platforms where larger core and PE nodes are required to participate with other NGE-capable nodes.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Seamless BFD Application — Auto-bind tunnel

This chapter provides information about seamless BFD application — auto-bind tunnel.

Topics in this chapter include:

Applicability

This chapter was initially written based on SR OS Release 19.10.R3, but the MD-CLI in the current edition corresponds to SR OS Release 23.3.R3.

A prerequisite is to read the "Seamless BFD for SR-TE LSPs" chapter in the Segment Routing and PCE volume in the *7450 ESS, 7750 SR, and 7950 XRS MD-CLI Advanced Configuration Guide - Part I*.

Overview

Bidirectional forwarding detection (BFD) is widely deployed in IP/MPLS networks to rapidly detect failures in the forwarding path between network elements.

Seamless BFD (S-BFD) is described in RFC 7880. S-BFD minimizes the time required to establish BFD sessions by removing the discovery of discriminators during the initial handshaking procedure, which contributes to its seamless operation. S-BFD relies on the fact that the discriminators needed to establish the BFD session are already known by the endpoints for each session, either through configuration or advertisement using unicast protocols.

Figure 22: S-BFD session establishment – continuity check shows the S-BFD session establishment between PE-1 and PE-4. The BFD discriminator used by the initiator is chosen by the system. On PE-1, the BFD (initiator) discriminator equals 123; on PE-4, the S-BFD (reflector) discriminator equals 524288. Through IGP advertisement or configuration, head-end router PE-1 is aware of the S-BFD discriminator of PE-4 (system ID 192.0.2.4; S-BFD discriminator 524288).

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 22: S-BFD session establishment – continuity check

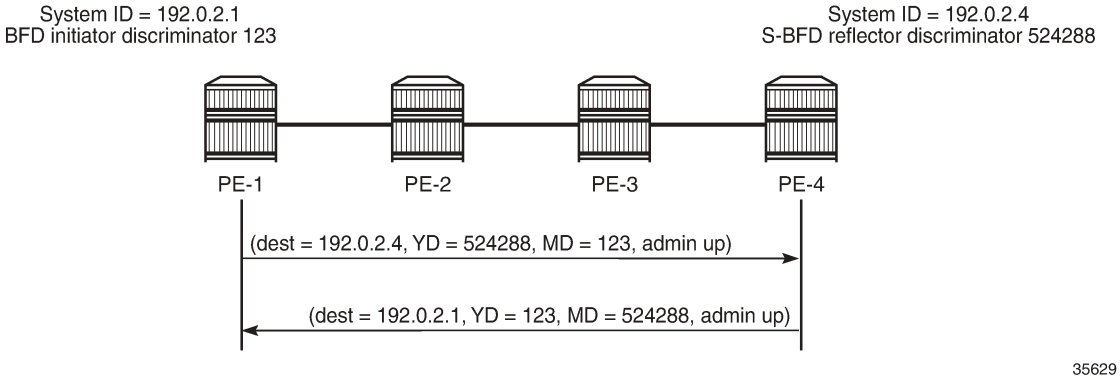

The state of the SR-TE LSP is linked to the state of the S-BFD session when failure action **failover-or-down** is configured. In the "Seamless BFD for SR-TE LSPs" chapter in the Segment Routing and PCE volume in the *7450 ESS, 7750 SR, and 7950 XRS MD-CLI Advanced Configuration Guide - Part I*, one of the examples illustrates the use of S-BFD with failure action **failover-or-down** in an SR-TE LSP with a primary path and a standby secondary path. When a link or node fails on the primary path, the S-BFD session goes down and the head-end node switches to a standby path that is operationally up.

In this chapter, S-BFD is configured in an SR-TE LSP with primary path only. Services such as VPRNs or EVPNs may have auto-bind tunnel configured with multiple tunnel resolution protocols, such as SR-TE and SR-ISIS. SR-TE tunnels are preferred to SR-ISIS tunnels. When a link or node fails on the primary path, the S-BFD session goes operationally down and the SR-TE LSP goes operationally down, and is removed from the tunnel table. The head-end node reverts to the best preference tunnel that is up; in this case, an SR-ISIS tunnel.

Configuration

Figure 23: Example topology shows the example topology. The VPRN and EVPN services will be configured on PE-2 and PE-5.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 23: Example topology

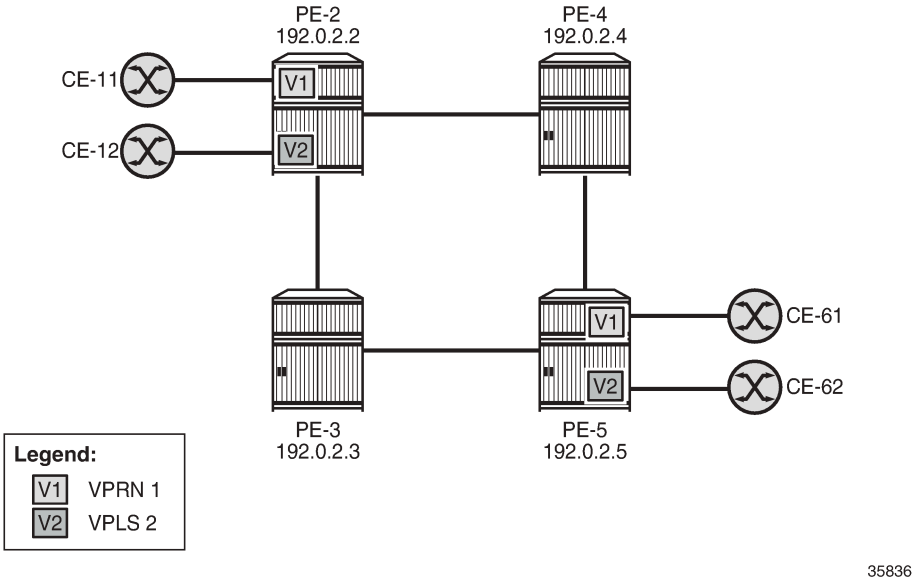

Initial configuration

The initial configuration on the PEs includes:

- Cards, MDAs, ports
- Router interfaces
- IS-IS as IGP (alternatively, OSPF can be used)
- SR-ISIS enabled
- Traffic engineering enabled on PE-2 and PE-5

The initial configuration on PE-2 is as follows:

```
# on PE-2:
configure {
    router "Base" {
        interface "int-PE-2-PE-3" {
            port 1/1/c2/1:1000
            ipv4 {
                primary {
                    address 192.168.23.1
                    prefix-length 30
                }
            }
        }
        interface "int-PE-2-PE-4" {
            port 1/1/c1/1:1000
            ipv4 {
                primary {
                    address 192.168.24.1
                    prefix-length 30
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
        interface "system" {
            ipv4 {
                primary {
                    address 192.0.2.2
                    prefix-length 32
                }
            }
        }
        mpls-labels {
            sr-labels {
                start 32000
                end 32999
            }
        }
        isis 0 {
            admin-state enable
            advertise-router-capability area
            traffic-engineering true
            area-address [49.0001]
            segment-routing {
                admin-state enable
                prefix-sid-range {
                    global
                }
            }
            interface "int-PE-2-PE-3" {
                interface-type point-to-point
            }
            interface "int-PE-2-PE-4" {
                interface-type point-to-point
            }
            interface "system" {
                ipv4-node-sid {
                    index 2
                }
            }
        }
```

S-BFD configuration

For S-BFD, the reflector BFD discriminator values must be configured in the range from 524288 to 526335. On far-end node PE-5, the global S-BFD configuration is as follows. This S-BFD discriminator will be advertised by IGP.

```
# on PE-5:
configure {
    bfd {
        seamless-bfd {
            reflector "PE-5" {
                admin-state enable
                discriminator 524291
            }
        }
```

For S-BFD, a BFD template of type CPM-NP must be configured. On PE-2, the following BFD template is configured:

```
# on PE-2:
configure {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    bfd {
        bfd-template "bfd-cpm-np-1s" {
            receive-interval 1000      # minimum value is 10 ms
            transmit-interval 1000     # minimum value is 10 ms
            type cpm-np
        }
```

Note:
Even though CPM-NP BFD can use intervals of minimum 10 ms, the used example setup has its limitations. The nodes in the used example setup are sims and the simulation for CPM-NP or central BFD sessions has the limitation that intervals that are configured with a value smaller than 1000 ms are always negotiated to intervals of 1000 ms. To avoid confusion when the configured intervals differ from the negotiated intervals on sims, a BFD template with intervals of 1000 ms is configured and used in this chapter.

On PE-2, the preceding BFD template is applied in the following SR-TE LSP to PE-5. For SR-TE LSPs, the only allowed failure action is **failover-or-down**.

```
# on PE-2:
configure {
    router "Base" {
        mpls {
            admin-state enable
            path "empty" {
                admin-state enable
            }
            lsp "LSP-PE-2-PE-5_empty_localCSPF" {
                admin-state enable
                type p2p-sr-te
                to 192.0.2.5
                path-computation-method local-cspf
                bfd {
                    bfd-liveness true
                    bfd-template "bfd-cpm-np-1s"
                    failure-action failover-or-down
                }
                primary "empty" {
                }
            }
```

The following tunnel table on PE-2 shows that two tunnels are available toward PE-5: an SR-TE tunnel with tunnel ID 655362 and default preference 8, and an SR-ISIS tunnel with tunnel ID 524293 and default preference 11. The SR-TE tunnel with preference 8 is preferred to the SR-ISIS tunnel with preference 11.

```
[/]
A:admin@PE-2# show router tunnel-table 192.0.2.5/32

===============================================================================
IPv4 Tunnel Table (Router: Base)
===============================================================================
Destination           Owner     Encap TunnelId  Pref   Nexthop        Metric
   Color
-------------------------------------------------------------------------------
192.0.2.5/32          sr-te     MPLS  655362    8      192.168.24.2   20
192.0.2.5/32          isis (0)  MPLS  524293    11     192.168.23.2   20
-------------------------------------------------------------------------------
Flags: B = BGP or MPLS backup hop available
       L = Loop-Free Alternate (LFA) hop available
       E = Inactive best-external BGP route
       k = RIB-API or Forwarding Policy backup hop
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
```

The SR-TE LSP with tunnel ID 655362 is "LSP-PE-2-PE-5_empty_localCSPF":

```
[/]
A:admin@PE-2# show router mpls sr-te-lsp detail

===============================================================================
MPLS SR-TE LSPs (Originating) (Detail)
===============================================================================
Legend :
    + - Inherited
===============================================================================
-------------------------------------------------------------------------------
Type : Originating
-------------------------------------------------------------------------------
LSP Name   : LSP-PE-2-PE-5_empty_localCSPF
LSP Type        : SrTeLsp                   LSP Tunnel ID        : 1
LSP Index       : 65536                     TTM Tunnel Id        : 655362
From            : 192.0.2.2
To              : 192.0.2.5
Adm State       : Up                        Oper State           : Up
---snip---
```

The S-BFD session for the SR-TE LSP is up, as follows:

```
[/]
A:admin@PE-2# show router bfd seamless-bfd session
                                      lsp-name "LSP-PE-2-PE-5_empty_localCSPF"

===============================================================================
Legend:
  Session Id = Interface Name | LSP Name | Prefix | RSVP Sess Name | Service Id
  wp = Working path   pp = Protecting path
===============================================================================
BFD Session
===============================================================================
Session Id                                        State      Tx Pkts    Rx Pkts
  Rem Addr/Info/SdpId:VcId                      Multipl     Tx Intvl   Rx Intvl
  Protocols                                        Type     LAG Port     LAG ID
  Loc Addr
-------------------------------------------------------------------------------
192.0.2.5/32                                         Up          N/A        N/A
  192.0.2.5                                           3         1000       1000
  mplsLsp                                        cpm-np          N/A        N/A
  192.0.2.2
-------------------------------------------------------------------------------
No. of BFD sessions: 1
===============================================================================
```

VPRN and EVPN services with auto-bind tunnel

Both VPRN "VPRN-1" and an EVPN VPLS "VPLS-2" will be configured on PE-2 and PE-5. For advertising VPN-IPv4 and EVPN routes, BGP is configured on PE-2 and PE-5 for the VPN-IPv4 and EVPN address families. Both VPRN "VPRN-1" and EVPN VPLS "VPLS-2" have auto-bind tunnel enabled with resolution filter allowing SR-ISIS and SR-TE.

```
# on PE-2:
configure {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        router "Base" {
        autonomous-system 64496
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            split-horizon true
            rapid-update {
                vpn-ipv4 true
                evpn true
            }
            group "internal" {
                peer-as 64496
                family {
                    vpn-ipv4 true
                    evpn true
                }
            }
            neighbor "192.0.2.5" {
                group "internal"
            }
        }
    }
    service {
        vpls "VPLS-2" {
            admin-state enable
            service-id 2
            customer "1"
            bgp 1 {
            }
            bgp-evpn {
                evi 2
                mpls 1 {
                    admin-state enable
                    auto-bind-tunnel {
                        resolution filter
                        resolution-filter {
                            sr-isis true
                            sr-te true
                        }
                    }
                }
            }
            sap 1/1/c3/1:2 {
            }
        }
        vprn "VPRN-1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:1"
                    vrf-target {
                        community "target:64496:1"
                    }
                    auto-bind-tunnel {
                        resolution filter
                        resolution-filter {
                            sr-isis true
                            sr-te true
                        }
                    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
        }
        interface "int-VPRN-1_PE-2_CE-11" {
            mac 00:00:5e:00:53:11
            ipv4 {
                primary {
                    address 172.31.2.2
                    prefix-length 30
                }
            }
            sap 1/1/c4/1:1 {
            }
        }
    }
```

The following route table for VPRN "VPRN-1" on PE-2 shows that the SR-TE tunnel with tunnel ID 655362 is used toward next-hop 192.0.2.5:

```
[/]
A:admin@PE-2# show router 1 route-table

===============================================================================
Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
172.31.2.0/30                                 Local   Local     00h01m53s  0
       int-VPRN-1_PE-2_CE-11                                        0
172.31.5.4/30                                 Remote  BGP VPN   00h01m39s  170
       192.0.2.5 (tunneled:SR-TE:655362)                            20
-------------------------------------------------------------------------------
No. of Routes: 2
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

Likewise, for the EVPN service, the SR-TE tunnel with tunnel ID 655362 is used toward 192.0.2.5, as follows:

```
[/]
A:admin@PE-2# show service id 2 fdb detail

===============================================================================
Forwarding Database, Service 2
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
2          00:00:5e:00:53:12 sap:1/1/c3/1:2          L/0      07/05/23 15:17:23
2          00:00:5e:00:53:62 mpls-1:                 Evpn     07/05/23 15:17:23
                             192.0.2.5:524285
           sr-te:655362
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
```

```
[/]
A:admin@PE-2# show router bgp next-hop evpn service-id 2
===============================================================================
 BGP Router ID:192.0.2.2        AS:64496       Local AS:64496
===============================================================================

===============================================================================
BGP VPN Next Hop
===============================================================================
VPN Next Hop                                                  Owner
   Autobind                                          FibProg  Reason
   Labels (User-labels)                              FlexAlgo Metric
   Admin-tag-policy (strict-tunnel-tagging)                   Last Mod.
-------------------------------------------------------------------------------
192.0.2.5                                                     SR_TE
   sr-isis sr-te                                     Y
   -- (3)                                            --       20
   -- (N)                                                     00h02m02s
-------------------------------------------------------------------------------
Next Hops : 1
===============================================================================
```

Failure of the SR-TE LSP

The following command shows that—without any failures—the primary path of the SR-TE LSP goes via PE-4:

```
[/]
A:admin@PE-2# show router mpls sr-te-lsp "LSP-PE-2-PE-5_empty_localCSPF" path detail
                                                | match "Actual Hops" post-lines 3
Actual Hops      :
    192.168.24.2(192.0.2.4)(A-SID)              Record Label        : 524286
 -> 192.168.45.2(192.0.2.5)(A-SID)              Record Label        : 524286
```

Figure 24: Primary path of SR-TE LSP via PE-4 shows the primary path of the SR-TE LSP.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 24: Primary path of SR-TE LSP via PE-4

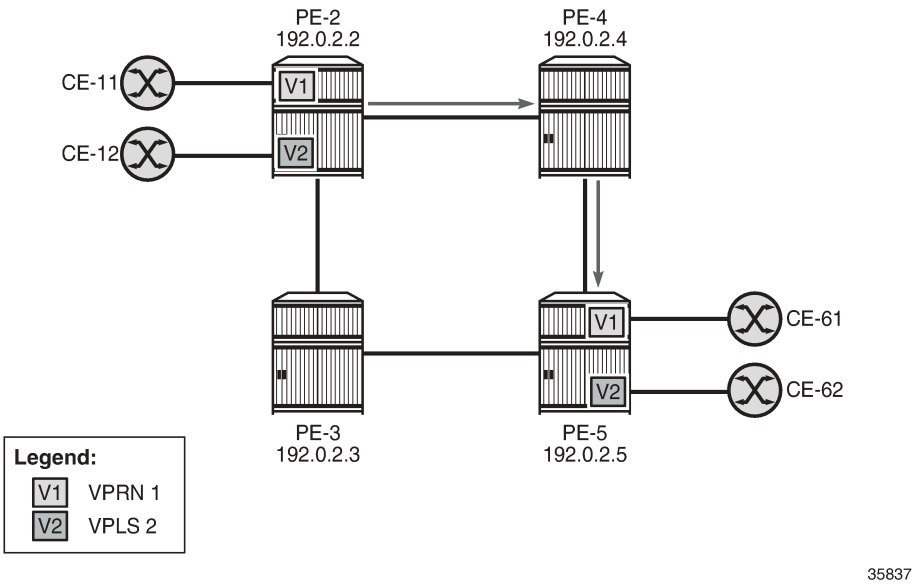

S-BFD is configured in the SR-TE LSP with failure action **failover-or-down**. If the SR-TE LSP fails, the S-BFD session will go down and it will bring the SR-TE tunnel down. The next-hop 192.0.2.5 cannot be resolved using the SR-TE tunnel, so an SR-ISIS tunnel will be used instead.

On PE-4, port 1/1/c1/1 to PE-5 is disabled to emulate a failure in the primary path of the SR-TE LSP, as follows:

```
# on PE-4:
configure {
    port 1/1/c1/1 {     # port to PE-5
        admin-state disable
```

Figure 25: Remote failure in the primary path of the SR-TE LSP shows that a remote failure occurs in the primary path of the SR-TE LSP.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 25: Remote failure in the primary path of the SR-TE LSP

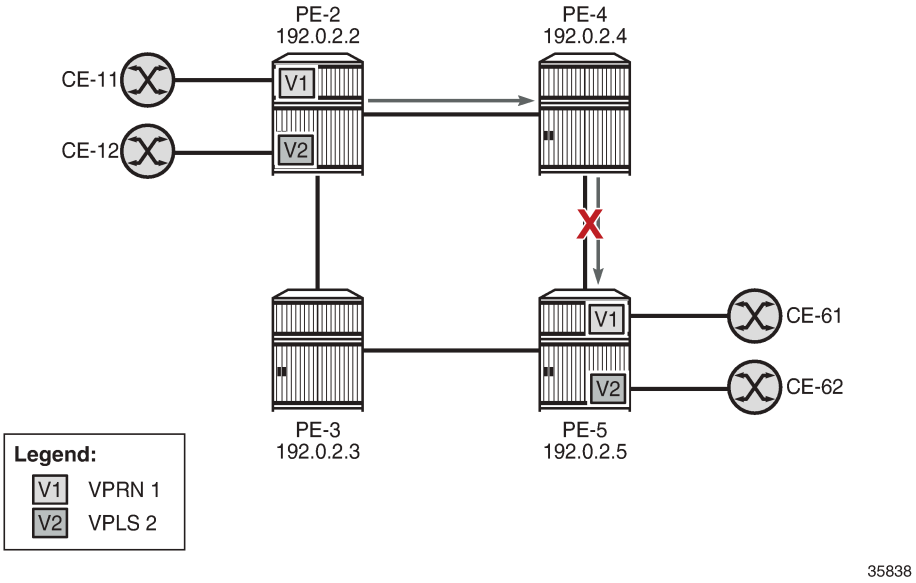

The S-BFD session goes operationally down, as follows:

```
[/]
A:admin@PE-2# show router bfd seamless-bfd session lsp-path detail prefix 192.0.2.5/32

===============================================================================
BFD Session
===============================================================================
Prefix         : 192.0.2.5/32
Local Address  : 192.0.2.2
LSP Name       : LSP-PE-2-PE-5_empty_localCSPF
LSP Index      : 65536                    Path LSP ID      : 4096
Fec Type       : srTe
Oper State     : Down                     Protocols        : mplsLsp
Last Up Time   : 0d 00:04:14              Up Transitions   : 1
Down Time      : 0d 00:00:01              Down Transitions : 1
                                          Version Mismatch : 0

Forwarding Information

Local Discr    : 1                        Local State      : Down
Local Diag     : 1 (Detect time expired)
Local Mode     : Demand
Local Min Tx   : 1000                     Local Mult       : 3
Last Sent (ms) : 0                        Local Min Rx     : 0
Type           : cpm-np
Remote         : Unheard                  Remote Discr     : 524291
===============================================================================
===============================================================================
```

When the S-BFD session goes down, the SR-TE LSP goes operationally down, as follows:

```
[/]
A:admin@PE-2# show router mpls sr-te-lsp

===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
MPLS SR-TE LSPs (Originating)
================================================================================
LSP Name                                          Tun     Protect   Adm  Opr
  To                                              Id      Path
--------------------------------------------------------------------------------
LSP-PE-2-PE-5_empty_localCSPF                     1       N/A       Up   Dwn
  192.0.2.5
--------------------------------------------------------------------------------
LSPs : 1
================================================================================
```

Because the SR-TE tunnel is operationally down, the only available tunnel to 192.0.2.5 is the SR-ISIS tunnel, as follows:

```
[/]
A:admin@PE-2# show router tunnel-table 192.0.2.5/32

================================================================================
IPv4 Tunnel Table (Router: Base)
================================================================================
Destination           Owner     Encap TunnelId  Pref   Nexthop        Metric
   Color
--------------------------------------------------------------------------------
192.0.2.5/32          isis (0)  MPLS  524293    11     192.168.23.2   20
--------------------------------------------------------------------------------
Flags: B = BGP or MPLS backup hop available
       L = Loop-Free Alternate (LFA) hop available
       E = Inactive best-external BGP route
       k = RIB-API or Forwarding Policy backup hop
================================================================================
```

The route table for VPRN "VPRN-1" shows that an SR-ISIS tunnel is used toward next-hop 192.0.2.5:

```
[/]
A:admin@PE-2# show router 1 route-table

================================================================================
Route Table (Service: 1)
================================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
--------------------------------------------------------------------------------
172.31.2.0/30                                   Local   Local     00h03m17s  0
       int-VPRN-1_PE-2_CE-11                                          0
172.31.5.4/30                                   Remote  BGP VPN   00h00m12s  170
       192.0.2.5 (tunneled:SR-ISIS:524293)                            20
--------------------------------------------------------------------------------
No. of Routes: 2
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
================================================================================
```

Likewise, the FDB for the EVPN VPLS "VPLS-2" shows that an SR-ISIS tunnel with tunnel ID 524293 is used toward next-hop 192.0.2.5:

```
[/]
A:admin@PE-2# show service id 2 fdb detail

================================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Forwarding Database, Service 2
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
2          00:00:5e:00:53:12 sap:1/1/c3/1:2          L/60     07/05/23 15:17:23
2          00:00:5e:00:53:62 mpls-1:                 Evpn     07/05/23 15:17:23
                             192.0.2.5:524285
           isis:524293
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

SR-TE LSP reconnects after retry timer expires

When the SR-TE LSP retry timer expires, the primary path is recalculated and it will go via PE-3 (192.0.2.3), as follows:

```
[/]
A:admin@PE-2# show router mpls sr-te-lsp "LSP-PE-2-PE-5_empty_localCSPF" path detail
                                                  | match "Actual Hops" post-lines 3
Actual Hops      :
    192.168.23.2(192.0.2.3)(A-SID)                Record Label        : 524287
 -> 192.168.35.2(192.0.2.5)(A-SID)                Record Label        : 524286
```

Figure 26: SR-TE LSP reconnects after retry timer expires show that the primary path of the SR-TE tunnel goes via PE-3.

Figure 26: SR-TE LSP reconnects after retry timer expires

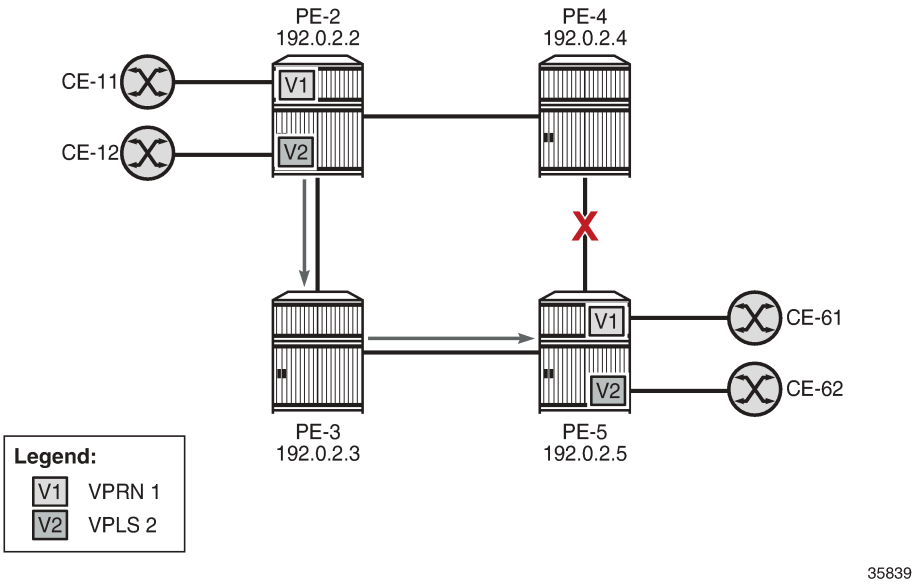

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The tunnel table shows two tunnels to 192.0.2.5: one SR-TE tunnel with tunnel ID 655362 and one SR-ISIS tunnel with tunnel ID 524293:

```
[/]
A:admin@PE-2# show router tunnel-table 192.0.2.5/32

===============================================================================
IPv4 Tunnel Table (Router: Base)
===============================================================================
Destination           Owner     Encap TunnelId  Pref   Nexthop        Metric
   Color
-------------------------------------------------------------------------------
192.0.2.5/32          sr-te     MPLS  655362    8      192.168.23.2   20
192.0.2.5/32          isis (0)  MPLS  524293    11     192.168.23.2   20
-------------------------------------------------------------------------------
Flags: B = BGP or MPLS backup hop available
       L = Loop-Free Alternate (LFA) hop available
       E = Inactive best-external BGP route
       k = RIB-API or Forwarding Policy backup hop
===============================================================================
```

Again, the SR-TE LSP will be preferred to the SR-ISIS LSP and both VPRN "VPRN-1" and EVPN VPLS "VPLS-2" will use the SR-TE tunnel to 192.0.2.5.

Conclusion

S-BFD can be used to determine the state of SR-TE LSPs that only have a primary path. The resiliency is at the service level for VPRN and EVPN services with auto-bind tunnel where several resolution protocols are configured and SR-TE has the lowest preference. When the S-BFD session for the SR-TE tunnel goes operationally down, the SR-TE tunnel goes operationally down. The VPRN and EVPN services will then use the best tunnel that is available; in this example, an SR-ISIS tunnel.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Layer 2 Services and EVPN

This section provides configuration information for the following topics:

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- LDP VPLS Using BGP Auto-Discovery
- LDP VPLS Using BGP Auto-Discovery — Prefer Provisioned SDP
- Mobility for EVPN Hosts Within an R-VPLS
- Operational Groups for EVPN-VXLAN VPWS Services
- Operational Groups in EVPN Services
- P2MP mLDP FEC Resolution for BGP-LU in EVPN
- P2MP mLDP Inter-AS Model C for EVPN-MPLS Services
- P2MP mLDP Tunnels for BUM Traffic in EVPN-MPLS Services
- PBB-Epipe
- PBB-EVPN ISID-based CMAC Flush
- PBB-EVPN ISID-based Route Targets
- PBB-VPLS
- PIM Snooping for IPv4 in EVPN-MPLS Services
- PIM Snooping for IPv4 in PBB-EVPN Services
- Preference-based and Non-revertive EVPN DF Election
- Proxy-ARP/ND MAC List for Dynamic Entries
- Static VXLAN Termination in Epipe Services
- Three-byte EVI in EVPN Services
- VCCV BFD for Epipe Services
- Virtual Ethernet Segments
- VLAN Range SAPs for VPLS and Epipe Services
- VXLAN Forwarding Path Extension

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

AC-Influenced DF Election on an ES

This chapter provides information about Attachment Circuit (AC) influenced Designated Forwarder (DF) election on an Ethernet Segment (ES).

Topics in this chapter include:

Applicability

The information and configuration in this chapter are based on SR OS Release 22.5.R1. Attachment Circuit (AC) influenced Designated Forwarder (DF) election on an Ethernet Segment (ES) is always enabled in SR OS releases earlier than 21.5.R1. The AC-DF election capability can be disabled in SR OS Release 21.5.R1 and later.

Overview

RFC 8584, section "The AC-Influenced DF Election Capability", describes the AC-DF capability that modifies the EVPN DF election process in RFC 7432. RFC 8584 states that when PEs build their candidate DF election list, they do not include PEs when no Auto-Discovery (AD) per-ES or per-EVI routes for those PEs are present. In SR OS, this behavior is default for all ESs, configured as **ac-df-capability include**.

The **ac-df-capability** command is configurable in the **configure service system bgp evpn ethernet-segment** context:

```
[ex:/configure service system bgp evpn ethernet-segment "SA-ESI-23"]
A:admin@PE-2# ac-df-capability ?

 ac-df-capability <keyword>
 <keyword>  - (include|exclude)
 Default    - include

    AC-influenced DF election capability

    Warning: Modifying this element toggles
    'configure service system bgp evpn ethernet-segment "SA-ESI-23" admin-state'
    automatically for the new value to take effect.
```

The command **ac-df-capability exclude** disables AC-DF on the ES, so the presence of an AD per-ES or per-EVI does not influence the candidate DF election list. When **ac-df-capability exclude** is configured:

- The candidate DF election list is not influenced by the presence or absence of AD per-ES/EVI routes (type 1) from the ES peers.
- PEs are only removed from the candidate DF election list when their ES route (type 4) is not present.
- The local ES route is active if there are active SAPs on the ES.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- When the local AC is operationally down, due to **admin-state disable** or reason other than Multi Homing (MH) standby, this does not trigger a DF switchover.

The **ac-df-capability exclude** option:

- is supported with any type of service-carving (DF Election)
- is recommended in ESs that use an operational group monitored by the access LAG to signal standby LACP or power-off
- must be configured consistently on all PEs attached to the same ES

AC-DF enabled – default

The following example illustrates the default behavior, where a PE builds the list of DF candidates with nodes that have sent EVPN AD per-ES/EVI routes. This behavior is compatible with the behavior in SR OS releases earlier than 21.5.R1.

Figure 27: PE-4 as the DF on a single-active ES for three VPLSs shows a topology with MTU-6 connected via SDPs to the single-active ES "SA-vESI-45". PE-4 is the DF for three services: VPLS 1, VPLS 2, and VPLS 3. Traffic for these services passes via PE-4, while PE-5 is standby.

Figure 27: PE-4 as the DF on a single-active ES for three VPLSs

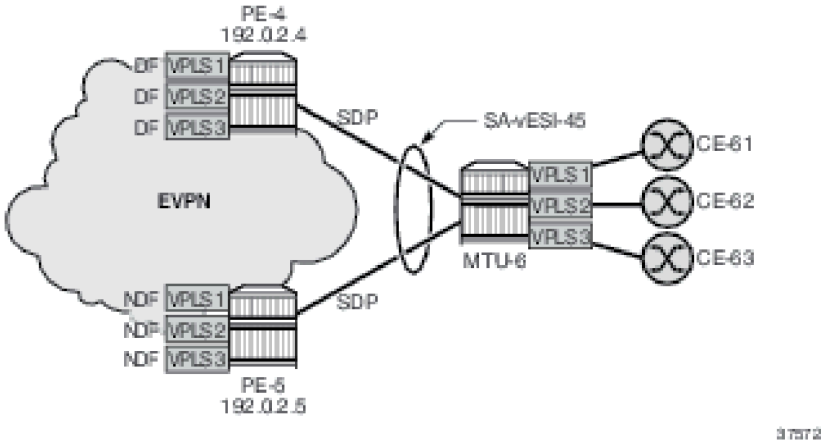

When a failure occurs on the spoke-SDP in VPLS 2 on PE-4, PE-4 sends an EVPN-AD per-EVI withdrawal and PE-4 becomes the Non-Designated Forwarder (NDF) for VPLS 2, while remaining the DF for VPLS 1 and VPLS 3, as shown in Figure 28: AC failure in VPLS 2 on PE-4 causes PE-5 to become the DF for VPLS 2.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 28: AC failure in VPLS 2 on PE-4 causes PE-5 to become the DF for VPLS 2

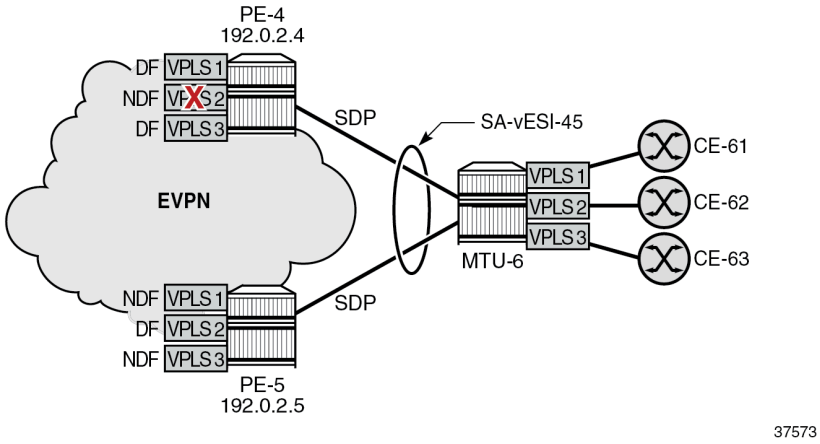

VPLS 2 traffic to and from MTU-6 passes via DF PE-5, while VPLS 1 and VPLS 3 traffic will pass via DF PE-4. No traffic is dropped. The AC failure in VPLS 2 does not have an impact on the other services.

Problem with AC-DF on ES with the operational group monitored by LAG

In this example, a failure in an access circuit of a particular service also impacts other services when the AC-DF capability is enabled.

Figure 29: PE-2 is DF on single-active ES for three VPLSs shows a single-active ES with LAG 1 associated with it. An operational group is assigned to the ES and monitored by the LAG to signal standby LACP (default) or power off. Three VPLSs are configured on PE-2 and PE-3. PE-2 is the DF for each of these VPLSs.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 29: PE-2 is DF on single-active ES for three VPLSs

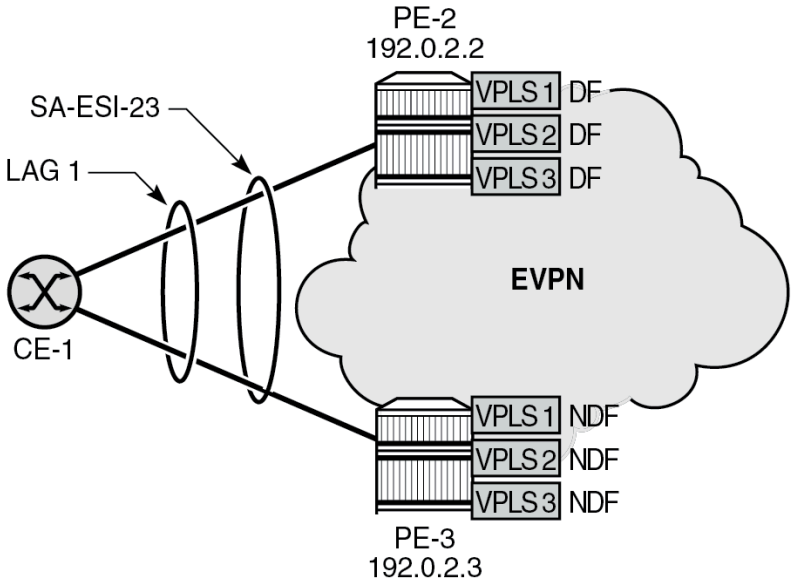

On NDF PE-3, the ES is inactive which causes the operational group in the ES to go down. LAG 1 monitors this operational group, so the LAG goes standby on NDF PE-3. LAG 1 has LACP standby-signaling enabled (default). On CE-1, only the LAG port to DF PE-2 is up and all traffic for the VPLSs goes via PE-2.

When the single-active ES has the default AC-DF setting (**ac-df-capability include**), a failure (or an unintended **admin-state disable**) on SAP lag-1:2 in VPLS 2 (or on the VPLS 2 service) on PE-2 can have an impact on all three services that share LAG 1. Figure 30: AC failure in VPLS 2 on PE-2 causes PE-3 to become DF for VPLS 2 shows that such an AC failure in VPLS 2 on PE-2 causes PE-3 to become the DF for VPLS 2 (after receiving an AD per-EVI withdrawal from PE-2).

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 30: AC failure in VPLS 2 on PE-2 causes PE-3 to become DF for VPLS 2

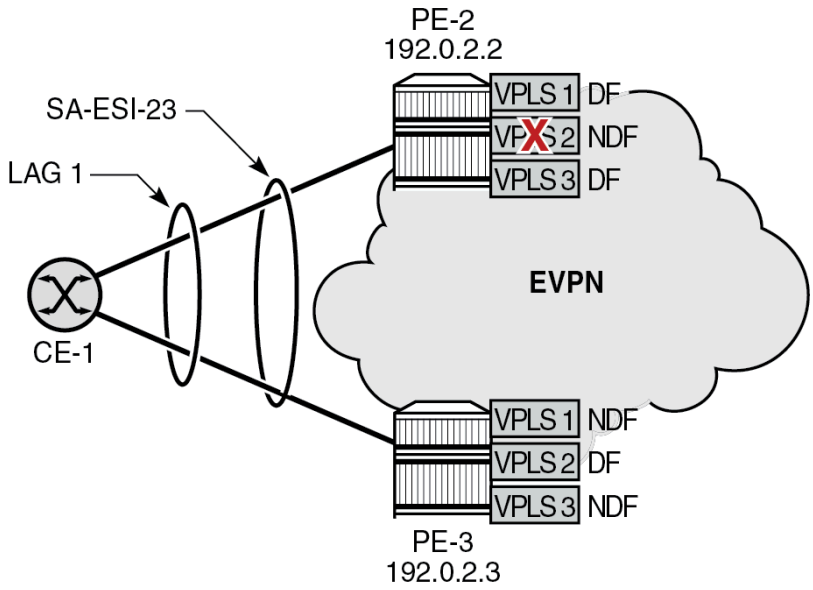

When PE-3 is the DF for VPLS 2, the ES operational group on PE-3 goes up. Therefore, the monitoring LAG is up on PE-3. On CE-1, both LAG ports to PE-2 and PE-3 are up. CE-1 can now send all VPLS traffic via either LAG port: DF PE-2 forwards the VPLS 1 and VPLS 3 traffic whereas NDF PE-3 drops it. PE-3 accepts VPLS 2 traffic, but PE-2 drops it. Approximately 50% of the traffic is lost.

AC-DF capability disabled

Nokia recommends disabling the AC-DF capability in ESs where the operational group is monitored by the LAG. Figure 31: AC failure in VPLS 2 on PE-2 has no impact on DF election shows the situation with the AC-DF disabled (**ac-df-capability exclude**): the PEs ignore the AD per-EVI withdrawal and PE-2 remains the DF for VPLS 2.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 31: AC failure in VPLS 2 on PE-2 has no impact on DF election

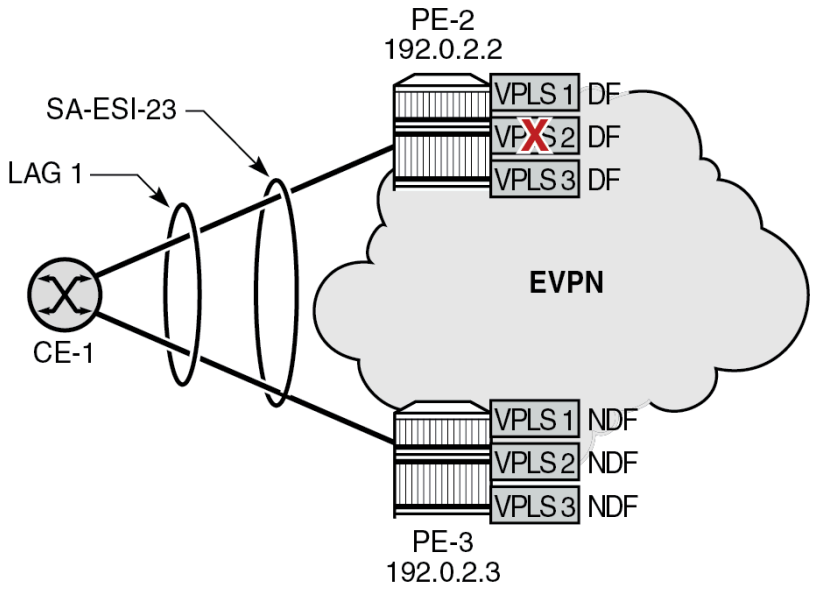

VPLS 2 traffic is dropped by PE-2, but the other services are not impacted.

Configuration

Figure 32: Example topology shows the example topology with four PEs in an EVPN-MPLS network.

Figure 32: Example topology

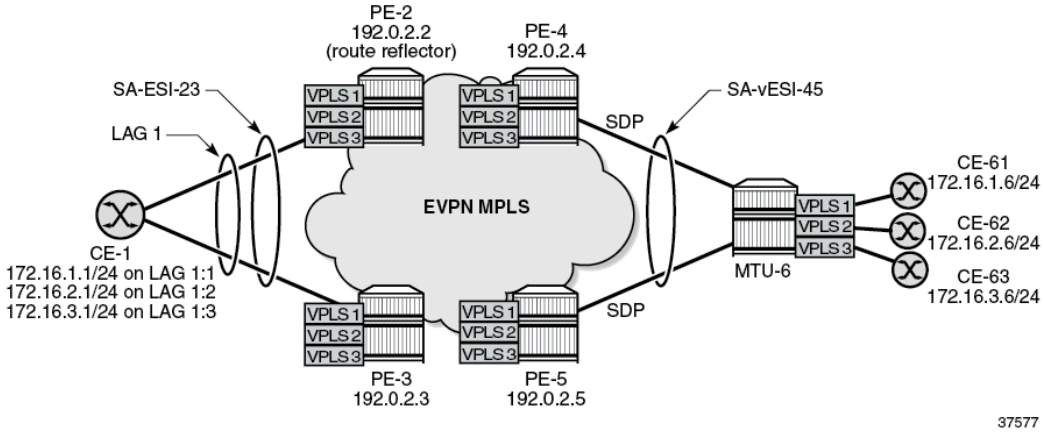

The initial configuration includes:

- cards, MDAs, ports
- router interfaces on the PEs and on MTU-6

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- IS-IS on the router interfaces (alternatively, OSPF can be configured)
- LDP on the router interfaces

On the PEs, BGP is configured for the EVPN address family. In this example, PE-2 is the Route Reflector (RR) with the following BGP configuration:

```
# on PE-2:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            rapid-update {
                evpn true
            }
            group "internal" {
                peer-as 64500
                family {
                    evpn true
                }
                cluster {
                    cluster-id 192.0.2.2
                }
            }
            neighbor "192.0.2.3" {
                group "internal"
            }
            neighbor "192.0.2.4" {
                group "internal"
            }
            neighbor "192.0.2.5" {
                group "internal"
            }
        }
```

The BGP configuration on the clients PE-3, PE-4, and PE-5 is as follows:

```
# on PE-3, PE-4, PE-5:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            rapid-update {
                evpn true
            }
            group "internal" {
                peer-as 64500
                family {
                    evpn true
                }
            }
            neighbor "192.0.2.2" {
                group "internal"
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
```

AC-DF capability enabled – default

On PE-2 and PE-3, operational group "op-grp-sa-es-23" is configured. This operational group is assigned to the single-active ES "SA-ESI-23" and monitored on LAG 1.

On PE-2, LAG 1 is configured as follows. The LAG configuration on PE-3 is similar, but with port 1/1/1 instead.

```
# on PE-2:
configure {
    lag "lag-1" {
        admin-state enable
        encap-type dot1q
        mode access
        # standby-signaling lacp      # default
        monitor-oper-group "op-grp-sa-es-23"
        max-ports 64
        lacp {
            mode active
            system-id 00:00:00:00:23:01
            administrative-key 1
        }
        port 1/1/2 {
        }
    }
```

On PE-2 and PE-3, three VPLS services are configured with SAPs from LAG 1, which is associated with single-active ES "SA-ESI-23". This ES is configured with the operational group "op-grp-sa-es-23" that is monitored by LAG 1. The operational group triggers the LACP standby signaling from the NDF PE to CE-1 to avoid attracting traffic.

The service configuration on PE-2 and PE-3 is similar; only the preference value for the service carving in the ES is different.

Note:
When an operational group is associated with an ES, the hold timers for the operational group must be zero (the default value).

```
# on PE-2:
configure {
    service {
        oper-group "op-grp-sa-es-23" {
            hold-time {
             ## down    # default 0
                up 0
            }
        }
        system {
            bgp {
                evpn {
                    ethernet-segment "SA-ESI-23" {
                        admin-state enable
                        esi 01:00:00:00:00:23:01:00:00:01
                        multi-homing-mode single-active
                        oper-group "op-grp-sa-es-23"
                        # ac-df-capability include       # default
                        df-election {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        service-carving-mode manual
                        manual {
                            preference {
                                mode non-revertive
                                value 200         # on PE-3: preference value 100
                            }
                        }
                    }
                    association {
                        lag "lag-1" {
                        }
                    }
                }
            }
        }
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
            }
            bgp-evpn {
                evi 1
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    ecmp 2
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            sap lag-1:1 {
            }
        }
        vpls "VPLS 2" {
            admin-state enable
            service-id 2
            customer "1"
            bgp 1 {
            }
            bgp-evpn {
                evi 2
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    ecmp 2
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            sap lag-1:2 {
            }
        }
        vpls "VPLS 3" {
            admin-state enable
            service-id 3
            customer "1"
            bgp 1 {
            }
            bgp-evpn {
                evi 3
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    ecmp 2
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            sap lag-1:3 {
            }
        }
```

On PE-4 and PE-5, single-active virtual ES "SA-vESI-45" is configured. No operational group is configured here. The service configuration on PE-4 is as follows. The configuration on PE-5 is similar, but with a different SDP and a different preference value for service carving.

```
# on PE-4:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "SA-vESI-45" {
                        admin-state enable
                        type virtual
                        esi 0x01000000004501000001
                        multi-homing-mode single-active
                        # ac-df-capability include      # default
                        df-election {
                            service-carving-mode manual
                            manual {
                                preference {
                                    value 200        # on PE-5: value 100
                                }
                            }
                        }
                        association {
                            sdp 46 {
                                virtual-ranges {
                                    vc-id 1 {
                                        end 3
                                    }
                                }
                            }
                        }
                    }
                }
            }
        }
        sdp 46 {                # on PE-5: sdp 56
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.6
            }
        }
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
            bgp-evpn {
                evi 1
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    ecmp 2
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            spoke-sdp 46:1 {            # on PE-5: spoke-sdp 56:1
            }
        }
        vpls "VPLS 2" {
            admin-state enable
            service-id 2
            customer "1"
            bgp 1 {
            }
            bgp-evpn {
                evi 2
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    ecmp 2
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            spoke-sdp 46:2 {            # on PE-5: spoke-sdp 56:2
            }
        }
        vpls "VPLS 3" {
            admin-state enable
            service-id 3
            customer "1"
            bgp 1 {
            }
            bgp-evpn {
                evi 3
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    ecmp 2
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            spoke-sdp 46:3 {            # on PE-5: spoke-sdp 56:3
            }
        }
```

With the AC-DF capability enabled (default), the PEs send ES routes with **AC:1** in the extended community for DF election. The following ES route is received by PE-3 from PE-2:

```
10 2022/06/08 15:38:15.005 CEST MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Received BGP UPDATE:
    Withdrawn Length = 0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    Total Path Attr Length = 71
    Flag: 0x90 Type: 14 Len: 34 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
        Type: EVPN-ETH-SEG Len: 23 RD: 192.0.2.2:0 ESI: 01:00:00:00:00:23:01:00:00:01, IP-Len:
 4 Orig-IP-Addr: 192.0.2.2
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        df-election::DF-Type:Preference/DP:1/DF-Preference:200/AC:1
        target:00:00:00:00:23:01
```

The remainder of the chapter focuses on PE-2 and PE-3, where an AC failure in one of the services can have an impact on the other services using the same LAG.

DF election

PE-2 is the highest-preference PE in the ES and becomes the DF (preference value 200 on PE-2 versus preference value 100 on PE-3). In case of equal preference value between PE-2 and PE-3, the Don't Preempt (DP) bit is the tiebreaker (DP = 1 for non-revertive wins over DP = 0); if that is also a tie, the lowest PE IP address is the tiebreaker.

The following command shows that PE-2 is the DF for all three VPLSs. The candidate list contains both PE-2 and PE-3 for each of these VPLSs.

```
[/]
A:admin@PE-2# show service system bgp-evpn ethernet-segment name "SA-ESI-23" all

===============================================================================
Service Ethernet Segment
===============================================================================
Name                    : SA-ESI-23
Eth Seg Type            : None
Admin State             : Enabled            Oper State         : Up
ESI                     : 01:00:00:00:00:23:01:00:00:01
Oper ESI                : 01:00:00:00:00:23:01:00:00:01
Auto-ESI Type           : None
AC DF Capability        : Include
Multi-homing            : singleActive       Oper Multi-homing  : singleActive
ES SHG Label            : 524276
Source BMAC LSB         : None
Lag                     : lag-1
ES Activation Timer     : 3 secs (default)
Oper Group              : op-grp-sa-es-23
Svc Carving             : manual             Oper Svc Carving   : manual
Cfg Range Type          : lowest-pref

-------------------------------------------------------------------------------
DF Pref Election Information
-------------------------------------------------------------------------------
Preference     Preference     Last Admin Change        Oper Pref      Do No
Mode           Value                                   Value          Preempt
-------------------------------------------------------------------------------
non-revertive  200            06/08/2022 15:38:15      200            Enabled
-------------------------------------------------------------------------------
EVI Ranges: <none>
ISID Ranges: <none>
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
EVI Information
===============================================================================
EVI                 SvcId               Actv Timer Rem      DF
-------------------------------------------------------------------------------
1                   1                   0                   yes
2                   2                   0                   yes
3                   3                   0                   yes
-------------------------------------------------------------------------------
Number of entries: 3
===============================================================================

-------------------------------------------------------------------------------
DF Candidate list
-------------------------------------------------------------------------------
EVI                                     DF Address
-------------------------------------------------------------------------------
1                                       192.0.2.2
1                                       192.0.2.3
2                                       192.0.2.2
2                                       192.0.2.3
3                                       192.0.2.2
3                                       192.0.2.3
-------------------------------------------------------------------------------
Number of entries: 6
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
---snip---
```

The same command on PE-3 shows that PE-3 is NDF for the three VPLSs and the DF candidate list is identical to the one on PE-2:

```
[/]
A:admin@PE-3# show service system bgp-evpn ethernet-segment name "SA-ESI-23" all

===============================================================================
Service Ethernet Segment
===============================================================================
Name                    : SA-ESI-23
Eth Seg Type            : None
Admin State             : Enabled            Oper State         : Up
ESI                     : 01:00:00:00:00:23:01:00:00:01
Oper ESI                : 01:00:00:00:00:23:01:00:00:01
Auto-ESI Type           : None
AC DF Capability        : Include
Multi-homing            : singleActive       Oper Multi-homing  : singleActive
ES SHG Label            : 524276
Source BMAC LSB         : None
Lag                     : lag-1
ES Activation Timer     : 3 secs (default)
Oper Group              : op-grp-sa-es-23
Svc Carving             : manual             Oper Svc Carving   : manual
Cfg Range Type          : lowest-pref

-------------------------------------------------------------------------------
DF Pref Election Information
-------------------------------------------------------------------------------
Preference     Preference     Last Admin Change        Oper Pref      Do No
Mode           Value                                   Value          Preempt
-------------------------------------------------------------------------------
non-revertive  100            06/08/2022 15:38:44      100            Enabled
-------------------------------------------------------------------------------
EVI Ranges: <none>
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
ISID Ranges: <none>
===============================================================================

===============================================================================
EVI Information
===============================================================================
EVI                 SvcId               Actv Timer Rem      DF
-------------------------------------------------------------------------------
1                   1                   0                   no
2                   2                   0                   no
3                   3                   0                   no
-------------------------------------------------------------------------------
Number of entries: 3
===============================================================================

-------------------------------------------------------------------------------
DF Candidate list
-------------------------------------------------------------------------------
EVI                                     DF Address
-------------------------------------------------------------------------------
1                                       192.0.2.2
1                                       192.0.2.3
2                                       192.0.2.2
2                                       192.0.2.3
3                                       192.0.2.2
3                                       192.0.2.3
-------------------------------------------------------------------------------
Number of entries: 6
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
---snip---
```

Operational group status

PE-2 is the DF, so the ES "SA-ESI-23" is active, the operational group "op-grp-sa-es-23" is operationally up, and the monitoring LAG 1 is operationally up.

```
[/]
A:admin@PE-2# show service oper-group "op-grp-sa-es-23" detail

===============================================================================
Service Oper Group Information
===============================================================================
Oper Group       : op-grp-sa-es-23
Creation Origin  : manual                          Oper Status: up
Hold DownTime    : 0 secs                          Hold UpTime: 0 secs
Members          : 1                               Monitoring : 1
===============================================================================

===============================================================================
Member Ethernet-Segment for OperGroup: op-grp-sa-es-23
===============================================================================
Ethernet-Segment                        Status
-------------------------------------------------------------------------------
SA-ESI-23                               Active
-------------------------------------------------------------------------------
Ethernet-Segment Entries found: 1
===============================================================================

===============================================================================
Monitoring LAG for OperGroup: op-grp-sa-es-23
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Lag-id         Adm       Opr       Weighted  Threshold Up-Count  Act/Stdby
    name
-------------------------------------------------------------------------------
1              up        up        No        0         1         N/A
    lag-1
-------------------------------------------------------------------------------
LAG Entries found: 1
===============================================================================
port option not supported with monitoring
```

PE-3 is NDF, so the ES "SA-ESI-23" is inactive, the operational group "op-grp-sa-es-23" is operationally down, and the monitoring LAG 1 is operationally down:

```
[/]
A:admin@PE-3# show service oper-group "op-grp-sa-es-23" detail

===============================================================================
Service Oper Group Information
===============================================================================
Oper Group       : op-grp-sa-es-23
Creation Origin  : manual                          Oper Status: down
Hold DownTime    : 0 secs                          Hold UpTime: 0 secs
Members          : 1                               Monitoring : 1
===============================================================================

===============================================================================
Member Ethernet-Segment for OperGroup: op-grp-sa-es-23
===============================================================================
Ethernet-Segment                    Status
-------------------------------------------------------------------------------
SA-ESI-23                           Inactive
-------------------------------------------------------------------------------
Ethernet-Segment Entries found: 1
===============================================================================

===============================================================================
Monitoring LAG for OperGroup: op-grp-sa-es-23
===============================================================================
Lag-id         Adm       Opr       Weighted  Threshold Up-Count  Act/Stdby
    name
-------------------------------------------------------------------------------
1              up        down      No        0         0         N/A
    lag-1
-------------------------------------------------------------------------------
LAG Entries found: 1
===============================================================================
port option not supported with monitoring
```

LAG port status

On DF PE-2, LAG port 1/1/2 toward CE-1 is operationally up:

```
[/]
A:admin@PE-2# show lag "lag-1" port

===============================================================================
Lag Port States
LACP Status: e - Enabled, d - Disabled
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Name
Id     Port-id        Adm   Act/     Opr   Primary Sub-group      Forced Prio
                            Stdby
------------------------------------------------------------------------------
lag-1
1(e)   1/1/2          up    active   up    yes     1              -      32768
==============================================================================
```

On NDF PE-3, LAG port 1/1/1 toward CE-1 is operationally down:

```
[/]
A:admin@PE-3# show lag "lag-1" port

==============================================================================
Lag Port States
LACP Status: e - Enabled, d - Disabled
==============================================================================
Name
Id     Port-id        Adm   Act/     Opr   Primary Sub-group      Forced Prio
                            Stdby
------------------------------------------------------------------------------
lag-1
1(e)   1/1/1          up    active   down  yes     1              -      32768
==============================================================================
```

On CE-1, LAG port 1/1/1 toward DF PE-2 is operationally up while LAG port 1/1/2 toward NDF PE-3 is down:

```
[/]
A:admin@CE-1# show lag "lag-1" port

==============================================================================
Lag Port States
LACP Status: e - Enabled, d - Disabled
==============================================================================
Name
Id     Port-id        Adm   Act/     Opr   Primary Sub-group      Forced Prio
                            Stdby
------------------------------------------------------------------------------
lag-1
1(e)   1/1/1          up    active   up    yes     1              -      32768
       1/1/2          up    active   down          1              -      32768
==============================================================================
```

AD per-EVI route withdrawal

A failure is simulated by disabling SAP lag-1:2 in VPLS 2 on PE-2:

```
# on PE-2:
configure {
    service {
        vpls "VPLS 2" {
            sap lag-1:2 {
                admin-state disable
```

PE-2 withdraws the EVPN-AD per-EVI route. The following withdrawal is received by PE-3:

```
77 2022/06/08 15:44:59.536 CEST MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 69
    Flag: 0x90 Type: 15 Len: 65 Multiprotocol Unreachable NLRI:
        Address Family EVPN
        Type: EVPN-MAC Len: 33 RD: 192.0.2.2:2 ESI: ESI-0, tag: 0, mac len: 48 mac:
 00:00:00:00:02:01, IP len: 0, IP: NULL, label1: 0
        Type: EVPN-AD Len: 25 RD: 192.0.2.2:2 ESI: 01:00:00:00:00:23:01:00:00:01, tag: 0 Label:
 0 (Raw Label: 0x0) PathId:
"
```

The following command on PE-3 shows that the list of DF candidates no longer includes PE-2 in the DF candidate list for VPLS 2 and that PE-3 is the DF for VPLS 2, while remaining the NDF for VPLS 1 and VPLS 3.

```
[/]
A:admin@PE-3# show service system bgp-evpn ethernet-segment name "SA-ESI-23" all | match "EVI
 Information" pre-lines 2 post-lines 24
===============================================================================
===============================================================================
EVI Information
===============================================================================
EVI                 SvcId               Actv Timer Rem      DF
-------------------------------------------------------------------------------
1                   1                   0                   no
2                   2                   0                   yes
3                   3                   0                   no
-------------------------------------------------------------------------------
Number of entries: 3
===============================================================================
-------------------------------------------------------------------------------
DF Candidate list
-------------------------------------------------------------------------------
EVI                                     DF Address
-------------------------------------------------------------------------------
1                                       192.0.2.2
1                                       192.0.2.3
2                                       192.0.2.3
3                                       192.0.2.2
3                                       192.0.2.3
-------------------------------------------------------------------------------
Number of entries: 5
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
===============================================================================
```

When PE-3 becomes the DF for one of the services, the ES "SA-ESI-23" is active and the operational group "op-grp-sa-es-23" and LAG 1 are up, as follows:

```
[/]
A:admin@PE-3# show service oper-group "op-grp-sa-es-23" detail

===============================================================================
Service Oper Group Information
===============================================================================
Oper Group       : op-grp-sa-es-23
Creation Origin  : manual                          Oper Status: up
Hold DownTime    : 0 secs                          Hold UpTime: 0 secs
Members          : 1                               Monitoring : 1
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Member Ethernet-Segment for OperGroup: op-grp-sa-es-23
===============================================================================
Ethernet-Segment                        Status
-------------------------------------------------------------------------------
SA-ESI-23                               Active
-------------------------------------------------------------------------------
Ethernet-Segment Entries found: 1
===============================================================================

===============================================================================
Monitoring LAG for OperGroup: op-grp-sa-es-23
===============================================================================
Lag-id         Adm       Opr       Weighted  Threshold Up-Count  Act/Stdby
    name
-------------------------------------------------------------------------------
1              up        up        No        0         1         N/A
    lag-1
-------------------------------------------------------------------------------
LAG Entries found: 1
===============================================================================
port option not supported with monitoring
```

On PE-3, LAG port 1/1/1 toward CE-1 is up:

```
[/]
A:admin@PE-3# show lag "lag-1" port

===============================================================================
Lag Port States
LACP Status: e - Enabled, d - Disabled
===============================================================================
Name
Id     Port-id        Adm   Act/     Opr   Primary Sub-group      Forced Prio
                            Stdby
-------------------------------------------------------------------------------
lag-1
1(e)   1/1/1          up    active   up    yes     1              -      32768
===============================================================================
```

PE-2 remains the DF for VPLS 1 and VPLS 3:

```
[/]
A:admin@PE-2# show service system bgp-evpn ethernet-segment name "SA-ESI-23" all | match "EVI
 Information" pre-lines 2 post-lines 24
===============================================================================
===============================================================================
EVI Information
===============================================================================
EVI                 SvcId               Actv Timer Rem      DF
-------------------------------------------------------------------------------
1                   1                   0                   yes
2                   2                   0                   no
3                   3                   0                   yes
-------------------------------------------------------------------------------
Number of entries: 3
===============================================================================
-------------------------------------------------------------------------------
DF Candidate list
-------------------------------------------------------------------------------
EVI                                     DF Address
-------------------------------------------------------------------------------
1                                       192.0.2.2
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
1                                   192.0.2.3
2                                   192.0.2.3
3                                   192.0.2.2
3                                   192.0.2.3
-------------------------------------------------------------------------------
Number of entries: 5
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
===============================================================================
```

On PE-2, ES "SA-ESI-23" remains active, so the operational group "op-grp-sa-es-23" is up and the monitoring LAG is also up:

```
[/]
A:admin@PE-2# show service oper-group "op-grp-sa-es-23" detail

===============================================================================
Service Oper Group Information
===============================================================================
Oper Group       : op-grp-sa-es-23
Creation Origin  : manual                          Oper Status: up
Hold DownTime    : 0 secs                          Hold UpTime: 0 secs
Members          : 1                               Monitoring : 1
===============================================================================

===============================================================================
Member Ethernet-Segment for OperGroup: op-grp-sa-es-23
===============================================================================
Ethernet-Segment                    Status
-------------------------------------------------------------------------------
SA-ESI-23                           Active
-------------------------------------------------------------------------------
Ethernet-Segment Entries found: 1
===============================================================================

===============================================================================
Monitoring LAG for OperGroup: op-grp-sa-es-23
===============================================================================
Lag-id         Adm       Opr       Weighted  Threshold Up-Count  Act/Stdby
    name
-------------------------------------------------------------------------------
1              up        up        No        0         1         N/A
    lag-1
-------------------------------------------------------------------------------
LAG Entries found: 1
===============================================================================
port option not supported with monitoring
```

The following commands on PE-2 shows that SAP lag-1:1 in VPLS 1 is up, SAP lag-1:2 in VPLS 2 is down (as it might be due to a failure or misconfiguration), and SAP lag-1:3 in VPLS 3 is up:

```
[/]
A:admin@PE-2# show service id 1 sap

===============================================================================
SAP(Summary), Service 1
===============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.  Egr.   Adm  Opr
                                           QoS   Fltr    QoS   Fltr
-------------------------------------------------------------------------------
lag-1:1                         1          1     none    1     none   Up   Up
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Number of SAPs : 1
------------------------------------------------------------------------------
==============================================================================
```

```
[/]
A:admin@PE-2# show service id 2 sap

==============================================================================
SAP(Summary), Service 2
==============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.  Egr.   Adm  Opr
                                           QoS   Fltr    QoS   Fltr
------------------------------------------------------------------------------
lag-1:2                         2          1     none    1     none   Down Down
------------------------------------------------------------------------------
Number of SAPs : 1
------------------------------------------------------------------------------
==============================================================================
```

```
[/]
A:admin@PE-2# show service id 3 sap

==============================================================================
SAP(Summary), Service 3
==============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.  Egr.   Adm  Opr
                                           QoS   Fltr    QoS   Fltr
------------------------------------------------------------------------------
lag-1:3                         3          1     none    1     none   Up   Up
------------------------------------------------------------------------------
Number of SAPs : 1
------------------------------------------------------------------------------
==============================================================================
```

On PE-3, lag-1:2 is up while lag-1:1 and lag-1:3 are down, as follows:

```
[/]
A:admin@PE-3# show service sap-using sap lag-1

==============================================================================
Service Access Points
==============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.  Egr.   Adm  Opr
                                           QoS   Fltr    QoS   Fltr
------------------------------------------------------------------------------
lag-1:1                         1          1     none    1     none   Up   Down
lag-1:2                         2          1     none    1     none   Up   Up
lag-1:3                         3          1     none    1     none   Up   Down
------------------------------------------------------------------------------
Number of SAPs : 3
------------------------------------------------------------------------------
==============================================================================
```

On CE-1, both ports in LAG 1 are up:

```
[/]
A:admin@CE-1# show lag "lag-1" port

==============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Lag Port States
LACP Status: e - Enabled, d - Disabled
===============================================================================
Name
Id     Port-id        Adm   Act/     Opr   Primary Sub-group      Forced Prio
                            Stdby
-------------------------------------------------------------------------------
lag-1
1(e)   1/1/1          up    active   up    yes     1              -      32768
       1/1/2          up    active   up            1              -      32768
===============================================================================
```

All traffic can take either LAG port, but PE-2 only forwards traffic for VPLS 1 and VPLS 3, while PE-3 only forwards traffic for VPLS 2. Traffic from VPLS 1 or VPLS 3 via port 1/1/2 to PE-3 is dropped by PE-3 because it is the NDF for VPLS 1 and VPLS 3. VPLS 2 traffic via LAG port 1/1/1 to PE-2 is dropped because SAP lag-1:2 is down (failure). This means that approximately 50% of the traffic is lost.

Potential loss on a single service under maintenance is acceptable but affecting other services on the same node is not acceptable. The solution is to disable the AC-DF capability.

AC-DF capability disabled

The default use of the AC-DF capability in SR OS is disabled on PE-2 and PE-3:

```
# on PE-2, PE-3:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "SA-ESI-23" {
                        ac-df-capability exclude
```

With AC-DF disabled, ES routes contain AC:0 in the DF-election extended community, as follows:

```
# on PE-3:
142 2022/06/08 15:54:10.390 CEST MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 71
    Flag: 0x90 Type: 14 Len: 34 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
        Type: EVPN-ETH-SEG Len: 23 RD: 192.0.2.2:0 ESI: 01:00:00:00:00:23:01:00:00:01, IP-Len:
 4 Orig-IP-Addr: 192.0.2.2
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        df-election::DF-Type:Preference/DP:1/DF-Preference:200/AC:0
        target:00:00:00:00:23:01
"
```

With the AC-DF capability disabled, the withdrawal of EVPN-AD routes does not influence the DF election. In this example, PE-2 remains the DF for all services, including VPLS 2, even when traffic for that service

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

is dropped by PE-2. The following command shows that the DF candidate list on PE-3 contains six entries: even for VPLS 2, PE-2 is included in the list. PE-3 is the NDF for all three services.

```
[/]
A:admin@PE-3# show service system bgp-evpn ethernet-segment name "SA-ESI-23" all

===============================================================================
Service Ethernet Segment
===============================================================================
Name                    : SA-ESI-23
Eth Seg Type            : None
Admin State             : Enabled            Oper State         : Up
ESI                     : 01:00:00:00:00:23:01:00:00:01
Oper ESI                : 01:00:00:00:00:23:01:00:00:01
Auto-ESI Type           : None
AC DF Capability        : Exclude
Multi-homing            : singleActive       Oper Multi-homing  : singleActive
ES SHG Label            : 524275
Source BMAC LSB         : None
Lag                     : lag-1
ES Activation Timer     : 3 secs (default)
Oper Group              : op-grp-sa-es-23
Svc Carving             : manual             Oper Svc Carving   : manual
Cfg Range Type          : lowest-pref

-------------------------------------------------------------------------------
DF Pref Election Information
-------------------------------------------------------------------------------
Preference     Preference     Last Admin Change        Oper Pref      Do No
Mode           Value                                   Value          Preempt
-------------------------------------------------------------------------------
non-revertive  100            06/08/2022 15:38:44      100            Enabled
-------------------------------------------------------------------------------
EVI Ranges: <none>
ISID Ranges: <none>
===============================================================================

===============================================================================
EVI Information
===============================================================================
EVI                 SvcId               Actv Timer Rem      DF
-------------------------------------------------------------------------------
1                   1                   0                   no
2                   2                   0                   no
3                   3                   0                   no
-------------------------------------------------------------------------------
Number of entries: 3
===============================================================================

-------------------------------------------------------------------------------
DF Candidate list
-------------------------------------------------------------------------------
EVI                                     DF Address
-------------------------------------------------------------------------------
1                                       192.0.2.2
1                                       192.0.2.3
2                                       192.0.2.2
2                                       192.0.2.3
3                                       192.0.2.2
3                                       192.0.2.3
-------------------------------------------------------------------------------
Number of entries: 6
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
---snip---
```

On NDF PE-3, the single-active ES "SA-ESI-23" is inactive and the ES operational group is down. The monitoring LAG is also operationally down.

On CE-1, LAG port 1/1/2 toward PE-3 is down:

```
[/]
A:admin@CE-1# show lag "lag-1" port

===============================================================================
Lag Port States
LACP Status: e - Enabled, d - Disabled
===============================================================================
Name
Id     Port-id        Adm   Act/     Opr   Primary Sub-group      Forced Prio
                            Stdby
-------------------------------------------------------------------------------
lag-1
1(e)   1/1/1          up    active   up    yes     1              -      32768
       1/1/2          up    active   down          1              -      32768
===============================================================================
```

CE-1 sends all traffic via LAG port 1/1/1 to PE-2. VPLS 1 and VPLS 3 traffic is forwarded by DF PE-2, whereas VPLS 2 traffic is dropped. Therefore, the failure does not have an impact on the other services.

On PE-2, SAP lag-1:1 in VPLS 1 and SAP lag-1:3 in VPLS 3 are operationally up:

```
[/]
A:admin@PE-2# show service id 1 sap

===============================================================================
SAP(Summary), Service 1
===============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.  Egr.   Adm  Opr
                                           QoS   Fltr    QoS   Fltr
-------------------------------------------------------------------------------
lag-1:1                         1          1     none    1     none   Up   Up
-------------------------------------------------------------------------------
Number of SAPs : 1
-------------------------------------------------------------------------------
===============================================================================


[/]
A:admin@PE-2# show service id 3 sap

===============================================================================
SAP(Summary), Service 3
===============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.  Egr.   Adm  Opr
                                           QoS   Fltr    QoS   Fltr
-------------------------------------------------------------------------------
lag-1:3                         3          1     none    1     none   Up   Up
-------------------------------------------------------------------------------
Number of SAPs : 1
-------------------------------------------------------------------------------
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

On PE-3, all SAPs in the VPLSs are down:

```
[/]
A:admin@PE-3# show service id 2 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 2                   Vpn Id            : 0
Service Type      : VPLS
MACSec enabled    : no
Name              : VPLS 2
---snip---

Admin State       : Up                  Oper State        : Up
---snip---

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                             Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:lag-1:2                            q-tag        1518    1518    Up   Down
===============================================================================
* indicates that the corresponding row element may have been truncated.


[/]
A:admin@PE-3# show service id 1 sap

===============================================================================
SAP(Summary), Service 1
===============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.  Egr.   Adm  Opr
                                           QoS   Fltr    QoS   Fltr
-------------------------------------------------------------------------------
lag-1:1                         1          1     none    1     none   Up   Down
-------------------------------------------------------------------------------
Number of SAPs : 1
-------------------------------------------------------------------------------
===============================================================================


[/]
A:admin@PE-3# show service id 3 sap

===============================================================================
SAP(Summary), Service 3
===============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.  Egr.   Adm  Opr
                                           QoS   Fltr    QoS   Fltr
-------------------------------------------------------------------------------
lag-1:3                         3          1     none    1     none   Up   Down
-------------------------------------------------------------------------------
Number of SAPs : 1
-------------------------------------------------------------------------------
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Conclusion

By default, the AC-DF capability is enabled. Disabling the AC-DF capability is recommended in ESs that use an operational group monitored by the access LAG to signal standby LACP or power-off.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

ARP-ND Host Routes in Data Centers

This chapter provides information about ARP-ND Host Routes in Data Centers.

Topics in this chapter include:

Applicability

This chapter was initially written based on SR OS Release 16.0.R1, but the MD-CLI in the current edition is based on SR OS Release 21.10.R3. Address Resolution Protocol - Neighbor Discovery (ARP-ND) host routes in VPRN and base router interfaces are supported in SR OS Release 15.0.R6 and later, but Nokia recommends using the feature in SR OS Release 15.0.R9, or later.

Chapters EVPN for MPLS Tunnels, EVPN for VXLAN Tunnels (Layer 2), EVPN for VXLAN Tunnels (Layer 3), and EVPN for MPLS Tunnels in Routed VPLS are prerequisite reading.

Overview

Inter-subnet forwarding (or simply routing) for a tenant domain in a Data Center (DC) must be efficient and avoid forwarding over the same path as arriving, known as tromboning or hairpinning. Figure 33: L2 broadcast domain extension across DCs shows an L2 broadcast domain (VPLS 1) extended across two DCs. This example is used to explain the requirement of upstream and downstream efficiency.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 33: L2 broadcast domain extension across DCs

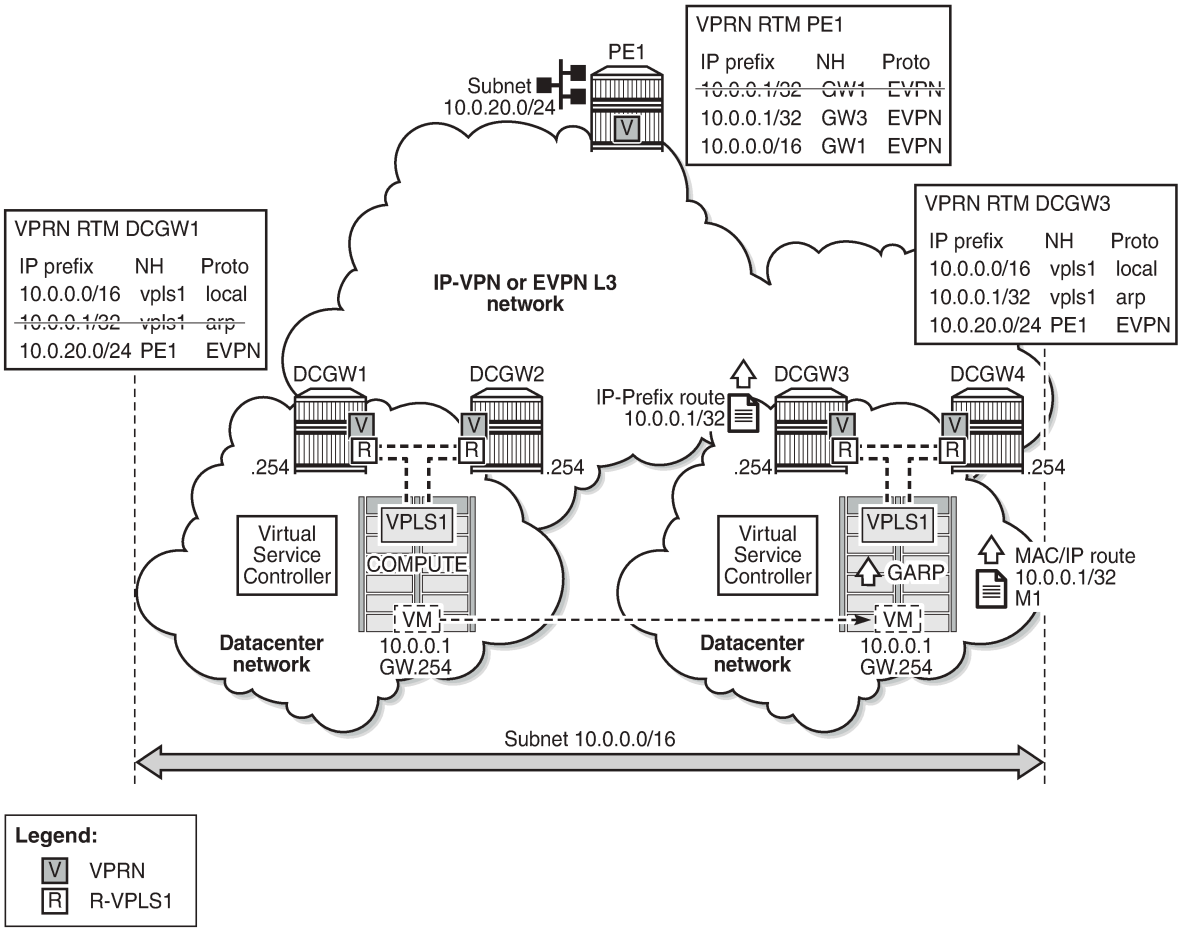

In Figure 33: L2 broadcast domain extension across DCs, subnet 10.0.0.0/16 is extended across two DCs and four DC Gateways (DCGWs), using VPLS 1 or R-VPLS 1 in the network nodes. The DCGWs are connected to the users of subnet 10.0.20.0/24 on PE1 via IP-VPN (or EVPN). In this scenario, there are two network characteristics that allow an efficient upstream and downstream routing:

- Anycast gateways
- ARP-ND host routes

Anycast Gateways provide upstream routing efficiency for the hosts connected to subnet 10.0.0.0/16, regardless of the DCGW to which they are connected. For example, if host 10.0.0.1 is in DC-1 and needs to forward traffic to subnet 10.0.20.0, DCGW1 and DCGW2 should be able to route the traffic upstream, without the need to go to DCGW3 or DCGW4. In the same way, if host 10.0.0.1 moves to DC-2, the upstream traffic to subnet 10.0.20.0 must be routed by the local DCGWs without changing the existing host default gateway IP and MAC configuration. To achieve this local default gateway routing, all the DCGWs of the extended broadcast domain need to have the same IP and MAC addresses in the R-VPLS interface (Integrated Routing and Bridging (IRB) interface in industry-standard terminology).

Anycast Gateways are implemented in SR OS by using passive VRRP. See EVPN for MPLS Tunnels in Routed VPLS for more information about passive VRRP.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

ARP-ND host routes learning and advertising are required to provide an efficient downstream routing from remote subnets to the hosts in the extended broadcast domain. Assuming virtual machine VM 10.0.0.1 (in Figure 33: L2 broadcast domain extension across DCs) is connected to DC-1 (left-side DC), when PE1 needs to send traffic to host 10.0.0.1, it will do a Longest Prefix Match (LPM) lookup on the VPRN route table. If the only IP prefix advertised by the four DCGWs were 10.0.0.0/16, PE1 could send the packets to a DC where the VM is not present. This would result in unnecessary tromboning; for example, PE1 could send the traffic to DCGW3, then DCGW3 would send it to DCGW2 to get to VM 10.0.0.1. However, PE1 could have forwarded directly to DCGW2.

To provide efficient downstream routing to the DC where the VM is located, DCGW1 and DCGW2 need to generate host routes for the VMs to which they are attached. Furthermore, when the VM moves to the other DC, DCGW3 and DCGW4 must be able to learn the VM host route and advertise it to PE1. Also, DCGW1 and DCGW2 will have to withdraw the route for 10.0.0.1, because the VM is no longer in the local DC.

To address this and other use cases, SR OS can learn the VM host route from the ARP or ND messages that it generates when it boots or when it moves. The host route can also be learned from EVPN routes type 2 (MAC/IP routes) that are installed in the ARP/ND caches, or in general, any ARP/ND entry can generate an ARP/ND host route.

A route owner type called "ARP-ND" is supported in the base router or a VPRN route table. The ARP-ND host routes have a preference of 1 and they are automatically created out of the ARP or ND Neighbor entries in the router instance. Figure 34: ARP-ND module and generated ARP-ND host routes shows how the ARP/ND software modules can generate ARP-ND host routes in the route table.

Figure 34: ARP-ND module and generated ARP-ND host routes

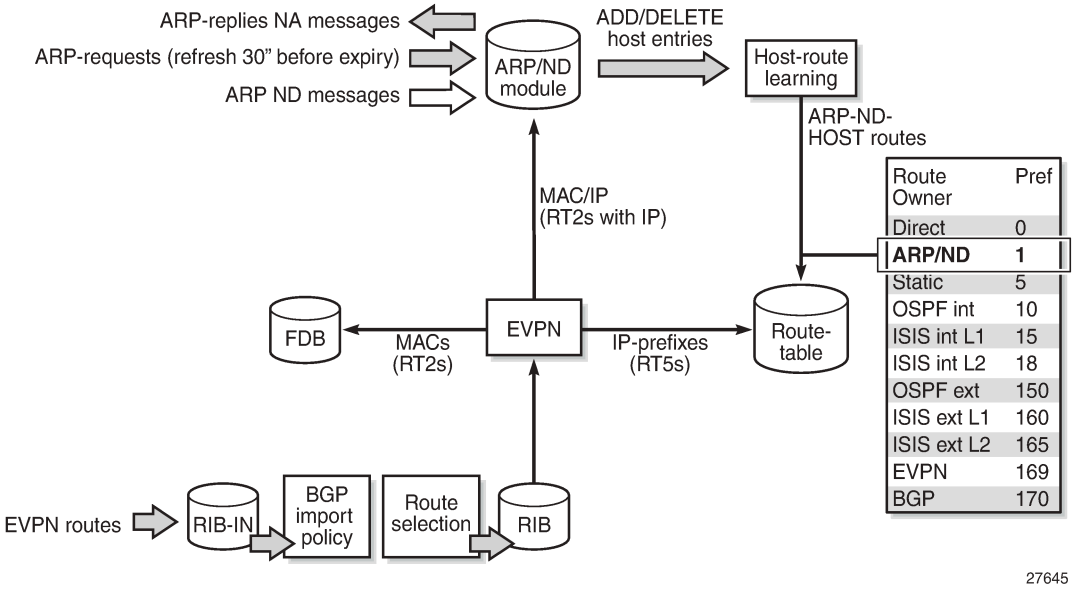

When **configure>service>vprn/ies>interface>ipv4>neighbor-discovery>host-route>populate [static | dynamic | evpn]** is enabled, the static, dynamic, and EVPN ARP entries of the routing context will create ARP-ND host routes in the route table. In the same way, ARP-ND host routes are created in the IPv6 route table out of static, dynamic, and EVPN neighbor entries, if **configure>service>vprn/ies>interface>ipv6>neighbor-discovery>host-route>populate[static | dynamic | evpn]** is enabled.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 34: ARP-ND module and generated ARP-ND host routes shows how the ARP/ND module populates its database from the usual dynamic and static entries, as well as from EVPN routes type 2 that include an IP address. Through the host-route learning action, ARP-ND host routes are handed over to the route table.

Figure 34: ARP-ND module and generated ARP-ND host routes also shows that the preference assigned to ARP-ND host routes is 1, which means that ARP-ND routes will be preferred over any other route owner, except for direct routes. For example, if the same host route gets to the route table from ARP-ND and VPN-IPv4 or EVPN, the ARP-ND host route will be preferred and added to the route table. Although they are added to the route table and advertised to routing protocols, ARP-ND host routes are never installed in the FIB. That helps preserve the FIB scale in the router.

The **neighbor-discovery>host-route>populate [static | dynamic | evpn]** commands are typically used along with other features:

- A route tag can be added to ARP-ND hosts by the command **route-tag**. This tag can be matched on BGP **vrf-export** and peer export policies.
- The ARP-ND host route will be kept in the route table while the corresponding ARP or Neighbor entry is active. The command **proactive-refresh** helps keep the entries active (even if there is no traffic destined to them) by sending an ARP refresh 30 seconds before the **timeout** or starting Neighbor Unreachable Detection (NUD) when the **stale-time** expires.
- To speed up the learning of the ARP-ND host routes, the command **learn-unsolicited** can be configured. When **learn-unsolicited** is enabled, received unsolicited ARP messages (typically, Gratuitous Address Resolution Protocol (GARP) messages) create an ARP entry, and therefore an ARP-ND route if **ipv4>neighbor-discovery>host-route>populate [static | dynamic | evpn]** is added. Similarly, unsolicited Neighbor Advertisement messages will create a "stale" neighbor. If **ipv6>neighbor-discovery>host-route>populate [static | dynamic | evpn]** is enabled, a confirmation message (NUD) is sent for all the neighbor entries created as stale, and, if confirmed, the corresponding ARP-ND routes are added to the route table.

In the example of Figure 33: L2 broadcast domain extension across DCs, **ipv4>neighbor-discovery>host-route>populate [static | dynamic | evpn]** on the DCGWs allows them to learn/advertise the ARP-ND host route 10.0.0.1/32 when the VM is locally connected, and remove/withdraw it when the VM is no longer present in the local DC.

The following sections describe three typical DC scenarios in which the use of Anycast gateways and ARP-ND host routes is needed. The examples are focused on IPv4 and ARP; however, there is equivalent functionality for IPv6 and ND.

Configuration

The initial configuration includes the following:

- Cards, MDAs, ports
- Router interfaces
- IS-IS as an IGP

The following three scenarios are configured and presented in this document:

- DC inter-subnet forwarding with Anycast GWs (and no ARP-ND hosts)
- DC inter-subnet forwarding with Anycast GWs and ARP-ND hosts
- Data Center Interconnect (DCI) inter-subnet forwarding with Anycast GWs and ARP-ND hosts

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

DC inter-subnet forwarding with Anycast GWs

Figure 35: DC inter-subnet forwarding with Anycast GWs shows a typical DC network, where PE-1, PE-2, and PE-3 are leaf switches that use EVPN-VXLAN services to provide connectivity between two subnets of a tenant domain. Those two subnets are 10.0.0.0/24 and 10.0.20.0/24, respectively, and while the three PEs are attached to hosts in the 10.0.0.0/24 subnet, only PE-1 is attached to the 10.0.20.0/24 subnet. Subnet 10.0.0.0/24 uses R-VPLS 17 in the three PEs and subnet 10.0.20.0/24 uses R-VPLS 22 in PE-1. The distribution of the R-VPLS services does not have to be uniform in all the PEs, and those R-VPLS services are only created if there are hosts attached to them.

To provide inter-subnet forwarding for the tenant, each PE must be configured with a VPRN instance (VPRN 16) that has an interface to the subnet R-VPLS. In industry-standard terms, VPRN 16 represents the IP-VRF for the tenant, and R-VPLS 17 and R-VPLS 22 are user Broadcast Domains (BDs). R-VPLS 15 is not a user BD, but rather a backhaul R-VPLS that provides EVPN connectivity among the VPRN instances.

Figure 35: DC inter-subnet forwarding with Anycast GWs

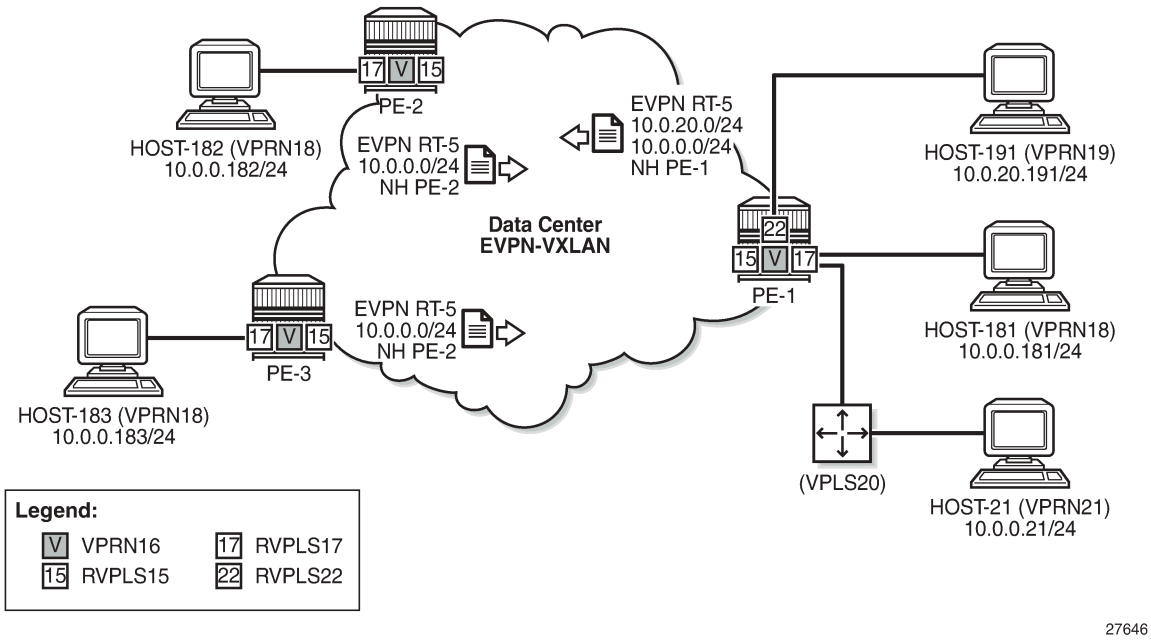

The BGP configuration in the PEs is similar. As an example, the BGP configuration in PE-1 is as follows:

```
# on PE-1:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            family {
                evpn true
            }
            rapid-update {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                evpn true
            }
            group "dc" {
                type internal
            }
            neighbor "192.0.2.2" {
                group "dc"
            }
            neighbor "192.0.2.3" {
                group "dc"
            }
        }
```

PE-2 has the following service configuration. The service configuration on PE-3 is similar.

```
# on PE-2:
configure {
    service {
        vpls "sbd-15" {
            admin-state enable
            description "R-VPLS 15"
            service-id 15
            customer "1"
            vxlan {
                instance 1 {
                    vni 15
                }
            }
            routed-vpls {
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 15
                routes {
                    mac-ip {
                        advertise false
                    }
                    ip-prefix {
                        advertise true
                    }
                }
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
        }
        vprn "ip-vrf-16" {
            admin-state enable
            service-id 16
            customer "1"
            ecmp 2
            interface "evi-15" {
                mac 00:00:00:00:00:02
                vpls "sbd-15" {
                    evpn-tunnel {
                    }
                }
            }
            interface "evi-17" {
                ipv4 {
                    primary {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        address 10.0.0.2
                        prefix-length 24
                    }
                    vrrp 1 {
                        backup [10.0.0.254]
                        passive true
                        ping-reply true
                        traceroute-reply true
                    }
                }
                vpls "evi-17" {
                }
            }

        }
        vpls "evi-17" {
            admin-state enable
            description "R-VPLS 17"
            service-id 17
            customer "1"
            vxlan {
                instance 1 {
                    vni 17
                }
            }
            routed-vpls {
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 17
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
            sap pxc-10.a:17 {
            }
        }
```

R-VPLS 17, "evi-17" in the configuration, is the BD used by subnet 10.0.0.0/24 in all the PEs. On the evi-17 interface in VPRN 16, a real IP address as well as a virtual (passive VRRP) IP address are configured. The real IP address is a unique address across the three PEs in R-VPLS 17 (10.0.0.2 in PE-2). This IP address will not be used by the R-VPLS 17 hosts as a default gateway, but rather will be used for troubleshooting purposes (ICMP or similar).

The backup IP address in the passive VRRP instance (10.0.0.254) is the Anycast Gateway IP address, and the same IP address is configured in all the PEs attached to R-VPLS 17. Because the virtual MAC is auto-derived from the VRRP instance, all the PEs will also have the same virtual MAC for this Anycast Gateway:

```
[/]
A:admin@PE-2# show router 16 vrrp instance interface "evi-17"

===============================================================================
VRRP Instances for interface "evi-17"
===============================================================================
-------------------------------------------------------------------------------
VRID 1
-------------------------------------------------------------------------------
Owner               : No                  VRRP State        : Master
Passive             : Yes
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Primary IP of Master: 10.0.0.2 (Self)
Primary IP          : 10.0.0.2            Standby-Forwarding: Disabled
VRRP Backup Addr    : 10.0.0.254
Admin State         : Up                  Oper State        : Up
Up Time             : 02/21/2022 16:38:17 Virt MAC Addr     : 00:00:5e:00:01:01
---snip---
```

```
[/]
A:admin@PE-3# show router 16 vrrp instance interface "evi-17"

===============================================================================
VRRP Instances for interface "evi-17"
===============================================================================
-------------------------------------------------------------------------------
VRID 1
-------------------------------------------------------------------------------
Owner               : No                  VRRP State        : Master
Passive             : Yes
Primary IP of Master: 10.0.0.3 (Self)
Primary IP          : 10.0.0.3            Standby-Forwarding: Disabled
VRRP Backup Addr    : 10.0.0.254
Admin State         : Up                  Oper State        : Up
Up Time             : 02/21/2022 16:38:33 Virt MAC Addr     : 00:00:5e:00:01:01
---snip---
```

```
[/]
A:admin@PE-1# show router 16 vrrp instance interface "evi-17"

===============================================================================
VRRP Instances for interface "evi-17"
===============================================================================
-------------------------------------------------------------------------------
VRID 1
-------------------------------------------------------------------------------
Owner               : No                  VRRP State        : Master
Passive             : Yes
Primary IP of Master: 10.0.0.1 (Self)
Primary IP          : 10.0.0.1            Standby-Forwarding: Disabled
VRRP Backup Addr    : 10.0.0.254
Admin State         : Up                  Oper State        : Up
Up Time             : 02/21/2022 16:38:06 Virt MAC Addr     : 00:00:5e:00:01:01
---snip---
```

All the hosts attached to R-VPLS 17, such as host-181, host-182, and host-183, are configured with the Anycast Gateway as default gateway (10.0.0.254). The use of passive VRRP (or Anycast Gateway in standard terminology) has the following benefits:

- All the hosts use the same default gateway configuration, regardless of what PE they are attached to.
- When the hosts send traffic destined to a remote subnet, the local PE can route it directly, without any tromboning.
- In the case of a host moving to a different leaf switch, the host does not need to change its IP or default gateway, or even its ARP cache.

For completeness, the service configuration in PE-1 follows:

```
# on PE-1:
configure {
    service {
        vpls "sbd-15" {
            admin-state enable
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            description "R-VPLS 15"
            service-id 15
            customer "1"
            vxlan {
                instance 1 {
                    vni 15
                }
            }
            routed-vpls {
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 15
                routes {
                    mac-ip {
                        advertise false
                    }
                    ip-prefix {
                        advertise true
                    }
                }
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
        }
        vprn "ip-vrf-16" {
            admin-state enable
            service-id 16
            customer "1"
            ecmp 2
            interface "evi-15" {
                mac 00:00:00:00:00:01
                vpls "sbd-15" {
                    evpn-tunnel {
                    }
                }
            }
            interface "evi-17" {
                ipv4 {
                    primary {
                        address 10.0.0.1
                        prefix-length 24
                    }
                    vrrp 1 {
                        backup [10.0.0.254]
                        passive true
                        ping-reply true
                        traceroute-reply true
                    }
                }
                vpls "evi-17" {
                }
            }
            interface "evi-22" {
                ipv4 {
                    primary {
                        address 10.0.20.1
                        prefix-length 24
                    }
                    vrrp 1 {
                        backup [10.0.20.254]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    passive true
                    ping-reply true
                    traceroute-reply true
                }
            }
            vpls "evi-22" {
            }
        }
    }
    vpls "evi-17" {
        admin-state enable
        description "R-VPLS 17"
        service-id 17
        customer "1"
        vxlan {
            instance 1 {
                vni 17
            }
        }
        routed-vpls {
        }
        bgp 1 {
        }
        bgp-evpn {
            evi 17
            vxlan 1 {
                admin-state enable
                vxlan-instance 1
            }
        }
        sap pxc-10.a:17 {
        }
        sap pxc-10.b:20 {
        }
    }
    vpls "evi-22" {
        admin-state enable
        description "R-VPLS 22"
        service-id 22
        customer "1"
        routed-vpls {
        }
        sap pxc-10.b:19 {
        }
    }
```

See the EVPN for VXLAN Tunnels (Layer 3) for more information about the EVPN-related configuration in the R-VPLS services. When there is no need for a recursive resolution of the EVPN IP prefix routes to a MAC/IP route, **mac-ip>advertisement false** is configured in the R-VPLS 15, compared to the examples in EVPN for VXLAN Tunnels (Layer 3).

With the described configuration, as an example, the intra-subnet and inter-subnet forwarding connectivity from host-182 is tested (host-182 is simulated with VPRN "VM-test-anycast-gw" that is connected to R-VPLS 17 via PXC SAP):

```
[/]
A:admin@PE-2# traceroute 10.0.0.183 router-instance "VM-test-anycast-gw"
                                                        source-address 10.0.0.182
traceroute to 10.0.0.183 from 10.0.0.182, 30 hops max, 40 byte packets
  1  10.0.0.183 (10.0.0.183)    6.61 ms  3.90 ms  3.79 ms
```

```
[/]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
A:admin@PE-2# traceroute 10.0.20.191 router-instance "VM-test-anycast-gw"
                                                        source-address 10.0.0.182
traceroute to 10.0.20.191 from 10.0.0.182, 30 hops max, 40 byte packets
  1  10.0.0.2 (10.0.0.2)    1.71 ms  2.31 ms  2.34 ms
  2  10.0.20.1 (10.0.20.1)    3.09 ms  4.56 ms  2.99 ms
  3  10.0.20.191 (10.0.20.191)    3.91 ms  3.85 ms  3.78 ms
```

When host-182 sends traffic to host-191, it will ARP for the Anycast Gateway IP and will receive the virtual MAC as a reply. The virtual MAC is always associated with the local CPM on the local PE; therefore, the local PE can always route the traffic directly while it has a route for the IP destination.

Host-182 (VPRN 18) resolves the Anycast Gateway to the virtual MAC:

```
[/]
A:admin@PE-2# show router 18 arp 10.0.0.254

===============================================================================
ARP Table (Service: 18)
===============================================================================
IP Address      MAC Address       Expiry    Type   Interface
-------------------------------------------------------------------------------
10.0.0.254      00:00:5e:00:01:01 03h59m18s Dyn[I] local
===============================================================================
```

In PE-2, the virtual MAC is associated with a local IP interface:

```
[/]
A:admin@PE-2# show service id 17 fdb mac 00:00:5e:00:01:01

===============================================================================
Forwarding Database, Service 17
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
17         00:00:5e:00:01:01 cpm                     Intf     02/21/22 16:38:17
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

The following route table of VPRN 16 on PE-2 shows that subnet 10.0.20.0/24 from host-191 is learned via EVPN:

```
[/]
A:admin@PE-3# show router 16 route-table

===============================================================================
Route Table (Service: 16)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
10.0.0.0/24                                   Local   Local     00h03m10s  0
       evi-17                                                        0
10.0.20.0/24                                  Remote  EVPN-IFF  00h03m08s  169
       evi-15 (ET-00:00:00:00:00:01)                                 0
-------------------------------------------------------------------------------
No. of Routes: 2
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
       S = Sticky ECMP requested
===============================================================================
```

DC inter-subnet forwarding with Anycast GWs and ARP-ND host routes

While the configuration shown in the preceding section is common in DCs, there is a variation that eliminates the flooding among PEs that are attached to the same BD, typically caused by ARP messages and ND. The configuration described in this section is recommended only if all the following conditions are met:

- All the hosts are directly connected to the leaf switches (PEs in Figure 35: DC inter-subnet forwarding with Anycast GWs).
- All the hosts announce themselves by issuing a GARP (or unsolicited NA for IPv6) whenever they boot up or move to a different leaf switch.

  Note: This is the case for virtual machines.
- All the traffic among hosts is IP unicast or non-IP unicast (if the hosts are in the same BD), and there is no Broadcast, Unknown unicast, or Multicast (BUM) traffic from the hosts in the tenant domain, other than ARP/ND.

If the preceding conditions are true, the ARP-ND host route feature can help eliminate BUM traffic completely.

Figure 36: DC inter-subnet forwarding with Anycast GWs and ARP-ND host routes shows the scenario used in this section.

Figure 36: DC inter-subnet forwarding with Anycast GWs and ARP-ND host routes

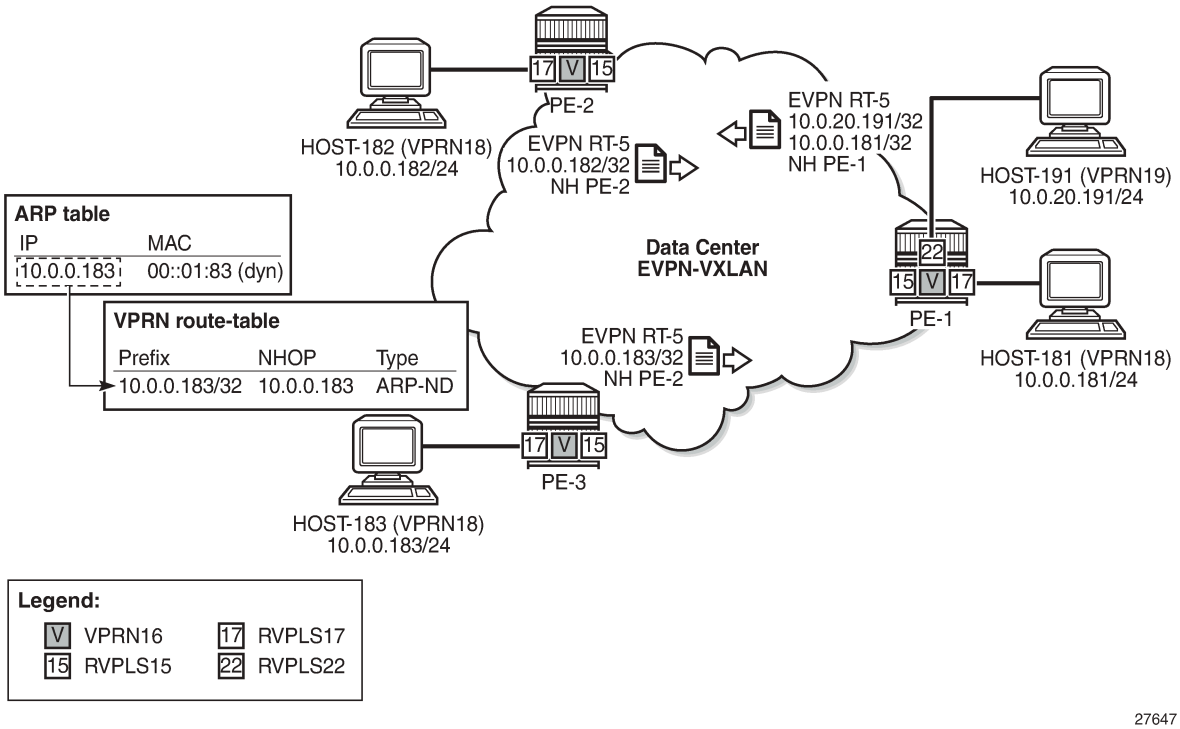

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Compared to the configuration used in the preceding section, VPRN 16 is modified in the three PEs as follows (changes in **neighbor-discovery** context):

```
# on PE-2:
configure {
    service {
        vprn "ip-vrf-16" {
            admin-state enable
            service-id 16
            customer "1"
            ecmp 2
            interface "evi-15" {
                mac 00:00:00:00:00:02
                vpls "sbd-15" {
                    evpn-tunnel {
                    }
                }
            }
            interface "evi-17" {
                mac 00:00:00:00:2e:17
                ipv4 {
                    primary {
                        address 10.0.0.2
                        prefix-length 24
                    }
                    neighbor-discovery {
                        timeout 300
                        learn-unsolicited true
                        proactive-refresh true
                        remote-proxy-arp true
                        host-route {
                            populate static {
                            }
                            populate dynamic {
                            }
                            populate evpn {
                            }
                        }
                    }
                    vrrp 1 {
                        backup [10.0.0.254]
                        passive true
                        ping-reply true
                        traceroute-reply true
                    }
                }
                vpls "evi-17" {
                }
            }
```

```
# on PE-3:
configure {
    service {
        vprn "ip-vrf-16" {
            admin-state enable
            service-id 16
            customer "1"
            ecmp 2
            interface "evi-15" {
                mac 00:00:00:00:00:03
                vpls "sbd-15" {
                    evpn-tunnel {
                    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
            }
            interface "evi-17" {
                mac 00:00:00:00:3e:17
                ipv4 {
                    primary {
                        address 10.0.0.3
                        prefix-length 24
                    }
                    neighbor-discovery {
                        timeout 300
                        learn-unsolicited true
                        proactive-refresh true
                        remote-proxy-arp true
                        host-route {
                            populate static {
                            }
                            populate dynamic {
                            }
                            populate evpn {
                            }
                        }
                    }
                    vrrp 1 {
                        backup [10.0.0.254]
                        passive true
                        ping-reply true
                        traceroute-reply true
                    }
                }
                vpls "evi-17" {
                }
            }
```

```
# on PE-1:
configure {
    service {
        vprn "ip-vrf-16" {
            admin-state enable
            service-id 16
            customer "1"
            ecmp 2
            interface "evi-15" {
                mac 00:00:00:00:00:01
                vpls "sbd-15" {
                    evpn-tunnel {
                    }
                }
            }
            interface "evi-17" {
                mac 00:00:00:00:1e:17
                ipv4 {
                    primary {
                        address 10.0.0.1
                        prefix-length 24
                    }
                    neighbor-discovery {
                        timeout 300
                        learn-unsolicited true
                        proactive-refresh true
                        remote-proxy-arp true
                        host-route {
                            populate static {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        }
                        populate dynamic {
                        }
                        populate evpn {
                        }
                    }
                }
                vrrp 1 {
                    backup [10.0.0.254]
                    passive true
                    ping-reply true
                    traceroute-reply true
                }
            }
            vpls "evi-17" {
            }
        }
        interface "evi-22" {
            mac 00:00:00:00:1e:22
            ipv4 {
                primary {
                    address 10.0.20.1
                    prefix-length 24
                }
                neighbor-discovery {
                    timeout 300
                    learn-unsolicited true
                    proactive-refresh true
                    remote-proxy-arp true
                    host-route {
                        populate static {
                        }
                        populate dynamic {
                        }
                        populate evpn {
                        }
                    }
                }
                vrrp 1 {
                    backup [10.0.20.254]
                    passive true
                    ping-reply true
                    traceroute-reply true
                }
            }
            vpls "evi-22" {
            }
        }
```

The behavior due to the added commands in the **neighbor-discovery** context is as follows:

- **host-route>populate [static | dynamic | evpn]** makes the router create an ARP-ND host route per ARP entry in the route table of VPRN "ip-vrf-16".
- **learn-unsolicited** makes the router learn ARP entries for the hosts out of the GARP messages that they send when they boot up or move. Without this command, ARP entries are only created after the router receives packets with the host as the destination, issues an ARP request, and the host replies to this solicited ARP request.
- **proactive-refresh** makes the router refresh every dynamic ARP entry even if there is no traffic destined to the owner. Without the command, host IP addresses will not be maintained in the ARP cache unless they receive traffic from remote hosts.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- **timeout 300** is the timeout selected in this example (in seconds). The ARP timeout has an impact on how often the router will try to refresh an entry (30 seconds before the timeout expires). In environments where the hosts are subject to mobility (VMs moving between leaves), having a shorter ARP timeout will speed up the removal of the old ARP entry, that is, the old ARP-ND host route entry. However, in scaled environments with tens of thousands of ARP entries, Nokia does not recommend lowering the ARP timeout under 10 minutes.
- **remote-proxy-arp** allows the router to reply to any ARP request looking for an IP address in the same subnet as the source, with its virtual MAC (00:00:5e:00:01:01), and route the traffic, as long as there is a route for the destination in the route table.

In addition, the following command will be executed in the three PEs:

```
# on PE-1, PE-2, PE-3:
configure {
    service {
        vpls "evi-17" {
            bgp-evpn {
                routes {
                    incl-mcast {
                        advertise-ingress-replication false
                    }
                }
            }
```

By disabling the advertisement of the Inclusive Multicast Ethernet Tag (IMET) route in R-VPLS 17, the PEs will not create a VXLAN BUM destination among each other, preventing the exchange of BUM traffic. Only known unicast traffic can be now exchanged in the context of R-VPLS 17. The three PEs will show VXLAN destinations that have Mcast "-", as opposed to "BUM":

```
[/]
A:admin@PE-3# show service id 17 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
Instance    VTEP Address                          Egress VNI  EvpnStatic Num
Mcast       Oper State                            L2 PBR     SupBcasDom MACs
-------------------------------------------------------------------------------
1           192.0.2.1                             17          evpn       3
-           Up                                    No          No
1           192.0.2.2                             17          evpn       2
-           Up                                    No          No
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 2
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN-VXLAN Ethernet Segment Dest
===============================================================================
Instance  Eth SegId                      Num. Macs     Last Change
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

With the described configuration, when the hosts boot up and generate a GARP message, the ARP entries will be created, and subsequently ARP-ND hosts and EVPN IP-prefix advertisements for them. The host

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

bootup is simulated by disabling and re-enabling the VPRN that emulates the host. As an example, some debug commands are used to see the behavior when host-181 boots up and sends a GARP:

```
1 2022/02/21 16:47:36.844 CET MINOR: DEBUG #2001 vprn18 PIP
"PIP: ARP
instance 2 (18), interface index 4 (local),
ARP egressing on local
   Who has 10.0.0.181 ? Tell 10.0.0.181
"

2 2022/02/21 16:47:36.845 CET MINOR: DEBUG #2001 vprn16 PIP
"PIP: ARP
instance 5 (16), interface index 8 (evi-17),
ARP ingressing on evi-17
   Who has 10.0.0.181 ? Tell 10.0.0.181
"

4 2022/02/21 16:47:36.845 CET MINOR: DEBUG #2001 vprn21 PIP
"PIP: ARP
instance 4 (21), interface index 6 (local),
ARP ingressing on local
   Who has 10.0.0.181 ? Tell 10.0.0.181
"
```

The GARP creates an ARP entry and, subsequently, an ARP-ND host route in the route table of VPRN 16. Host-181 MAC/IP and IP-prefix routes are advertised too:

```
3 2022/02/21 16:47:36.845 CET MINOR: DEBUG #2001 vprn16 PIP
"PIP: ROUTE
instance 5 (16), RTM ADD event
   New Route Info
      prefix: 10.0.0.181/32 (0xb8fcb278)  preference: 1   metric: 0  backup metric: 0
             owner: ARP-ND ownerId: 0
      1 ecmp hops  0 backup hops:
         hop 0:  10.0.0.181 @ if 8, weight 0
"

5 2022/02/21 16:47:36.845 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 90
    Flag: 0x90 Type: 14 Len: 45 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.1
        Type: EVPN-IP-PREFIX Len: 34 RD: 192.0.2.1:15, tag: 0,
              ip_prefix: 10.0.0.181/32 gw_ip 0.0.0.0 Label: 15 (Raw Label: 0xf)
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:15
        mac-nh:00:00:00:00:00:01
        bgp-tunnel-encap:VXLAN
"

6 2022/02/21 16:47:36.845 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 81
    Flag: 0x90 Type: 14 Len: 44 Multiprotocol Reachable NLRI:
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.1
        Type: EVPN-MAC Len: 33 RD: 192.0.2.1:17 ESI: ESI-0, tag: 0, mac len: 48
                       mac: 00:00:00:00:01:81, IP len: 0, IP: NULL, label1: 17
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:17
        bgp-tunnel-encap:VXLAN
"
```

As an example, following are the ARP and route tables in PE-1:

```
[/]
A:admin@PE-1# show router 16 arp

===============================================================================
ARP Table (Service: 16)
===============================================================================
IP Address      MAC Address       Expiry    Type   Interface
-------------------------------------------------------------------------------
10.0.0.1        00:00:00:00:1e:17 00h00m00s Oth[I] evi-17
10.0.0.2        00:00:00:00:2e:17 00h00m00s Evp[I] evi-17
10.0.0.3        00:00:00:00:3e:17 00h00m00s Evp[I] evi-17
10.0.0.181      00:00:00:00:01:81 00h02m19s Dyn[I] evi-17
10.0.0.254      00:00:5e:00:01:01 00h00m00s Oth[I] evi-17
10.0.20.1       00:00:00:00:1e:22 00h00m00s Oth[I] evi-22
10.0.20.254     00:00:5e:00:01:01 00h00m00s Oth[I] evi-22
-------------------------------------------------------------------------------
No. of ARP Entries: 7
===============================================================================
```

```
[/]
A:admin@PE-1# show router 16 route-table

===============================================================================
Route Table (Service: 16)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
10.0.0.0/24                                   Local   Local     00h03m40s  0
       evi-17                                                        0
10.0.0.2/32                                   Remote  ARP-ND    00h03m34s  1
       10.0.0.2                                                      0
10.0.0.3/32                                   Remote  ARP-ND    00h03m34s  1
       10.0.0.3                                                      0
10.0.0.181/32                                 Remote  ARP-ND    00h03m23s  1
       10.0.0.181                                                    0
10.0.0.182/32                                 Remote  EVPN-IFF  00h03m31s  169
       evi-15 (ET-00:00:00:00:00:02)                                 0
10.0.0.183/32                                 Remote  EVPN-IFF  00h03m34s  169
       evi-15 (ET-00:00:00:00:00:03)                                 0
10.0.20.0/24                                  Local   Local     00h03m40s  0
       evi-22                                                        0
-------------------------------------------------------------------------------
No. of Routes: 7
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
```

As discussed, the ARP-ND host routes are installed in the route table, but not in the FIB:

```
[/]
A:admin@PE-1# show router 16 fib 1

===============================================================================
FIB Display
===============================================================================
Prefix [Flags]                                                  Protocol
  NextHop
-------------------------------------------------------------------------------
10.0.0.0/24                                                     LOCAL
  10.0.0.0 (evi-17)
10.0.0.182/32                                                   EVPN-IFF
  (evi-15 (ET-00:00:00:00:00:02))
10.0.0.183/32                                                   EVPN-IFF
  (evi-15 (ET-00:00:00:00:00:03))
10.0.20.0/24                                                    LOCAL
  10.0.20.0 (evi-22)
-------------------------------------------------------------------------------
Total Entries : 4
-------------------------------------------------------------------------------
===============================================================================
```

A side effect of this scenario is that traffic between hosts in the same BD (R-VPLS "evi-17") is routed instead of switched. This can be shown on the traceroute from host-181 to host-182 (there are three hops instead of two), or the TTL on the ping packets (62 instead of 64):

```
[/]
A:admin@PE-1# traceroute 10.0.0.182 router-instance "H-181" source-address 10.0.0.181
traceroute to 10.0.0.182 from 10.0.0.181, 30 hops max, 40 byte packets
  1  10.0.0.1 (10.0.0.1)    1.53 ms  2.37 ms  2.42 ms
  2  10.0.0.2 (10.0.0.2)    3.42 ms  3.26 ms  3.29 ms
  3  10.0.0.182 (10.0.0.182)    3.80 ms  3.54 ms  3.76 ms
```

```
[/]
A:admin@PE-1# ping 10.0.0.182 router-instance "H-181" source-address 10.0.0.181
PING 10.0.0.182 56 data bytes
64 bytes from 10.0.0.182: icmp_seq=1 ttl=62 time=3.40ms.
64 bytes from 10.0.0.182: icmp_seq=2 ttl=62 time=3.49ms.
64 bytes from 10.0.0.182: icmp_seq=3 ttl=62 time=3.43ms.
64 bytes from 10.0.0.182: icmp_seq=4 ttl=62 time=3.58ms.
64 bytes from 10.0.0.182: icmp_seq=5 ttl=62 time=3.21ms.

---- 10.0.0.182 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 3.21ms, avg = 3.42ms, max = 3.58ms, stddev = 0.121ms
```

This extension of a subnet across a pure routing domain is compliant with the virtual subnet concept described in RFC 7814.

DCI inter-subnet forwarding with Anycast GWs and ARP-ND hosts

Figure 37: DCI inter-subnet forwarding with Anycast GWs and ARP-ND host routes shows a DCI scenario where the use of Anycast GWs, ARP-ND hosts, and some additional configuration provide efficient inter-subnet forwarding within the tenant domain.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 37: DCI inter-subnet forwarding with Anycast GWs and ARP-ND host routes

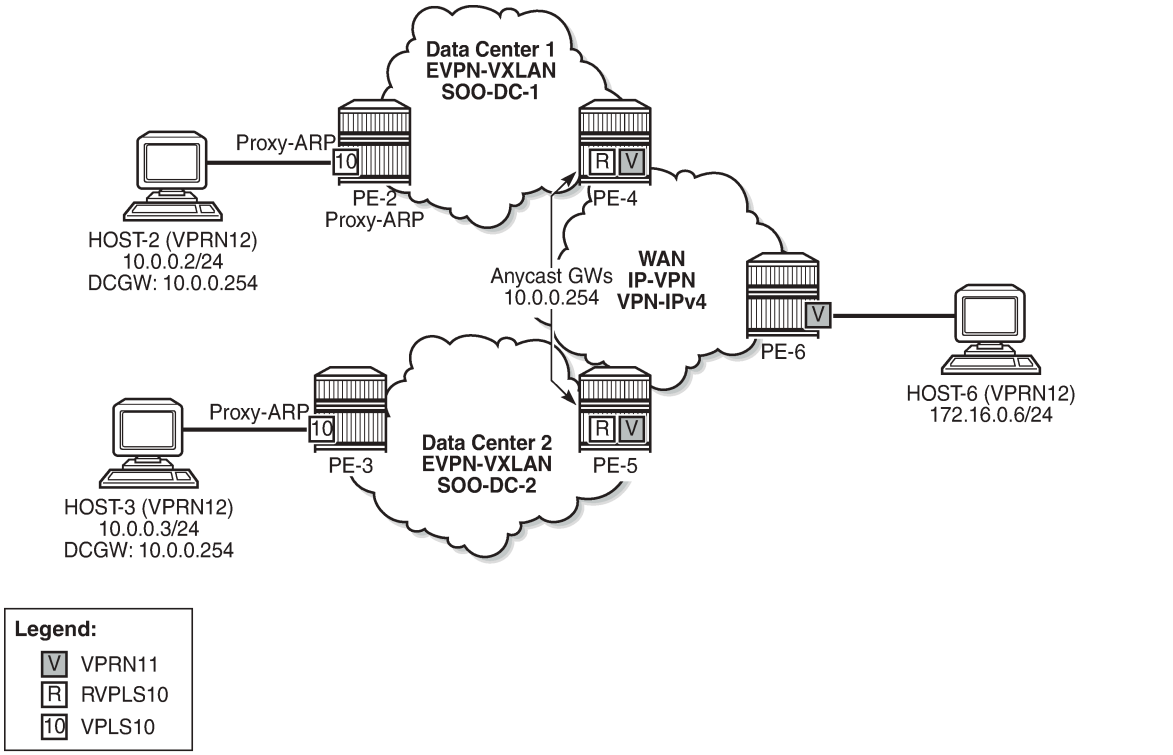

In this example, VPLS 10 is extended across DC-1 and DC-2, via PE-4 and PE-5 (which are DC GWs). PE-4 and PE-5 are also connected to the WAN and use IP-VPN for inter-subnet forwarding connectivity to the remote host-6. In this network, PE-4 and PE-5 provide the Anycast GW functionality to host-2 and host-3, so that they can move between the two DCs without having to change their IP/MAC/default GW or ARP cache, and efficient upstream forwarding is provided.

PE-4 and PE-5 learn the ARP-ND host route of their respective host and advertise it to the WAN, so that downstream routing from PE-6 can be efficient and without tromboning.

To avoid unnecessary ARP flooding between DCs, proxy-ARP is used in PE-2 and PE-3. The configuration of VPLS 10 in the PE-2 and PE-3 is as follows:

```
# on PE-2:
configure {
    service {
        vpls "centralized-gw-bd" {
            admin-state enable
            service-id 10
            customer "1"
            vxlan {
                instance 1 {
                    vni 10
                }
            }
            proxy-arp {
                admin-state enable
                dynamic-populate true
                send-refresh 120
                evpn {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    route-tag 1
                    flood {
                        unknown-arp-req false
                        gratuitous-arp false
                    }
                }
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 10
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
            sap pxc-10.a:10 {
            }
        }
```

```
# on PE-3:
configure {
    service {
        vpls "centralized-gw-bd" {
            admin-state enable
            service-id 10
            customer "1"
            vxlan {
                instance 1 {
                    vni 10
                }
            }
            proxy-arp {
                admin-state enable
                dynamic-populate true
                send-refresh 120
                evpn {
                    route-tag 1
                    flood {
                        unknown-arp-req false
                        gratuitous-arp false
                    }
                }
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 10
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
            sap pxc-10.a:10 {
            }
        }
```

Because the hosts are directly connected to PE-2 and PE-3, and they announce themselves to the network through a GARP when they boot up or move, the proxy-ARP configuration includes the parameters **evpn>flood>unknown-arp-req false** and **evpn>flood>gratuitous-arp false**. Those two commands prevent unnecessary ARP flooding between DCs.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The two PEs also include the **proxy-arp>evpn>route-tag 1** command. This command allows the proxy-ARP module to tag the routes when sent to BGP for advertisement of a MAC/IP route with non-zero IP. In this example, the tag is used in an export policy to add a Site-Of-Origin (SOO) extended community to the MAC/IP routes with non-zero IP. This, for example, allows PE-4 to accept MAC/IP routes from its own DC-1 and drop MAC/IP routes from DC-2 so that PE-4 only advertises ARP-ND host routes attached to DC-1. Vice versa for PE-5. The MAC/IP routes with zero-IP (that are also sent for every MAC) will not be tagged with the SOO and, therefore, will be imported by all the PEs in VPLS 10. This allows normal L2 connectivity among the four PEs, while the ARP-ND routes are only generated for the local hosts.

On PE-2, BGP is configured with the export policy "export-add-SOO" and import policy "import-prefer-DC-1", as follows:

```
# on PE-2:
configure {
    policy-options {
        community "SOO-DC-1" {
            member "origin:64500:1" { }
        }
        policy-statement "export-add-SOO" {
            entry 10 {
                from {
                    tag 1
                }
                action {
                    action-type accept
                    community {
                        add ["SOO-DC-1"]
                    }
                }
            }
        }
        policy-statement "import-prefer-DC-1" {
            entry 10 {
                from {
                    community {
                        name "SOO-DC-1"
                    }
                }
                action {
                    action-type accept
                    local-preference 200
                }
            }
        }
    }
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            family {
                ipv4 false
                vpn-ipv4 true
                vpn-ipv6 true
                evpn true
            }
            rapid-update {
                evpn true
            }
            import {
                policy ["import-prefer-DC-1"]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
            export {
                policy ["export-add-SOO"]
            }
            group "dc" {
                type internal
            }
            group "dcgws" {
                type internal
            }
            neighbor "192.0.2.3" {
                group "dc"
            }
            neighbor "192.0.2.4" {
                group "dcgws"
            }
            neighbor "192.0.2.5" {
                group "dcgws"
            }
        }
```

On PE-3, BGP is configured as follows:

```
# on PE-3:
configure {
    policy-options {
        community "SOO-DC-2" {
            member "origin:64500:2" { }
        }
        policy-statement "export-add-SOO" {
            entry 10 {
                from {
                    tag 1
                }
                action {
                    action-type accept
                    community {
                        add ["SOO-DC-2"]
                    }
                }
            }
        }
        policy-statement "import-prefer-DC-2" {
            entry 10 {
                from {
                    community {
                        name "SOO-DC-2"
                    }
                }
                action {
                    action-type accept
                    local-preference 200
                }
            }
        }
    }
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            family {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                ipv4 false
                vpn-ipv4 true
                vpn-ipv6 true
                evpn true
            }
            rapid-update {
                evpn true
            }
            import {
                policy ["import-prefer-DC-2"]
            }
            export {
                policy ["export-add-SOO"]
            }
            group "dc" {
                type internal
                peer-as 64500
            }
            group "dcgws" {
                type internal
            }
            neighbor "192.0.2.2" {
                group "dc"
            }
            neighbor "192.0.2.4" {
                group "dcgws"
            }
            neighbor "192.0.2.5" {
                group "dcgws"
            }
        }
```

As an example, the following show commands prove that PE-2 does not add an SOO to MAC/IP routes with zero-IP, but it does add SOO-DC-1 for MAC/IP routes with non-zero IP:

```
[/]
A:admin@PE-2# show router bgp routes evpn mac rd 192.0.2.2:10 hunt
===============================================================================
 BGP Router ID:192.0.2.2        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN MAC Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
---snip---

Network        : n/a
Nexthop        : 192.0.2.2
Path Id        : None
To             : 192.0.2.3
Res. Nexthop   : n/a
Local Pref.    : 100                    Interface Name : NotAvailable
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Aggregator AS  : None                    Aggregator     : None
Atomic Aggr.   : Not Atomic              MED            : None
AIGP Metric    : None                    IGP Cost       : n/a
Connector      : None
Community      : origin:64500:1 target:64500:10
                 bgp-tunnel-encap:VXLAN
Cluster        : No Cluster Members
Originator Id  : None                    Peer Router Id : 192.0.2.3
Origin         : IGP
AS-Path        : No As-Path
EVPN type      : MAC
ESI            : ESI-0
Tag            : 0
IP Address     : 10.0.0.2
Route Dist.    : 192.0.2.2:10
Mac Address    : 00:00:00:00:00:02
MPLS Label1    : VNI 10                  MPLS Label2    : n/a
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                       Dest Class     : 0

Network        : n/a
Nexthop        : 192.0.2.2
Path Id        : None
To             : 192.0.2.3
Res. Nexthop   : n/a
Local Pref.    : 100                     Interface Name : NotAvailable
Aggregator AS  : None                    Aggregator     : None
Atomic Aggr.   : Not Atomic              MED            : None
AIGP Metric    : None                    IGP Cost       : n/a
Connector      : None
Community      : target:64500:10 bgp-tunnel-encap:VXLAN
                 mac-mobility:Seq:0/Static
Cluster        : No Cluster Members
Originator Id  : None                    Peer Router Id : 192.0.2.3
Origin         : IGP
AS-Path        : No As-Path
EVPN type      : MAC
ESI            : ESI-0
Tag            : 0
IP Address     : n/a
Route Dist.    : 192.0.2.2:10
Mac Address    : 02:13:ff:00:03:3a
MPLS Label1    : VNI 10                  MPLS Label2    : n/a
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                       Dest Class     : 0

---snip---
```

The VPLS 10 configuration on PE-4 and the corresponding import policy to drop non-local SOO follow. PE-5 has a similar configuration (not shown), including the same RD 64500:10 in VPLS 10 as PE-4. The policy will drop routes tagged with SOO-DC-1 instead of SOO-DC-2.

```
# on PE-4:
configure {
    service {
        vpls "centralized-gw-bd" {
            admin-state enable
            service-id 10
            customer "1"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            vxlan {
                instance 1 {
                    vni 10
                }
            }
            routed-vpls {
            }
            bgp 1 {
                route-distinguisher "64500:10"
            }
            bgp-evpn {
                evi 10
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
        }
```

On PE-4, the BGP configuration is as follows:

```
# on PE-4:
configure {
    policy-options {
        community "SOO-DC-1" {
            member "origin:64500:1" { }
        }
        community "SOO-DC-2" {
            member "origin:64500:2" { }
        }
        policy-statement "export-add-SOO" {
            entry 10 {
                from {
                }
                action {
                    action-type accept
                    community {
                        add ["SOO-DC-1"]
                    }
                }
            }
        }
        policy-statement "import-drop-DC-2" {
            entry 10 {
                from {
                    community {
                        name "SOO-DC-2"
                    }
                }
                action {
                    action-type reject
                }
            }
        }
    }
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            family {
                ipv4 false
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                vpn-ipv4 true
                vpn-ipv6 true
                evpn true
            }
            rapid-update {
                evpn true
            }
            import {
                policy ["import-drop-DC-2"]
            }
            export {
                policy ["export-add-SOO"]
            }
            group "dc" {
                type internal
            }
            group "wan" {
                type internal
            }
            neighbor "192.0.2.2" {
                group "dc"
            }
            neighbor "192.0.2.3" {
                group "dc"
            }
            neighbor "192.0.2.5" {
                group "wan"
            }
            neighbor "192.0.2.6" {
                group "wan"
            }
        }
```

There is another aspect for which policies are used: on PE-2 and PE-3, two MAC/IP routes with the Anycast GW virtual MAC are received (one from PE-4 and another from PE5). To provide efficient upstream routing with no tromboning, it is important that PE-2 prefers the PE-4 virtual MAC route (its own DGW) over that of PE-5, and vice versa for PE-3. This is achieved by:

- Configuring the same RD on PE-4 and PE-5 for VPLS10.
- Configuring an import policy on PE-2 and PE-3 that modifies the local preference of the routes, so that each one prefers the local DGW.

PE-2 and PE-3 could have dropped the routes from the non-local DCGW, but with this configuration, DCGW redundancy is provided in case of failure:

```
[/]
A:admin@PE-2# show router policy plcy-name "import-prefer-DC-1"
    entry 10
        from
            community "SOO-DC-1"
        exit
        action accept
            local-preference 200
        exit
    exit
```

```
[/]
A:admin@PE-2# show router bgp routes evpn mac community target:64500:10
                                                    mac-address 00:00:5e:00:01:01
===============================================================================
 BGP Router ID:192.0.2.2        AS:64500       Local AS:64500
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN MAC Routes
===============================================================================
Flag  Route Dist.         MacAddr           ESI
      Tag                 Mac Mobility      Label1
                          Ip Address
                          NextHop
-------------------------------------------------------------------------------
u*>i  64500:10            00:00:5e:00:01:01 ESI-0
      0                   Static            VNI 10
                          10.0.0.254
                          192.0.2.4

*i    64500:10            00:00:5e:00:01:01 ESI-0
      0                   Static            VNI 10
                          10.0.0.254
                          192.0.2.5

-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

```
[/]
A:admin@PE-3# show router policy plcy-name "import-prefer-DC-2"
    entry 10
        from
            community "SOO-DC-2"
        exit
        action accept
            local-preference 200
        exit
    exit
```

```
[/]
A:admin@PE-3# show router bgp routes evpn mac community target:64500:10
                                                    mac-address 00:00:5e:00:01:01
===============================================================================
 BGP Router ID:192.0.2.3        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN MAC Routes
===============================================================================
Flag  Route Dist.         MacAddr           ESI
      Tag                 Mac Mobility      Label1
                          Ip Address
                          NextHop
-------------------------------------------------------------------------------
u*>i  64500:10            00:00:5e:00:01:01 ESI-0
      0                   Static            VNI 10
                          10.0.0.254
                          192.0.2.5
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
*i    64500:10            00:00:5e:00:01:01 ESI-0
      0                   Static            VNI 10
                          10.0.0.254
                          192.0.2.4

-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

Finally, the VPRN 11 configuration on PE-4 and PE-5 is as follows:

```
# on PE-4:
configure {    service {
        vprn "wan-ip-vpn" {
            admin-state enable
            service-id 11
            customer "1"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher auto-rd
                    vrf-target {
                        community "target:64500:11"
                    }
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            interface "evi-10" {
                mac 00:00:00:00:00:04
                ipv4 {
                    primary {
                        address 10.0.0.4
                        prefix-length 16
                    }
                    neighbor-discovery {
                        timeout 600
                        learn-unsolicited true
                        host-route {
                            populate static {
                                route-tag 1
                            }
                            populate dynamic {
                                route-tag 1
                            }
                            populate evpn {
                                route-tag 1
                            }
                        }
                    }
                    vrrp 1 {
                        backup [10.0.0.254]
                        passive true
                        ping-reply true
                        traceroute-reply true
                    }
                }
                vpls "centralized-gw-bd" {
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
```

```
# on PE-5:
configure {
    service {
        vprn "wan-ip-vpn" {
            admin-state enable
            service-id 11
            customer "1"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher auto-rd
                    vrf-target {
                        community "target:64500:11"
                    }
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            interface "evi-10" {
                mac 00:00:00:00:00:05
                ipv4 {
                    primary {
                        address 10.0.0.5
                        prefix-length 16
                    }
                    neighbor-discovery {
                        timeout 600
                        learn-unsolicited true
                        host-route {
                            populate static {
                                route-tag 1
                            }
                            populate dynamic {
                                route-tag 1
                            }
                            populate evpn {
                                route-tag 1
                            }
                        }
                    }
                    vrrp 1 {
                        backup [10.0.0.254]
                        passive true
                        ping-reply true
                        traceroute-reply true
                    }
                }
                vpls "centralized-gw-bd" {
                }
            }
        }
```

The passive VRRP commands, as well as the ARP commands, have already been discussed in preceding sections. The only new command in the configuration is **route-tag 1**. This command tags all the ARP-ND host routes learned on the interface, so that export policies can match on that tag and modify the routes before they are advertised. The command is included for completeness, however, in this configuration, there is no export policy using this tag.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

When the configuration is in place and the hosts are connected, the FDBs, proxy-ARP, ARP caches, and route tables are checked with the following commands (example for host-2 and host-6).

When host-2 ARPs for its default gateway (10.0.0.254), PE-2 will reply with the information from its proxy-ARP table:

```
[/]
A:admin@PE-2# show service id 10 proxy-arp 10.0.0.254 detail
-------------------------------------------------------------------------------
Proxy Arp
-------------------------------------------------------------------------------
Admin State       : enabled
Dyn Populate      : enabled
Age Time          : disabled            Send Refresh      : 120 secs
Table Size        : 250                 Total             : 5
Static Count      : 0                   EVPN Count        : 4
Dynamic Count     : 1                   Duplicate Count   : 0

Dup Detect
-------------------------------------------------------------------------------
Detect Window     : 3 mins              Num Moves         : 5
Hold down         : 9 mins
Anti Spoof MAC    : None

EVPN
-------------------------------------------------------------------------------
Garp Flood        : disabled            Req Flood         : disabled
Static Black Hole : disabled
EVPN Route Tag    : 1
-------------------------------------------------------------------------------

===============================================================================
VPLS Proxy Arp Entries
===============================================================================
IP Address          Mac Address         Type      Status    Last Update
-------------------------------------------------------------------------------
10.0.0.254          00:00:5e:00:01:01   evpn      active    02/21/2022 16:57:46
-------------------------------------------------------------------------------
Number of entries : 1
===============================================================================
```

When host-2 sends traffic to the virtual MAC, it will forward it to PE-4 based on a lookup on the FDB:

```
[/]
A:admin@PE-2# show service id 10 fdb mac 00:00:5e:00:01:01

===============================================================================
Forwarding Database, Service 10
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
10         00:00:5e:00:01:01 vxlan-1:                EvpnS:P  02/21/22 16:57:46
                             192.0.2.4:10
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

If PE-4 receives packets with MAC Destination Address (DA) equal to the virtual MAC and IP DA of host-6 (172.16.0.6), the forwarding is based on the information in the R-VPLS FDB first, and afterward on the VPRN 11 route table, as follows.

```
[/]
A:admin@PE-4# show service id 10 fdb mac 00:00:5e:00:01:01

===============================================================================
Forwarding Database, Service 10
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
10         00:00:5e:00:01:01 cpm                     Intf     02/21/22 16:57:46
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

```
[/]
A:admin@PE-4# show router 11 route-table

===============================================================================
Route Table (Service: 11)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
10.0.0.0/16                                     Local   Local     00h07m15s  0
       evi-10                                                          0
10.0.0.2/32                                     Remote  ARP-ND    00h07m12s  1
       10.0.0.2                                                        0
172.16.0.0/24                                   Remote  BGP VPN   00h06m35s  170
       192.0.2.6 (tunneled)                                            10
-------------------------------------------------------------------------------
No. of Routes: 3
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

When the traffic goes back from host-6 to host-2, PE-6 will forward to PE-4 due to an LPM lookup on the VPRN route table. The advertisement of the ARP-ND routes on PE-4 and PE-6 ensures that PE-6 can forward downstream traffic to the correct PE:

```
[/]
A:admin@PE-6# show router 11 route-table

===============================================================================
Route Table (Service: 11)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
10.0.0.0/16                                     Remote  BGP VPN   00h06m57s  170
       192.0.2.4 (tunneled)                                            10
10.0.0.0/16                                     Remote  BGP VPN   00h06m57s  170
       192.0.2.5 (tunneled)                                            10
10.0.0.2/32                                     Remote  BGP VPN   00h06m57s  170
       192.0.2.4 (tunneled)                                            10
10.0.0.3/32                                     Remote  BGP VPN   00h06m57s  170
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
       192.0.2.5 (tunneled)                                              10
172.16.0.0/24                                 Local   Local     00h07m01s  0
       local                                                        0
-------------------------------------------------------------------------------
No. of Routes: 5
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

Traceroute commands from host-6 provide information about the path to each remote host (VPRN 12 in PE-6 simulates host-6):

```
[/]
A:admin@PE-6# traceroute 10.0.0.2 router-instance "CE-PE-6"
traceroute to 10.0.0.2, 30 hops max, 40 byte packets
  1  172.16.0.254 (172.16.0.254)    1.85 ms  2.30 ms  2.26 ms
  2  10.0.0.4 (10.0.0.4)    3.38 ms  3.20 ms  5.66 ms
  3  10.0.0.2 (10.0.0.2)    4.80 ms  4.41 ms  4.73 ms
```

```
[/]
A:admin@PE-6# traceroute 10.0.0.3 router-instance "CE-PE-6"
traceroute to 10.0.0.3, 30 hops max, 40 byte packets
  1  172.16.0.254 (172.16.0.254)    1.55 ms  2.34 ms  2.28 ms
  2  10.0.0.5 (10.0.0.5)    3.53 ms  3.58 ms  3.43 ms
  3  10.0.0.3 (10.0.0.3)    7.18 ms  5.05 ms  4.74 ms
```

Communication between host-2 and host-3 uses regular L2 switching, as expected, because there are EVPN-VXLAN destinations created between PE-2 and PE-3 for VPLS 10:

```
[/]
A:admin@PE-2# show service id 10 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
Instance    VTEP Address                          Egress VNI  EvpnStatic Num
 Mcast       Oper State                            L2 PBR     SupBcasDom MACs
-------------------------------------------------------------------------------
1           192.0.2.3                             10          evpn       2
 BUM         Up                                    No          No
1           192.0.2.4                             10          evpn       2
 BUM         Up                                    No          No
1           192.0.2.5                             10          evpn       1
 BUM         Up                                    No          No
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 3
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN-VXLAN Ethernet Segment Dest
===============================================================================
Instance  Eth SegId                      Num. Macs     Last Change
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

```
[/]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
A:admin@PE-2# ping 10.0.0.3 router-instance "VM-PE-2"
PING 10.0.0.3 56 data bytes
64 bytes from 10.0.0.3: icmp_seq=1 ttl=64 time=8.84ms.
64 bytes from 10.0.0.3: icmp_seq=2 ttl=64 time=3.77ms.
64 bytes from 10.0.0.3: icmp_seq=3 ttl=64 time=3.54ms.
64 bytes from 10.0.0.3: icmp_seq=4 ttl=64 time=3.38ms.
64 bytes from 10.0.0.3: icmp_seq=5 ttl=64 time=3.25ms.

---- 10.0.0.3 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 3.25ms, avg = 4.56ms, max = 8.84ms, stddev = 2.15ms
```

```
[/]
A:admin@PE-2# traceroute 10.0.0.3 router-instance "VM-PE-2"
traceroute to 10.0.0.3, 30 hops max, 40 byte packets
  1  10.0.0.3 (10.0.0.3)    2.99 ms  3.41 ms  3.76 ms
```

Troubleshooting and debugging

The following commands can be used when troubleshooting these scenarios:

- **show router** <*id*> **route table** and **show router** <*id*> **fib** <*id*> (and their corresponding commands for IPv6)
- **show router** <*id*> **arp / neighbor**
- **show service** <*id*> **fdb detail**
- **show service** <*id*> **proxy-arp/nd detail**
- **show router bgp routes evpn / vpn-ipv4 / vpn-ipv6**

The following debug commands—in classic CLI—are also important to analyze the scenarios:

```
debug
    router "Base"
        bgp
            update
        exit
    exit
    router service-name "ip-vrf-16"
        ip
            arp
            route-table
        exit
    exit
    router service-name "VM-test-anycast-gw"
        ip
            arp
        exit
    exit
    service
        id 10
            proxy-arp
                 all
            exit
        exit
    exit
exit
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Conclusion

ARP-ND host routes are generated out of ARP-ND entries in a router context. These ARP-ND host routes, along with passive VRRP (for Anycast GWs), provide the correct solution for efficient inter-subnet forwarding in DCs and DCI networks.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Auto-Learn MAC Protect in EVPN

This chapter provides information about Auto-Learn MAC Protect in EVPN.

Topics in this chapter include:

Applicability

This chapter was initially written for SR OS Release 14.0.R5, but the MD-CLI in the current edition is based on SR OS Release 21.2.R1. Auto-Learn MAC Protect (ALMP) is supported for EVPN in SR OS Release 14.0.R1, and later.

Overview

MAC protection is needed in Layer 2 services to safeguard business-critical MAC addresses against the possibility of being learned on the wrong SAP/SDP-binding. When a MAC address is learned on the wrong SAP/SDP-binding, traffic would be diverted from its intended destination. This could be caused by misconfiguration or by a malicious source launching a Denial of Service (DoS) attack. MAC protect can also be used to prevent loops in certain topologies.

Chapter EVPN for VXLAN Tunnels (Layer 2) describes MAC protection for static MAC addresses that are configured on SAPs or spoke-SDPs. The command to configure static MAC addresses in a VPLS service is as follows:

```
*[ex:configure service vpls "1" fdb static-mac]
A:admin@PE-4# mac ?

 [mac-address] <unicast-mac-address-no-zero>
 <unicast-mac-address-no-zero> - <xx:xx:xx:xx:xx:xx>

    Static MAC address to SAP/SDP-binding or black-hole
```

Configuring static MAC addresses is not scalable if large numbers of MAC addresses need to be protected. Also, configuring static MAC addresses is not an option when the MAC addresses are unknown. Auto-Learn MAC Protect (ALMP) offers the same protection for learned MAC addresses in services such as EVPN VPLS and EVPN R-VPLS. However, ALMP is not supported for PBB-EVPN.

ALMP can be enabled with the **auto-learn-mac-protect** command in EVPN with VXLAN or MPLS bindings on the following:

- SAPs
- Mesh-SDPs
- Spoke-SDPs
- Pseudowire (PW) templates

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- Split Horizon Groups (SHGs)
- SHGs in PW templates

When enabled, all MAC addresses learned on those objects become protected.

The following commands can be used to enable ALMP on objects in VPLS 1:

```
configure {
    service {
        pw-template "PW1" {
            fdb {
                auto-learn-mac-protect true
            }
        }
        vpls "VPLS 1" {
            split-horizon-group "SHG1" {
                fdb {
                    saps {
                        auto-learn-mac-protect true
                    }
                }
            }
            spoke-sdp 23:1 {
                fdb {
                    auto-learn-mac-protect true
                }
            }
            mesh-sdp 24:1 {
                fdb {
                    auto-learn-mac-protect true
                }
            }
            sap 1/2/1:1 {
                fdb {
                    auto-learn-mac-protect true
                }
            }
        }
```

When enabled on an SHG, it is only applicable to the SAPs within the SHG, not to spoke-SDPs. If ALMP is required on spoke-SDPs in the SHG, the parameter must be configured on each spoke-SDP individually. All MAC Source Addresses (SAs) learned on these objects will be protected and advertised with the sticky bit set. The sticky bit indicates that these MAC addresses should be treated as protected on the remote PEs, where these protected MAC addresses are considered to have been learned on the EVPN MPLS/VXLAN destinations. The remote EVPN peers then use the MAC protection functionality in the same way as the local peer to protect the MAC address.

By default, ALMP enables an implicit **protected-src-mac-violation-action discard** (restrict protected source discard frame (RPS-DF)) on SAPs and spoke/mesh-SDPs. When enabled, frames with a protected MAC SA are discarded if received on objects where they were not learned and protected. This configuration is the default and cannot be configured on objects where MAC addresses are learned, such as SAPs, spoke/mesh-SDPs, and SHGs.

However, protected-src-mac-violation-action discard can optionally be configured on destinations in EVPN MPLS or EVPN VXLAN, where data plane MAC learning is never performed for incoming traffic. The only configurable protected source MAC violation action is discard; it is not an option to bring down the entire EVPN destination. For EVPN MPLS, this configuration is in the BGP EVPN context, as follows:

```
*[ex:configure service vpls "VPLS 1" bgp-evpn mpls 1 fdb]
A:admin@PE-2# protected-src-mac-violation-action ?
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
 protected-src-mac-violation-action <keyword>
 <keyword> - discard

    Action when a relearn request for a protected MAC is received
```

```
configure {
    service {
        vpls "VPLS 1" {
            bgp-evpn {
                mpls 1 {
                    fdb {
                        protected-src-mac-violation-action discard
                    }
                }
```

For EVPN VXLAN, this configuration is in the VXLAN context, as follows:

```
*[ex:configure service vpls "VPLS 1" vxlan instance 1 fdb]
A:admin@PE-2# protected-src-mac-violation-action ?

 protected-src-mac-violation-action <keyword>
 <keyword> - discard

    Action when a relearn request for a protected MAC is received
```

```
configure {
    service {
        vpls "VPLS 1" {
            vxlan {
                instance 1 {
                    vni 1
                    fdb {
                        protected-src-mac-violation-action discard
                    }
                }
```

Instead of discarding the frame, the SAP or spoke/mesh-SDP can be brought operationally down when a frame is received with a protected MAC SA that has not been learned on the object, by configuring **protected-src-mac-violation-action sap-oper-down** or **sdp-bind-oper-down** on the object in EVPN services. After the object has been brought down, an operator needs to disable and re-enable the object in order to make it operational again.

The protected source MAC violation action can be configured as **sap-oper-down** on SAPs and SHGs, or **sdp-bind-oper-down** on spoke/mesh-SDPs, and PW templates, but not on EVPN MPLS/VXLAN destinations, using following commands:

```
configure {
    service {
        pw-template "PW1" {
            fdb {
                auto-learn-mac-protect true
                protected-src-mac-violation-action sdp-bind-oper-down
            }
        }
        pw-template "PW2" {
            split-horizon-group {
                name "SHG1"
                fdb {
                    saps {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    auto-learn-mac-protect true
                    protected-src-mac-violation-action sap-oper-down
                }
            }
        }
    }
    vpls "VPLS 1" {
        split-horizon-group "SHG1" {
            fdb {
                saps {
                    auto-learn-mac-protect true
                    protected-src-mac-violation-action sap-oper-down
                }
            }
        }
        spoke-sdp 23:1 {
            fdb {
                auto-learn-mac-protect true
                protected-src-mac-violation-action sdp-bind-oper-down
            }
        }
        mesh-sdp 24:1 {
            fdb {
                auto-learn-mac-protect true
                protected-src-mac-violation-action sdp-bind-oper-down
            }
        }
        sap 1/2/1:1 {
            fdb {
                auto-learn-mac-protect true
                protected-src-mac-violation-action sap-oper-down
            }
        }
```

Note:
The configuration of protected-src-mac-violation-action alarm-only is not allowed in BGP-EVPN.

Protection is provided at the point where a MAC address first enters the EVPN part of the network. Therefore, the preference for an auto-learned protected MAC address is higher than that of a MAC address received in a BGP update with the sticky bit set.

The following list shows the MAC learning priority, with the highest priority first:

1. Local MAC address (including AS-MAC without static-black-hole, es-bmac, src-bmac, OAM, and so on)
2. Conditional static MAC address (including AS-MAC with static-black-hole)
3. **Auto-learn protected MAC address**
4. EVPN MAC address with sticky/static bit set
5. Data plane learned MAC address (regular learning on SAP/SDP-binding)
6. EVPN MAC address without sticky/static bit set

Note:
ALMP MAC addresses have a higher priority but do not overwrite EVPN static MAC addresses.

Configuration

Figure 38: Example topology - no LAG shows the example topology with one MTU and three PEs.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 38: Example topology - no LAG

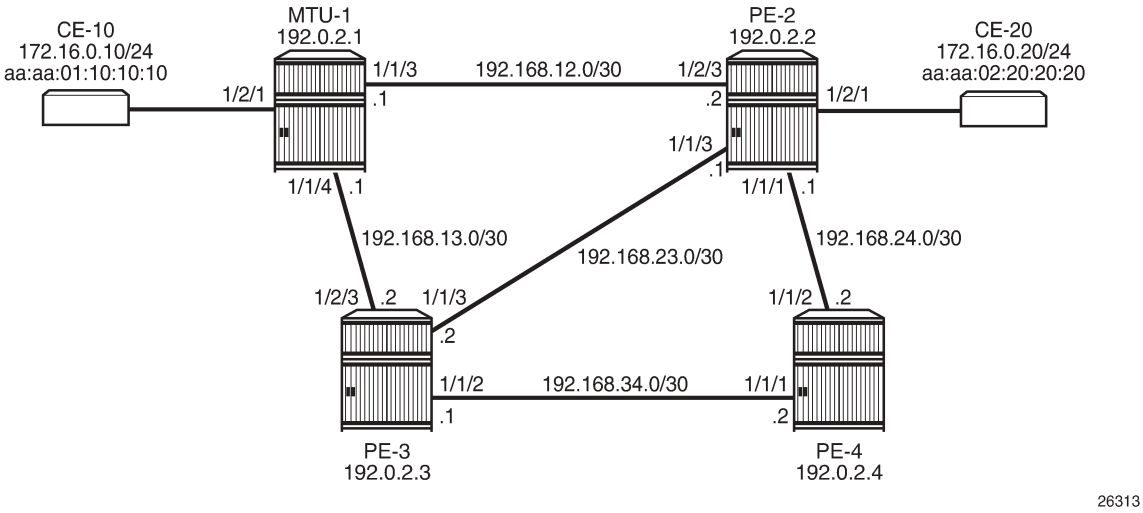

- Cards, MDAs
- The ports between the PEs are configured as network ports; the other ports are access ports. No LAG is configured initially.
- IGP (IS-IS is used in this example) between the PEs
- LDP between the PEs
- BGP with address family EVPN on the PEs

PE-2 is the BGP route reflector. The BGP configuration on the PEs is similar. The BGP configuration on the clients PE-3 and PE-4 is as follows:

```
# on PE-3, PE-4:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            split-horizon true
            rapid-update {
                evpn true
            }
            group "internal" {
                peer-as 64500
                family {
                    evpn true
                }
            }
            neighbor "192.0.2.2" {
                group "internal"
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

VPLS 1 is configured on all nodes. Initially, ALMP is disabled. On MTU-1, the VPLS 1 contains three SAPs: one toward CE-10, one toward PE-2, and one toward PE-3.

On PE-2, VPLS 1 is configured with EVPN MPLS and contains a SAP toward CE-20 and a SAP toward MTU-1, as follows:

```
# on PE-2:
configure {
    service {
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
            }
            bgp-evpn {
                evi 1
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            sap 1/2/1:1 {
            }
            sap 1/2/3:1 {
            }
        }
```

On PE-3, VPLS 1 is configured with EVPN MPLS and contains a SAP toward MTU-1, as follows:

```
# on PE-3:
configure {
    service {
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
            }
            bgp-evpn {
                evi 1
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            sap 1/2/3:1 {
            }
        }
```

The following use cases will be described in this section:

- EVPN MPLS without multi-homing.
  - Default behavior: no ALMP on SAPs, no protected MAC addresses
  - No ALMP on SAPs, RPS-DF on EVPN MPLS destinations

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

  - ALMP and implicit RPS-DF on SAPs.
    - RPS-DF on EVPN MPLS destinations, MAC first learned on PE-2
    - RPS-DF on EVPN MPLS destinations, MAC simultaneously learned on PE-2 and PE-3
    - No RPS-DF on EVPN MPLS destinations, MAC simultaneously learned on PE-2 and PE-3
  - ALMP and RPS on SAPs.
    - RPS-DF on EVPN MPLS destinations, MAC first learned on PE-2
    - RPS-DF on EVPN MPLS destinations, MAC simultaneously learned on PE-2 and PE-3
    - No RPS-DF on EVPN MPLS destinations, MAC simultaneously learned on PE-2 and PE-3
- EVPN MPLS with ALMP in all-active multi-homing.
  - RPS-DF on SAPs, RPS-DF on EVPN MPLS destinations

Default behavior: no protected MAC addresses

The following example is not a recommended configuration because it causes a loop. By default, ALMP is disabled and no static MAC addresses are configured. As described in chapter EVPN for VXLAN Tunnels (Layer 2), duplicate MAC addresses are detected in BGP EVPN and the MAC address will be put in a hold-down state on the EVPN destinations after a configurable threshold is reached. This applies to EVPN-MPLS as well as to EVPN-VXLAN. By default, the maximum number of MAC address moves is five in a time window of 3 minutes.

Figure 39: MAC address learned simultaneously on SAPs on PE-2 and PE-3 shows that the MAC address from CE-10 is learned simultaneously on the SAPs in VPLS 1 on PE-2 and PE-3.

Figure 39: MAC address learned simultaneously on SAPs on PE-2 and PE-3

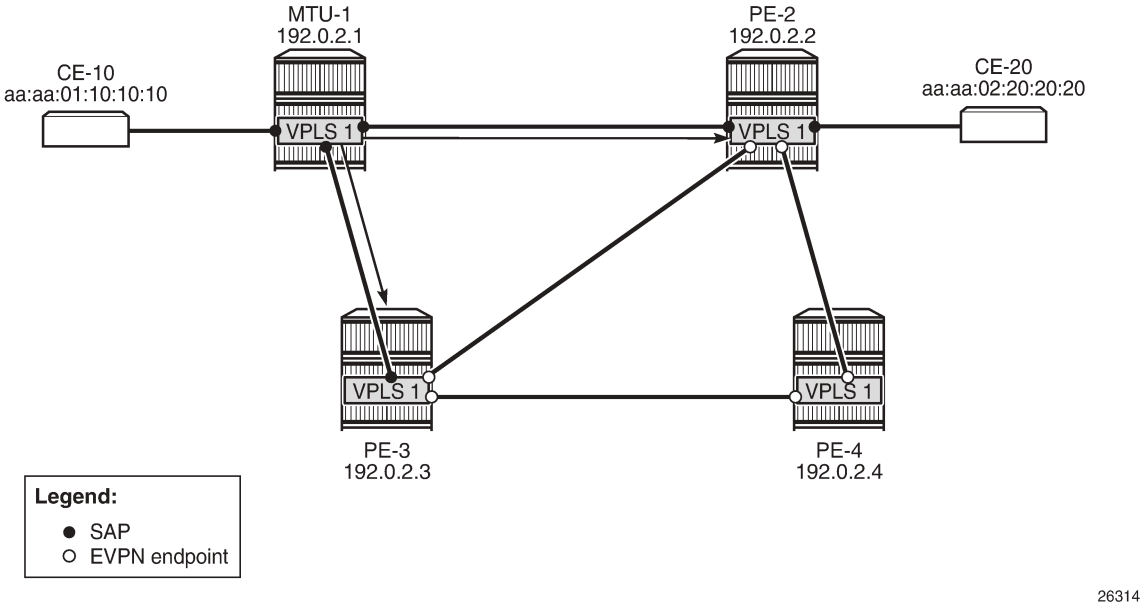

CE-10 sends frames to CE-20 with MAC Destination Address (DA) aa:aa:02:20:20:20. MTU-1 has not learned that MAC DA, so the frames are flooded to PE-2 and PE-3, where they enter the SAPs

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

simultaneously. PE-2 and PE-3 have not learned the MAC DA either, so the frames are flooded to all potential destinations. The frames received on PE-2 will be sent (among others) to PE-3, and vice versa. These frames are forwarded back out of the SAP toward MTU-1. This causes a loop.

Both PEs send a BGP update for the MAC SA aa:aa:01:10:10:10 to the other PEs with no sticky bit set. That MAC SA is learned, but not protected on the destination to the other PE. The stream of frames will cause the learned MAC SA to oscillate between the SAP and EVPN destinations on PE-2 and PE-3, and between the EVPN destinations on PE-4.

After a configurable number of BGP EVPN MAC address moves in a time span (by default, after five MAC address moves in a period of 3 minutes), the MAC address is put in a hold-down state on the EVPN destinations for a specific duration (until the next MAC address duplication detection retry; by default, after 9 minutes).

The following message in log 99 on PE-2 (and also on PE-3) indicates that duplicate MAC addresses have been detected:

```
77 2021/03/23 09:42:59.410 CET MINOR: SVCMGR #2331 Base
"VPLS Service 1 has MAC(s) detected as duplicates by EVPN mac-duplication detection."
```

The following shows the settings for EVPN MAC address duplication detection, which are the default. It also lists the detected duplicate MAC addresses of CE-10 and CE-20:

```
[/]
A:admin@PE-3# show service id 1 bgp-evpn

===============================================================================
BGP EVPN Table
===============================================================================
MAC Advertisement  : Enabled            Unknown MAC Route  : Disabled
CFM MAC Advertise  : Disabled
Creation Origin    : manual
MAC Dup Detn Moves : 5                  MAC Dup Detn Window: 3
MAC Dup Detn Retry : 9                  Number of Dup MACs : 2
MAC Dup Detn BH    : Disabled
IP Route Advert    : Disabled
Sel Mcast Advert   : Disabled
EVI                : 1
Ing Rep Inc McastAd: Enabled
Accept IVPLS Flush : Disabled

-------------------------------------------------------------------------------
Detected Duplicate MAC Addresses            Time Detected
-------------------------------------------------------------------------------
aa:aa:01:10:10:10                           03/23/2021 09:42:59
aa:aa:02:20:20:20                           03/23/2021 09:42:59
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN MPLS Information
===============================================================================
Admin Status       : Enabled            Bgp Instance       : 1
Force Vlan Fwding  : Disabled
Route NextHop Type : system-ipv4
Control Word       : Disabled
Max Ecmp Routes    : 1
Entropy Label      : Disabled
Default Route Tag  : none
Split Horizon Group: (Not Specified)
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Ingress Rep BUM Lbl: Enabled
Ingress Ucast Lbl  : 524284             Ingress Mcast Lbl  : 524283
RestProtSrcMacAct  : none
Evpn Mpls Encap    : Enabled            Evpn MplsoUdp      : Disabled
Oper Group         :
===============================================================================

===============================================================================
BGP EVPN MPLS Auto Bind Tunnel Information
===============================================================================
Allow-Flex-Algo-Fallback  : false
Resolution                : any                Strict Tnl Tag   : false
Max Ecmp Routes           : 1
Bgp Instance              : 1
Filter Tunnel Types       : (Not Specified)
===============================================================================
```

By default, there is no protected source MAC violation action on the EVPN destinations (**RestProtSrcMacAct : none**).

The MAC addresses are in a hold-down state on the EVPN destinations and no MAC address moves take place until the next MAC address duplication detection retry after 9 minutes. After 9 minutes, the EVPN MAC address duplication alarm is cleared, but after the next five MAC address moves within a time span of 3 minutes, the alarm is raised again and this threshold is reached soon after the alarm has been cleared.

The MAC addresses of both CEs are learned on the SAP of PE-3 (CE-20's MAC address is also learned on the SAP toward MTU-1), not on the EVPN destinations, because of the MAC address duplication detection and hold-down state in EVPN, as follows:

```
[/]
A:admin@PE-3# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
1          aa:aa:01:10:10:10 sap:1/2/3:1             L/0      03/23/21 09:42:59
1          aa:aa:02:20:20:20 sap:1/2/3:1             L/0      03/23/21 09:42:59
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

A similar output can be shown for PE-2.

Both PE-2 and PE-3 learn the MAC addresses locally and send BGP EVPN MAC address route updates to their BGP peers. PE-3 received the following BGP EVPN MAC address routes from PE-2, with the MAC address mobility sequence number representing the number of MAC address moves:

```
[/]
A:admin@PE-3# show router bgp routes evpn mac
===============================================================================
 BGP Router ID:192.0.2.3        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
BGP EVPN MAC Routes
===============================================================================
Flag  Route Dist.          MacAddr           ESI
      Tag                  Mac Mobility      Label1
                           Ip Address
                           NextHop
-------------------------------------------------------------------------------
u*>i  192.0.2.2:1          aa:aa:01:10:10:10 ESI-0
      0                    Seq:4             LABEL 524284
                           n/a
                           192.0.2.2

u*>i  192.0.2.2:1          aa:aa:02:20:20:20 ESI-0
      0                    Seq:4             LABEL 524284
                           n/a
                           192.0.2.2

-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

PE-3 does not use these BGP EVPN MAC address routes in its FDB, because locally learned MAC addresses are preferred.

The remote PE (PE-4) received the following BGP EVPN MAC routes from PE-2 and PE-3:

```
[/]
A:admin@PE-4# show router bgp routes evpn mac
===============================================================================
 BGP Router ID:192.0.2.4        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN MAC Routes
===============================================================================
Flag  Route Dist.          MacAddr           ESI
      Tag                  Mac Mobility      Label1
                           Ip Address
                           NextHop
-------------------------------------------------------------------------------
u*>i  192.0.2.2:1          aa:aa:01:10:10:10 ESI-0
      0                    Seq:4             LABEL 524284
                           n/a
                           192.0.2.2

u*>i  192.0.2.2:1          aa:aa:02:20:20:20 ESI-0
      0                    Seq:4             LABEL 524284
                           n/a
                           192.0.2.2

u*>i  192.0.2.3:1          aa:aa:01:10:10:10 ESI-0
      0                    Seq:3             LABEL 524284
                           n/a
                           192.0.2.3

u*>i  192.0.2.3:1          aa:aa:02:20:20:20 ESI-0
      0                    Seq:5             LABEL 524284
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                          n/a
                          192.0.2.3

-------------------------------------------------------------------------------
Routes : 4
===============================================================================
```

In the preceding output, MAC aa:aa:01:10:10:10 is learned from BGP peer 192.0.2.3 with MAC mobility sequence number 3, and from BGP peer 192.0.2.2 with sequence number 4. MAC aa:aa:02:20:20:20 is learned from BGP peer 192.0.2.2 with sequence number 4 and from BGP peer 192.0.2.3 with sequence number 5. The FDB for VPLS 1 on PE-4 contains the MAC addresses learned from BGP EVPN MAC updates with the highest MAC mobility sequence number, as follows:

```
[/]
A:admin@PE-4# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
1          aa:aa:01:10:10:10 mpls:                   Evpn     03/23/21 09:42:59
                             192.0.2.2:524284
           ldp:65537
1          aa:aa:02:20:20:20 mpls:                   Evpn     03/23/21 09:42:59
                             192.0.2.3:524284
           ldp:65538
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

VPLS 1 on MTU-1 does not have EVPN configured and no MAC address duplication detection mechanism implemented. The MAC address from CE-10 is last learned on the SAP toward PE-2 (it might equally have been the SAP toward PE-3) instead of the SAP toward CE-10, resulting from the loop, as follows:

```
[/]
A:admin@MTU-1# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
1          aa:aa:01:10:10:10 sap:1/1/4:1             L/0      03/23/21 09:45:56
1          aa:aa:02:20:20:20 sap:1/1/3:1             L/0      03/23/21 09:42:59
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

No ALMP on SAPs, RPS-DF on EVPN destinations

When there are no protected MAC addresses (ALMP is disabled and no static MAC addresses are configured), the behavior is as described earlier. RPS-DF discards frames with protected MAC addresses that were not learned on the object, but there are no protected MAC addresses, because ALMP is not configured. RPS-DF does not discard frames with MAC SAs that are not protected.

RPS-DF is enabled on EVPN destinations on all PEs, as follows:

```
# on PE-2, PE-3, PE-4:
configure {
    service {
        vpls "VPLS 1" {
            bgp-evpn {
                mpls 1 {
                    fdb {
                        protected-src-mac-violation-action discard
                    }
                }
            }
```

The state of RPS is now "discard-frame" instead of "none", as follows:

```
[/]
A:admin@PE-3# show service id 1 bgp-evpn | match RestProtSrcMacAct
RestProtSrcMacAct  : Discard-frame
```

It is also allowed to configure RPS (**protected-src-mac-violation-action sap-oper-down**) on the SAPs, but that does not change the behavior when ALMP is disabled and there are no protected MAC addresses. RPS will not bring down a SAP after receiving a frame with an unprotected MAC SA.

ALMP and implicit RPS-DF on SAPs

ALMP is enabled on the SAPs in PE-2 as follows:

```
# on PE-2:
configure {
    service {
        vpls "VPLS 1" {
            sap 1/2/1:1 {            # SAP toward CE-20
                fdb {
                    auto-learn-mac-protect true
                }
            }
            sap 1/2/3:1 {            # SAP toward MTU-1
                fdb {
                    auto-learn-mac-protect true
                }
            }
```

The configuration is similar on PE-3.

The following shows that ALMP is enabled on the SAP and that the default RPS-DF is used:

```
[/]
A:admin@PE-2# show service id 1 sap 1/2/3:1 detail

===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Service Access Points(SAP)
===============================================================================
Service Id         : 1
SAP                : 1/2/3:1                 Encap            : q-tag
Description        : (Not Specified)
Admin State        : Up                      Oper State       : Up
Flags              : None
---snip---

Restr MacUnpr Dst  : Disabled
Auto Learn Mac Prot: Enabled
ALMP Exclude List  : <none>
RestMacProtSrc Act : none (oper: Discard-frame)
---snip---
```

ALMP and RPS-DF on SAPs, RPS-DF on EVPN MPLS destinations, MAC first learned on PE-2

Initially, the SAP on PE-3 is disabled to ensure that the MAC address will first be learned on PE-2, then on PE-3, as follows:

```
# on PE-3:
configure {
    service {
        vpls "VPLS 1" {
            sap 1/2/3:1 {            # SAP toward MTU-1
                admin-state disable
            }
```

Each learned MAC address on the SAPs on PE-2 will be protected; therefore, a BGP update with the static/sticky bit set will be sent to the BGP EVPN peers. In this example, the MAC aa:aa:01:10:10:10 of CE-10 is learned first on SAP 1/2/3:1 on PE-2, and MAC aa:aa:02:20:20:20 is learned on SAP 1/2/1:1 on PE-2. Consequently, PE-2 sends BGP updates with the static/sticky bit set to PE-3 for both MAC aa:aa:01:10:10:10 and MAC aa:aa:02:20:20:20, as follows:

```
95 2021/03/23 09:55:45.486 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 89
    Flag: 0x90 Type: 14 Len: 44 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
        Type: EVPN-MAC Len: 33 RD: 192.0.2.2:1 ESI: ESI-0, tag: 0, mac len: 48
                       mac: aa:aa:01:10:10:10, IP len: 0, IP: NULL, label1: 8388544
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:1
        bgp-tunnel-encap:MPLS
        mac-mobility:Seq:0/Static
"
```

```
97 2021/03/23 09:55:45.486 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Send BGP UPDATE:
    Withdrawn Length = 0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    Total Path Attr Length = 89
    Flag: 0x90 Type: 14 Len: 44 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
        Type: EVPN-MAC Len: 33 RD: 192.0.2.2:1 ESI: ESI-0, tag: 0, mac len: 48
                       mac: aa:aa:02:20:20:20, IP len: 0, IP: NULL, label1: 8388544
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:1
        bgp-tunnel-encap:MPLS
        mac-mobility:Seq:0/Static
"
```

Note:
The MPLS label is label1 in the BGP update divided by 16 (2^4), as follows:

$$\frac{8388544}{16} = 524284$$

PE-2 sends similar BGP EVPN updates to peer PE-4.

After these BGP EVPN updates have been sent to PE-3 (and PE-4), the SAP on PE-3 is enabled again, as follows:

```
# on PE-3:
configure {
    service {
        vpls "VPLS 1" {
            sap 1/2/3:1 {            # SAP toward MTU-1
                admin-state enable
            }
```

The MAC addresses in the FDB on PE-2, where these MAC addresses are learned, get the indication "L" for learned and "P" for protected MAC address, as follows:

```
[/]
A:admin@PE-2# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
1          aa:aa:01:10:10:10 sap:1/2/3:1             LP/60    03/23/21 09:55:45
1          aa:aa:02:20:20:20 sap:1/2/1:1             LP/60    03/23/21 09:55:45
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The MAC addresses in the FDB on PE-3 are learned from the BGP EVPN updates and get the indication "S" for static (sticky bit) and "P" for protected MAC address, as follows

```
[/]
A:admin@PE-3# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
1          aa:aa:01:10:10:10 mpls:                   EvpnS:P  03/23/21 09:55:45
                             192.0.2.2:524284
           ldp:65537
1          aa:aa:02:20:20:20 mpls:                   EvpnS:P  03/23/21 09:55:45
                             192.0.2.2:524284
           ldp:65537
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

The FDB on the remote PE (PE-4) looks similar, as follows:

```
[/]
A:admin@PE-4# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
1          aa:aa:01:10:10:10 mpls:                   EvpnS:P  03/23/21 09:55:45
                             192.0.2.2:524284
           ldp:65537
1          aa:aa:02:20:20:20 mpls:                   EvpnS:P  03/23/21 09:55:45
                             192.0.2.2:524284
           ldp:65537
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

The BGP EVPN MAC address routes on PE-3 have MAC address mobility equal to "Static", as follows:

```
[/]
A:admin@PE-3# show router bgp routes evpn mac
===============================================================================
 BGP Router ID:192.0.2.3        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN MAC Routes
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Flag  Route Dist.          MacAddr           ESI
      Tag                  Mac Mobility      Label1
                           Ip Address
                           NextHop
-------------------------------------------------------------------------------
u*>i  192.0.2.2:1          aa:aa:01:10:10:10 ESI-0
      0                    Static            LABEL 524284
                           n/a
                           192.0.2.2

u*>i  192.0.2.2:1          aa:aa:02:20:20:20 ESI-0
      0                    Static            LABEL 524284
                           n/a
                           192.0.2.2

-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

The BGP EVPN MAC routes on PE-4 are similar.

When a stream of frames with MAC SA aa:aa:01:10:10:10 enters the SAP on PE-3, these frames will be dropped by this SAP because of the implicit RPS-DF behavior in the SAP for protected MAC addresses, as shown in Figure 40: Default RPS-DF on SAPs - MAC learned and protected on SAP on PE-2.

Figure 40: Default RPS-DF on SAPs - MAC learned and protected on SAP on PE-2

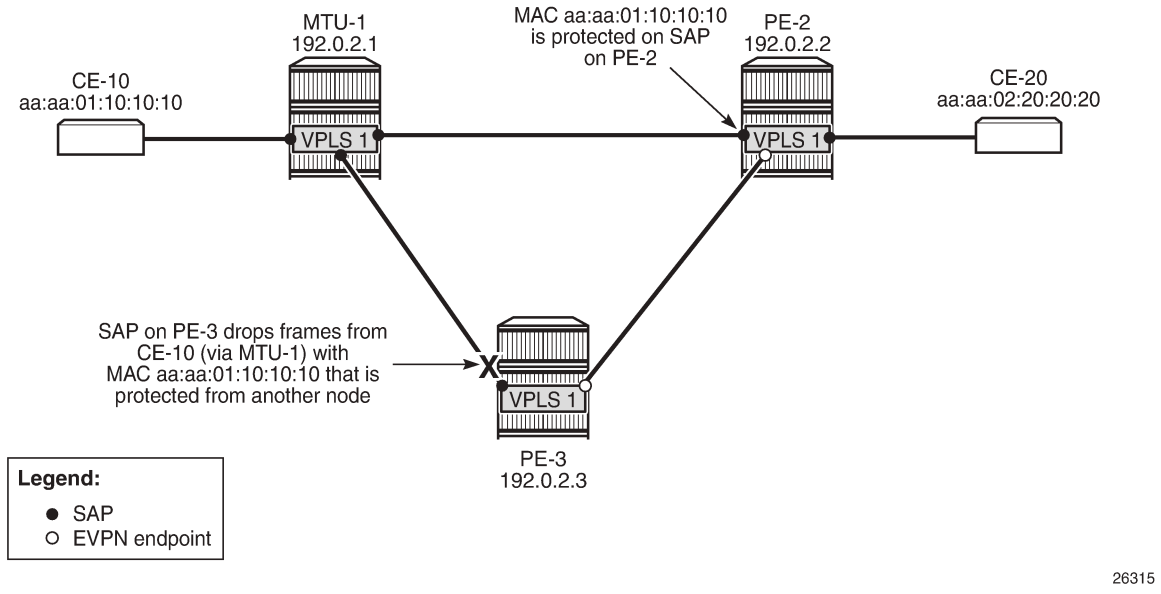

Because the MAC address was protected on the SAP on PE-2 and the BGP EVPN MAC route update had been received by PE-3 before any frame was received with this MAC SA, there will be no temporary loop. The frames with the protected MAC SA will be discarded at the SAP on PE-3, not on the EVPN MPLS destination on PE-2. In this case, there is no need to configure RPS-DF on the EVPN MPLS destinations, but it will make a difference when the MAC address is learned on both SAPs simultaneously.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

ALMP and RPS-DF on SAPs, RPS-DF on EVPN MPLS destinations, MAC simultaneously learned on PE-2 and PE-3

In the preceding example, the MAC addresses of CE-10 and CE-20 were first learned and protected on PE-2 and received on PE-3's SAP after the BGP update with static/sticky bit was received by PE-3. However, when the MAC address of CE-10 is learned simultaneously on both PEs, for example, because the MAC DA aa:aa:02:20:20:20 is unknown, there is a temporary loop until the MAC addresses are protected. Initially, the frames enter a SAP, are forwarded to the EVPN peer, and forwarded out of the remote SAP.

After the MAC addresses are learned and protected on the SAPs on both PEs, new frames received on a SAP with the protected MAC address will be sent to the other PE. However, they will be discarded due to RPS-DF on destination, as shown in Figure 41: MAC learned and protected simultaneously on PEs - RPS-DF on EVPN endpoints, because the destination PE has that same MAC address protected on its local SAP. This prevents a loop. BGP updates with the static/sticky bit set are sent to the BGP EVPN peer, but the locally learned and protected MAC address is preferred to the MAC address in a BGP update. Therefore, the FDB contains the locally learned MAC address aa:aa:01:10:10:10, not the BGP EVPN MAC address update for MAC address aa:aa:01:10:10:10.

Figure 41: MAC learned and protected simultaneously on PEs - RPS-DF on EVPN endpoints

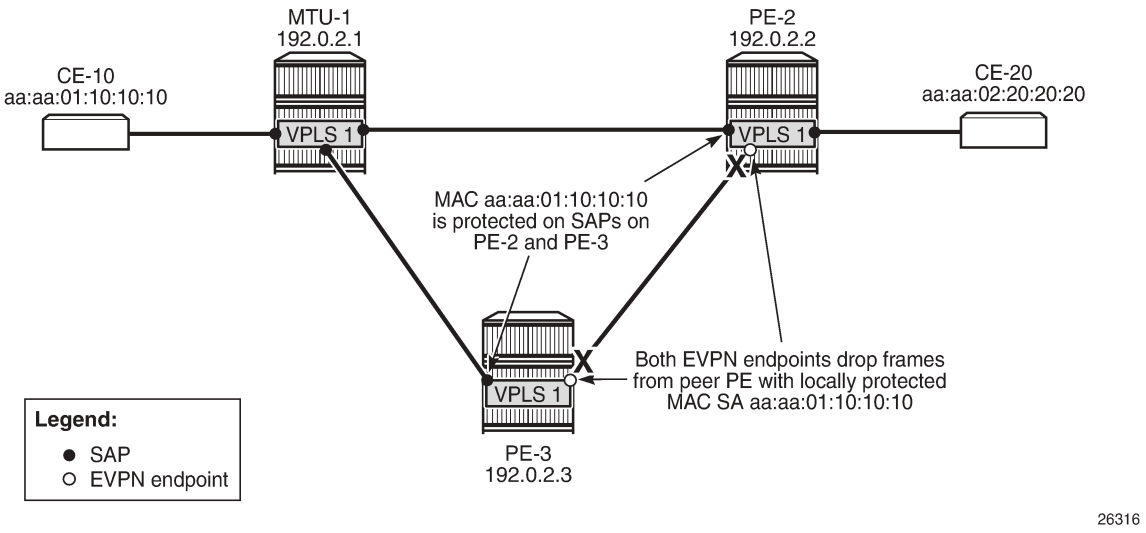

The MAC addresses of the CEs are cleared from the FDBs on all nodes, as follows:

```
clear service id 1 fdb mac aa:aa:01:10:10:10
clear service id 1 fdb mac aa:aa:02:20:20:20
```

This clear command for the FDB only works for auto-learned MAC addresses, not for BGP EVPN MAC address updates. BGP EVPN MAC address withdraw updates need to be sent. In this example, BGP is configured with **rapid-update evpn**, as shown previously.

When traffic is sent from CE-10 to CE-20, MAC address aa:aa:01:10:10:10 of CE-10 is learned simultaneously on SAP 1/2/3:1 in PE-2 and PE-3 and protected on both SAPs. MAC address aa:aa:02:20:20:20 is, in this case, first learned via MAC address learning on PE-2 and advertised via a BGP EVPN MAC address route update. However, it might happen that it was learned and protected on the SAP on PE-3 first, before the MAC address was learned and protected on PE-2 and the BGP EVPN

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

MAC address route update sent by PE-2 was received at PE-3. In the latter case, both MAC address aa:aa:01:10:10:10 and MAC address aa:aa:02:20:20:20 are learned and protected on the SAPs on both PE-2 and PE-3, and RPS-DF on the EVPN-MPLS destinations prevents loops.

However, in the present case, MAC address aa:aa:02:20:20:20 is only protected on the SAP on PE-2, because PE-3 received the EVPN MAC address update before it received a frame with MAC SA aa:aa:02:20:20:20. Therefore, the SAP on PE-3 will discard any frames with MAC SA aa:aa:02:20:20:20.

The FDB for VPLS 1 on PE-2 shows that both MAC addresses are learned locally and protected, as follows:

```
[/]
A:admin@PE-2# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
1          aa:aa:01:10:10:10 sap:1/2/3:1             LP/0     03/23/21 09:59:44
1          aa:aa:02:20:20:20 sap:1/2/1:1             LP/0     03/23/21 09:59:44
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

The FDB for VPLS 1 on PE-3 shows that MAC address aa:aa:01:10:10:10 is learned and protected locally, but MAC address aa:aa:02:20:20:20 is protected on PE-2, which has been advertised by PE-2 in a BGP EVPN MAC update, as follows:

```
[/]
A:admin@PE-3# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
1          aa:aa:01:10:10:10 sap:1/2/3:1             LP/0     03/23/21 09:59:44
1          aa:aa:02:20:20:20 mpls:                   EvpnS:P  03/23/21 09:59:44
                             192.0.2.2:524284
           ldp:65537
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

Both PE-2 and PE-3 send BGP EVPN MAC updates to their BGP peers for each locally learned and protected MAC address. The following BGP EVPN MAC update is sent by PE-2 to PE-3 for MAC address aa:aa:01:10:10:10:

```
# on PE-2:
103 2021/03/23 09:59:44.284 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Send BGP UPDATE:
    Withdrawn Length = 0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    Total Path Attr Length = 89
    Flag: 0x90 Type: 14 Len: 44 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
        Type: EVPN-MAC Len: 33 RD: 192.0.2.2:1 ESI: ESI-0, tag: 0, mac len: 48
                       mac: aa:aa:01:10:10:10, IP len: 0, IP: NULL, label1: 8388544
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:1
        bgp-tunnel-encap:MPLS
        mac-mobility:Seq:0/Static
"
```

Similar BGP EVPN updates are sent to the remote PE (PE-4). The FDB for VPLS 1 on PE-4 only contains entries learned from BGP EVPN updates, as follows:

```
[/]
A:admin@PE-4# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
1          aa:aa:01:10:10:10 mpls:                   EvpnS:P  03/23/21 09:59:44
                             192.0.2.2:524284
           ldp:65537
1          aa:aa:02:20:20:20 mpls:                   EvpnS:P  03/23/21 09:59:44
                             192.0.2.2:524284
           ldp:65537
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

PE-4 received BGP EVPN MAC address route updates from PE-2 and PE-3, but only installs the MAC address routes to PE-2 in its FDB, based on the lowest next-hop IP of the EVPN NLRI (192.0.2.2).

ALMP and RPS-DF on SAPs, no RPS-DF on EVPN MPLS destinations, MAC simultaneously learned on PE-2 and PE-3

RPS-DF is disabled on the EVPN MPLS destinations on the PEs, as follows:

```
# on PE-2, PE-3, PE-4:
configure {
    service {
        vpls "VPLS 1" {
            bgp-evpn {
                mpls 1 {
                    fdb {
                        delete protected-src-mac-violation-action
                    }
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

When a frame is received at SAP 1/2/3:1 on PE-3 with protected MAC SA aa:aa:01:10:10:10, it is not dropped by the SAP, because this MAC SA has been learned and protected on this SAP on PE-3. The frame is forwarded to PE-2 where it will not be discarded by the EVPN MPLS destination because RPS-DF is disabled. The frame will be forwarded to other objects in the VPLS in PE-2. For BUM traffic, there will be a loop, because all frames will be flooded to all objects in VPLS 1 on PE-2, including the SAP toward MTU-1.

ALMP and RPS on SAPs

When ALMP is enabled on an object, the default behavior is that frames with a protected MAC SA are discarded (RPS-DF). However, it is possible to configure RPS with **sap-oper-down** (or **sdp-bind-oper-down**), in this case on the SAPs on PE-2 and PE-3, as follows:

```
# on PE-2, PE-3:
configure {
    service {
        vpls "VPLS 1" {
            sap 1/2/3:1 {
                fdb {
                    protected-src-mac-violation-action sap-oper-down
                }
            }
```

Instead of discarding frames with MAC SAs that are protected on another object or node, the entire object (here: SAP) can be brought operationally down after a frame has been received with a MAC SA that is protected on another node.

The RPS configuration on the SAP can be shown as follows. The SAP has not been brought down yet.

```
[/]
A:admin@PE-2# show service id 1 sap "1/2/3:1" detail

===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 1
SAP                : 1/2/3:1                  Encap             : q-tag
Description        : (Not Specified)
Admin State        : Up                       Oper State        : Up
Flags              : None
---snip---

Restr MacUnpr Dst  : Disabled
Auto Learn Mac Prot: Enabled
ALMP Exclude List  : <none>
RestMacProtSrc Act : SAP-oper-down
---snip---
```

The **RestMacProtSrc Act** parameter is set to **SAP-oper-down**, meaning that RPS is configured, which causes the system to bring down the SAP when a duplicate MAC address is received that is protected on another object or node. When a SAP is brought down because of this, the **RxProtSrcMAC** flag will be raised and can be shown in the detailed SAP show output.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

ALMP and RPS on SAPs, RPS-DF on EVPN MPLS destinations, MAC first learned on PE-2

RPS-DF is enabled on the EVPN MPLS destinations on the PEs, as follows:

```
# on PE-2, PE-3, PE-4:
configure {
    service {
        vpls "VPLS 1" {
            bgp-evpn {
                mpls 1 {
                    fdb {
                        protected-src-mac-violation-action discard
                    }
                }
```

To simulate a scenario where the MAC addresses are first learned on PE-2, the SAP on PE-3 is disabled until the BGP EVPN MAC route updates are sent, as follows:

```
# on PE-3:
configure {
    service {
        vpls "VPLS 1" {
            sap 1/2/3:1 {
                admin-state disable
            }
```

The FDBs are cleared on the nodes, as follows:

```
clear service id 1 fdb mac aa:aa:01:10:10:10
clear service id 1 fdb mac aa:aa:02:20:20:20
```

Traffic is sent between CE-10 and CE-20, and the MAC addresses are learned and protected on the SAP on PE-2, as follows:

```
[/]
A:admin@PE-2# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
1          aa:aa:01:10:10:10 sap:1/2/3:1             LP/0     03/23/21 10:04:09
1          aa:aa:02:20:20:20 sap:1/2/1:1             LP/0     03/23/21 10:04:09
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

No MAC learning took place on the SAP on PE-3, and the FDB contains the MAC addresses from the BGP EVPN updates, as follows:

```
[/]
A:admin@PE-3# show service id 1 fdb detail

===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
1          aa:aa:01:10:10:10 mpls:                   EvpnS:P  03/23/21 10:04:09
                             192.0.2.2:524284
           ldp:65537
1          aa:aa:02:20:20:20 mpls:                   EvpnS:P  03/23/21 10:04:09
                             192.0.2.2:524284
           ldp:65537
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

The SAP on PE-3 is enabled, as follows:

```
# on PE-3:
configure {
    service {
        vpls "VPLS 1" {
            sap 1/2/3:1 {
                admin-state enable
            }
```

The operational state of the SAP is up, because no protected MAC addresses have been received yet:

```
[/]
A:admin@PE-3# show service id 1 sap

===============================================================================
SAP(Summary), Service 1
===============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.  Egr.   Adm  Opr
                                           QoS   Fltr    QoS   Fltr
-------------------------------------------------------------------------------
1/2/3:1                         1          1     none    1     none   Up   Up
-------------------------------------------------------------------------------
Number of SAPs : 1
-------------------------------------------------------------------------------
===============================================================================
```

The FDB is cleared for MAC address aa:aa:02:20:20:20 on MTU-1, as follows:

```
# on MTU-1:
clear service id 1 fdb mac aa:aa:02:20:20:20
```

Traffic from CE-10 toward the unknown MAC address aa:aa:02:20:20:20 reaches the SAPs on PE-2 and PE-3. When MAC SA aa:aa:01:10:10:10, which is protected on PE-2, is received on PE-3, SAP 1/2/3:1 will be brought operationally down, as shown in Figure 42: MAC learned and protected on SAP on PE-2 - RPS enabled on SAP on PE-3, and the following alarms will be raised in log 99:

```
89 2021/03/23 10:05:32.247 CET MINOR: SVCMGR #2208 Base
"Protected MAC aa:aa:01:10:10:10 received on SAP 1/2/3:1 in service 1. The SAP will be
 disabled."

90 2021/03/23 10:05:32.247 CET MINOR: SVCMGR #2203 Base
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
"Status of SAP 1/2/3:1 in service 1 (customer 1) changed to admin=up oper=down flags=RxProtSrc
Mac "
```

Figure 42: MAC learned and protected on SAP on PE-2 - RPS enabled on SAP on PE-3

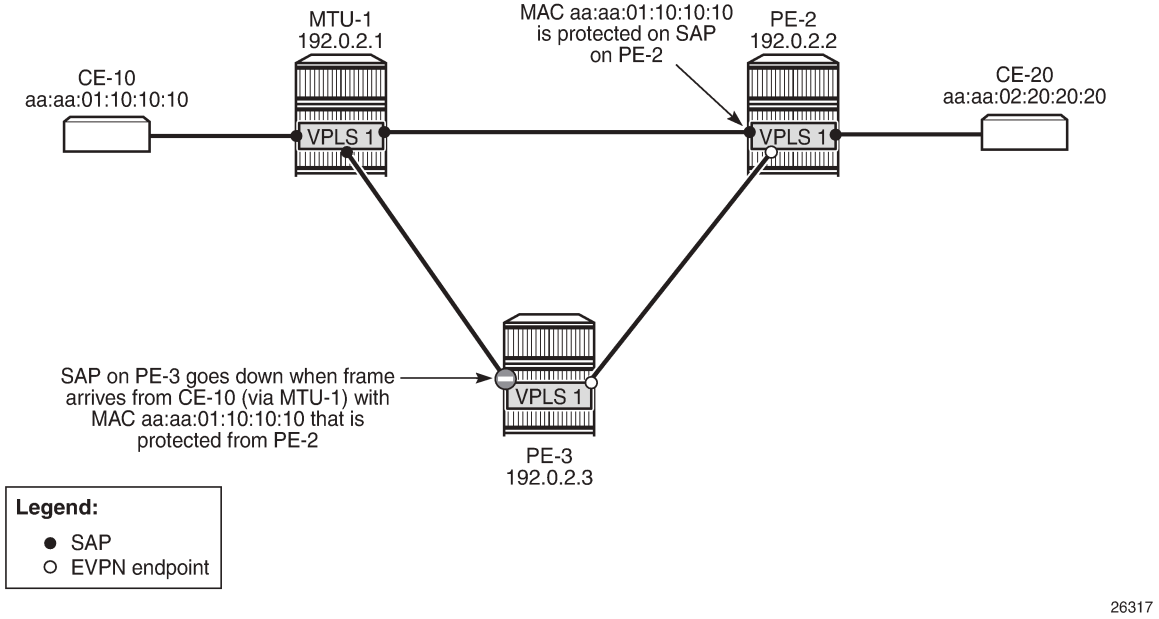

The operational state of SAP 1/2/3:1 is now down with flag **RxProtSrcMAC**, indicating that a duplicate MAC address that is protected on a remote node has been received, as follows:

```
[/]
A:admin@PE-3# show service id 1 sap 1/2/3:1

===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 1
SAP                : 1/2/3:1                  Encap             : q-tag
Description        : (Not Specified)
Admin State        : Up                       Oper State        : Down
Flags              : RxProtSrcMac
Multi Svc Site     : None
Last Status Change : 03/23/2021 10:05:32
Last Mgmt Change   : 03/23/2021 10:04:55
===============================================================================
```

The SAP is operationally down and will not come up automatically when the FDB is cleared. To bring the SAP up, an operator needs to disable and re-enable the SAP, as follows:

```
# on PE-3:
configure exclusive
    service {
        vpls "VPLS 1" {
            sap 1/2/3:1 {
                admin-state disable
                commit
                admin-state enable
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                commit
```

```
[/]
A:admin@PE-3# show service id 1 sap

===============================================================================
SAP(Summary), Service 1
===============================================================================
PortId                          SvcId      Ing.   Ing.    Egr.   Egr.   Adm   Opr
                                           QoS    Fltr    QoS    Fltr
-------------------------------------------------------------------------------
1/2/3:1                         1          1      none    1      none   Up    Up
-------------------------------------------------------------------------------
Number of SAPs : 1
-------------------------------------------------------------------------------
===============================================================================
```

ALMP and RPS on SAPs, RPS-DF on EVPN MPLS destinations, MAC simultaneously learned on PE-2 and PE-3

When CE-10 sends traffic to CE-20 and the destination MAC address is unknown, MAC address aa:aa:01:10:10:10 is simultaneously learned and protected on PE-2 and PE-3. No SAP will be brought down when MAC address aa:aa:01:10:10:10 is received on PE-2 or PE-3. This scenario is identical to the one with ALMP and (default) RPS-DF on the SAPs, as shown in Figure 43: RPS enabled on SAPs - RPS-DF on EVPN endpoints, MACs learned simultaneously (which is identical to Figure 41: MAC learned and protected simultaneously on PEs - RPS-DF on EVPN endpoints).

Figure 43: RPS enabled on SAPs - RPS-DF on EVPN endpoints, MACs learned simultaneously

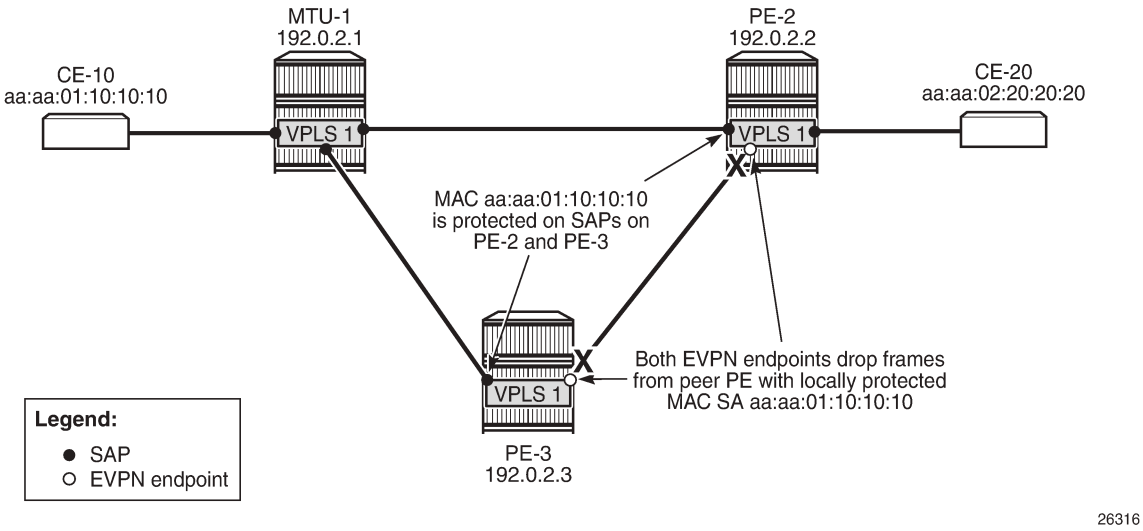

A temporary loop is possible until the MAC address is protected on the SAPs. Initially, the frames enter the SAP, are forwarded to the other PEs, and are forwarded out of the other SAP (unless the MAC address is protected). When the MAC address is protected, any other frames received on the SAP will be sent to the other PE (for example, from PE-3 to PE-2, or vice versa), but they will be discarded by the receiving PE, because RPS-DF is applied on the EVPN destination. BGP EVPN updates are sent to the peer PEs

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

with the sticky bit set. This MAC route will not be installed in the FDB of PE-2 and PE-3 because the MAC address has already been learned locally, which has a higher preference.

The FDB on PE-2 contains locally learned and protected MAC addresses, as follows:

```
[/]
A:admin@PE-2# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
1          aa:aa:01:10:10:10 sap:1/2/3:1             LP/0     03/23/21 10:09:54
1          aa:aa:02:20:20:20 sap:1/2/1:1             LP/0     03/23/21 10:09:54
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

The FDB on PE-3 contains MAC address aa:aa:01:10:10:10 that is locally learned and protected, and MAC address aa:aa:02:20:20:20 that is protected on PE-2, as follows:

```
[/]
A:admin@PE-3# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
1          aa:aa:01:10:10:10 sap:1/2/3:1             LP/0     03/23/21 10:09:54
1          aa:aa:02:20:20:20 mpls:                   EvpnS:P  03/23/21 10:09:54
                             192.0.2.2:524284
           ldp:65537
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

SAP 1/2/3:1 will not be brought down if frames are received with MAC address aa:aa:01:10:10:10 that is locally learned and protected. However, MAC address aa:aa:02:20:20:20 was learned and protected first on PE-2 and the BGP update was received by PE-3 before the MAC address was received on PE-3. Therefore, MAC address aa:aa:02:20:20:20 will not be learned and protected on PE-3 and, if frames with a MAC SA aa:aa:02:20:20:20 were received on SAP 1/2/3:1 on PE-3, the SAP would be brought down.

ALMP and RPS on SAPs, no RPS-DF on EVPN MPLS destinations, MAC simultaneously learned on PE-2 and PE-3

RPS-DF is disabled on the EVPN MPLS destinations on the PEs, as follows:

```
# on PE-2, PE-3, PE-4:
configure {
    service {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        vpls "VPLS 1" {
            bgp-evpn {
                mpls 1 {
                    fdb {
                        delete protected-src-mac-violation-action
                    }
                }
```

When frames are received at SAP 1/2/3:1 on PE-3 with protected MAC SA aa:aa:01:10:10:10, the SAP is not brought down, because this MAC SA has been learned and protected on this SAP. The frame is forwarded to PE-2 where it will not be discarded by the EVPN MPLS destination because RPS-DF is disabled. It will be forwarded to other objects in the VPLS. For BUM traffic, there will be a loop, because the frames will be flooded to all objects, including the SAP on PE-2 toward MTU-1.

ALMP in all-active multi-homing SAPs

All-active multi-homing for EVPN MPLS is explained in chapter EVPN for MPLS Tunnels. ALMP is not required on all-active multi-homing SAPs. The following example shows that traffic can be dropped when ALMP is enabled on the SAPs and RPS-DF is enabled on the EVPN-MPLS destinations.

Figure 44: ALMP in all-active multi-homing SAPs shows the example topology for all-active multi-homing.

Figure 44: ALMP in all-active multi-homing SAPs

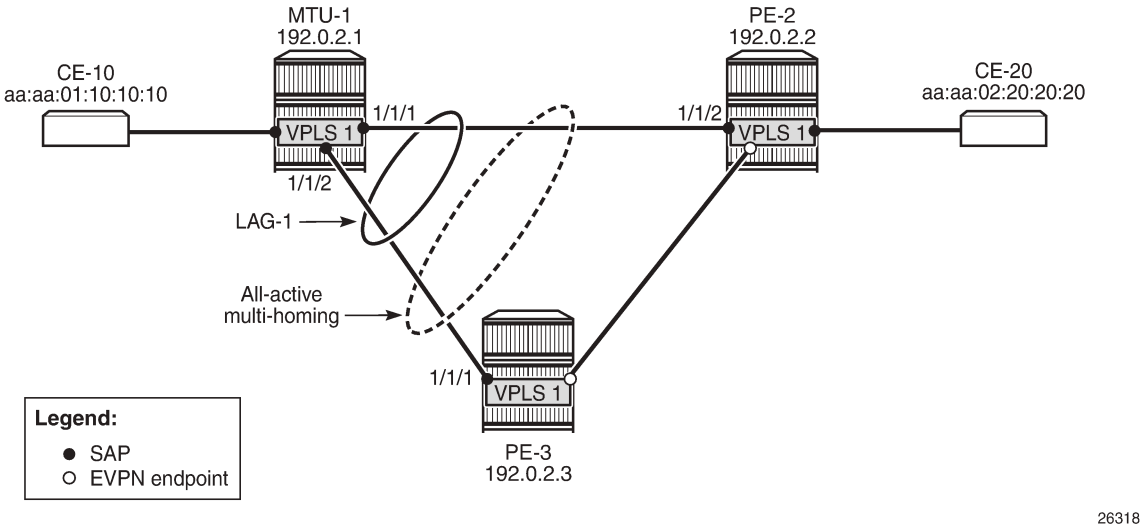

VPLS is configured with SAP lag-1:1 on the three nodes in the topology, as follows:

```
# on MTU-1, PE-2, PE-3:
configure {
    service {
         vpls "VPLS 1" {
            sap lag-1:1 {
            }
```

The SAPs used in the preceding scenarios are removed.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

All-active multi-homing is configured on PE-2 and PE-3, as follows:

```
# on PE-2, PE-3:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "ESI-12" {
                        admin-state enable
                        esi 01:00:00:00:00:23:00:00:00:01
                        multi-homing-mode all-active
                        df-election {
                            es-activation-timer 3
                        }
                        association {
                            lag "lag-1" {
                            }
                        }
                    }
                }
```

ALMP is enabled on the SAPs on PE-2 and PE-3, as follows:

```
# on PE-2, PE-3:
configure {
    service {
        vpls "VPLS 1" {
            sap lag-1:1 {
                fdb {
                    auto-learn-mac-protect true
                }
            }
```

MAC address aa:aa:01:10:10:10 is learned and protected on PE-2 and PE-3, as follows:

```
[/]
A:admin@PE-2# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
1          aa:aa:01:10:10:10 sap:lag-1:1             EvpnS:P  03/23/21 10:11:51
1          aa:aa:02:20:20:20 sap:1/2/1:1             LP/90    03/23/21 10:09:54
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

```
[/]
A:admin@PE-3# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
1          aa:aa:01:10:10:10 sap:lag-1:1             LP/0     03/23/21 10:11:51
1          aa:aa:02:20:20:20 mpls:                   EvpnS:P  03/23/21 10:09:54
                             192.0.2.2:524284
           ldp:65537
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

ALMP in all-active multi-homing, RPS-DF on EVPN MPLS destinations

ALMP is not recommended in all-active multi-homing because it can cause traffic loss. The following example shows when frames are dropped.

Figure 45: All-active multi-homing - RPS-DF on SAPs and EVPN endpoints shows the example setup with MAC address aa:aa:01:10:10:10 protected on SAP lag-1:1 on both PE-2 and PE-3, and RPS-DF enabled on the EVPN endpoints.

Figure 45: All-active multi-homing - RPS-DF on SAPs and EVPN endpoints

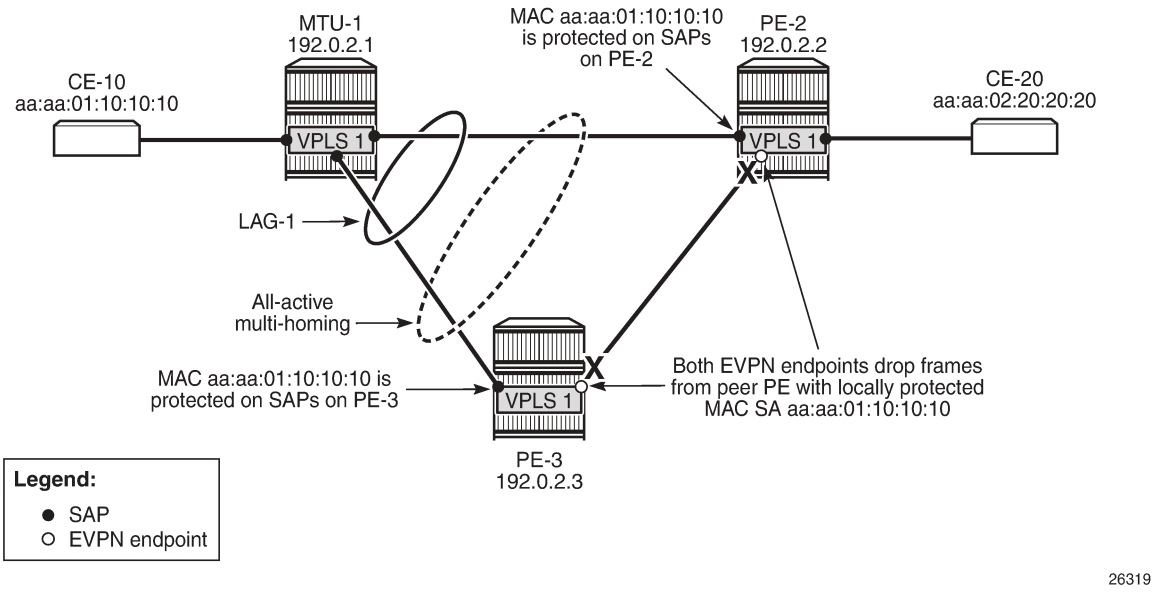

When frames with MAC address aa:aa:01:10:10:10 are sent between PE-2 and PE-3, these frames will be dropped by the EVPN MPLS destination that has RPS-DF enabled.

The traffic flows from CE-10 and CE-20 are hashed over both links in the LAG. When the frames are sent out on MTU-1 on port 1/1/1 toward PE-2, the traffic reaches CE-20, and traffic can be sent back from CE-20 to CE-10 via the direct link between PE-2 and MTU-1. However, when traffic is sent out from MTU-1 on port 1/1/2 toward PE-3, the frames will be forwarded from PE-3 to PE-2, where they will be discarded at the EVPN MPLS destination on PE-2 because of RPS-DF. No traffic flow is possible for frames with the protected MAC SA aa:aa:01:10:10:10 via PE-3 to PE-2, or vice versa. If the MAC address is not protected yet on PE-2, the first few messages get through until the MAC address is protected on PE-2. Both multi-homing PEs, PE-2 and PE-3, protect the MAC address aa:aa:01:10:10:10 on their local all-active SAP.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Therefore, PE-2 discards all frames with the MAC SA aa:aa:01:10:10:10 when they are received on the EVPN MPLS destination from the other multi-homing PE (PE-3).

An improved mechanism for EVPN loop protection in all-active multi-homing is black-hole MAC duplication, as described in chapter Black-hole MAC for EVPN Loop Protection.

For single-active multi-homing, this problem does not arise: only the designated forwarder in the Ethernet segment receives and forwards traffic. Therefore, the CE MAC addresses will not be learned and protected on different PEs in the same Ethernet segment.

Conclusion

For security, MAC addresses learned on objects, such as SAPs, spoke/mesh-SDPs, and SHGs in EVPN services can be protected and advertised by BGP with the sticky bit set. By default, frames with a protected MAC SA are discarded if received on objects where the MAC address was not learned. Objects can be configured to be brought operationally down when a frame is received with a protected MAC SA that has not been learned locally.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

BGP Multi-Homing for VPLS Networks

This chapter describes BGP Multi-Homing (BGP-MH) for VPLS network configurations.

Topics in this chapter include:

Applicability

Initially, the information in this chapter was based on SR OS Release 8.0.R5, with additions for SR OS Release 9.0.R1. The MD-CLI in the current edition corresponds to SR OS Release 20.10.R2.

Overview

SR OS supports the use of Border Gateway Protocol Multi-Homing for VPLS (hereafter called BGP-MH). BGP-MH is described in *draft-ietf-bess-vpls-multihoming, BGP based Multi-homing in Virtual Private LAN Service*, and provides a network-based resiliency mechanism (no interaction from the Provider Edge routers (PEs) to Multi-Tenant Units/Customer Equipment (MTU/CE)) that can be applied on service access points (SAPs) or network (pseudowires) topologies. The BGP-MH procedures will run between the PEs and will provide a loop-free topology from the network perspective (only one logical active path will be provided per VPLS among all the objects SAPs or pseudowires which are part of the same Multi-Homing site).

Each multi-homing site connected to two or more peers is represented by a site ID (2 bytes long) which is encoded in the BGP MH Network Layer Reachability Information (NLRI). The BGP peer holding the active path for a particular multi-homing site will be named as the Designated Forwarder (DF), whereas the rest of the BGP peers participating in the BGP MH process for that site will be named as non-DF and will block the traffic (in both directions) for all the objects belonging to that multi-homing site.

BGP MH uses the following rules to determine which PE is the DF for a particular multi-homing site:

1. A BGP MH NLRI with D flag = 0 (multi-homing object up) always takes precedence over a BGP MH NLRI with D flag = 1 (multi-homing object down). If there is a tie, then:
2. The BGP MH NLRI with the highest BGP Local Preference (LP) wins. If there is a tie, then:
3. The BGP MH NLRI issued from the PE with the lowest PE ID (system address) wins.

The main advantages of using BGP-MH as opposed to other resiliency mechanisms for VPLS are:

- Flexibility: BGP-MH uses a common mechanism for access and core resiliency. The designer has the flexibility of using BGP-MH to control the active/standby status of SAPs, spoke SDPs, Split Horizon Groups (SHGs) or even mesh SDP bindings.
- The standard protocol is based on BGP, a standard, scalable, and well-known protocol.
- Specific benefits at the access:

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

  - It is network-based, independent of the customer CE and, therefore, it does not need any customer interaction to determine the active path. Consequently, the operator will spend less effort on provisioning and will minimize both operation costs and security risks (in particular, this removes the requirement for spanning tree interaction between the PE and CE).
  - Easy load balancing per service (no service fate-sharing) on physical links.
- Specific benefits in the core:
  - It is a network-based mechanism, independent of the MTU resiliency capabilities and it does not need MTU interaction, therefore operational advantages are achieved as a result of the use of BGP-MH: less provisioning is required and there will be minimal risks of loops. In addition, simpler MTUs can be used.
  - Easy load balancing per service (no service fate-sharing) on physical links.
  - Less control plane overhead: there is no need for an additional protocol running the pseudowire redundancy when BGP is already used in the core of the network. BGP-MH just adds a separate NLRI in the L2-VPN family (AFI=25, SAFI=65).

This chapter describes how to configure and troubleshoot BGP-MH for VPLS

Knowledge of the LDP/BGP VPLS (RFC 4762, *Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling*, and RFC 4761, *Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling*) architecture and functionality is assumed throughout this chapter. For further information, see the relevant Nokia documentation.

Figure 46: Example topology shows the example topology that will be used throughout the rest of the chapter.

The initial configuration includes:

- IGP — IS-IS, Level 2 on all routers; area 49.0001
- RSVP-TE for transport tunnels
- Fast reroute (FRR) protection in the core; no FRR protection at the access.

Figure 46: Example topology

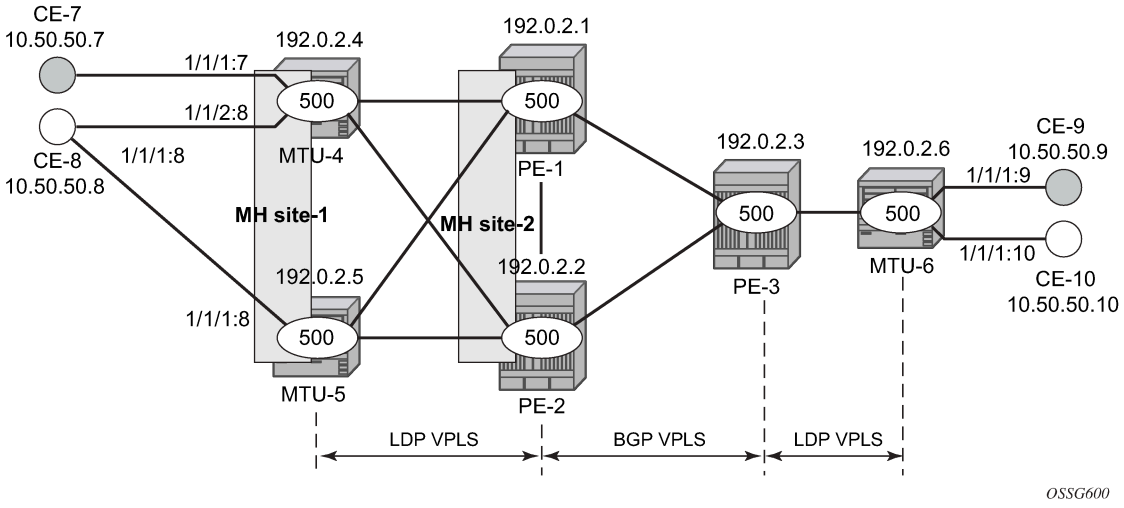

The topology consists of three core nodes (PE-1, PE-2, and PE-3) and three MTUs connected to the core.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The VPLS service VPLS-500 is configured on the six nodes with the following characteristics:

- The core VPLS instances are connected by a full mesh of BGP-signaled pseudowires (that is, pseudowires among PE-1, PE-2, and PE-3 will be signaled by BGP VPLS).
- As shown in Figure 46: Example topology, the MTUs are connected to the BGP VPLS core by T-LDP pseudowires. MTU-6 is connected to PE-3 by a single pseudowire, whereas MTU-4 and MTU-5 are dual-homed to PE-1 and PE-2. The following resiliency mechanisms are used on the dual-homed MTUs:
  - MTU-4 is dual-connected to PE-1 and PE-2 by an active/standby pseudowire (A/S pseudowire hereafter).
  - MTU-5 is dual-connected to PE-1 and PE-2 by two active pseudowires, one of them being blocked by BGP MH running between PE-1 and PE-2. The PE-1 and PE-2 pseudowires, set up from MTU-5, will be part of the BGP MH site MH-site-2.
  - MTU-4 and MTU-5 are running BGP MH, being SHG site-1 and SAP 1/1/1:8 on MTU-5 part of the same BGP MH site, MH-site-1.
- The CEs are connected to the network in the following way:
  - CE-7, CE-9, and CE-10 are single-connected to the network
  - CE-8 is dual connected to MTU-4 and MTU-5.
  - CE-7 and CE-8 are part of the split-horizon group (SHG) site-1(SAPs 1/1/4:500 and 1/1/3:500 on MTU-4). Assume that CE-7 and CE-8 have a backdoor link between them so that when MTU-5 is elected as DF, CE-7 does not get isolated. This configuration highlights the use of a SHG within a site configuration.

For each BGP MH site, MH-site-1 and MH-site-2, the BGP MH process will elect a DF, blocking the site objects for the non-DF nodes. In other words, based on the specific configuration described throughout the chapter:

- For MH-site-1, MTU-4 will be elected as the DF. The non-DF-MTU-5 will block the SAP 1/1/1:8.
- For MH-site-2, PE-1 will be elected as the DF. The non-DF PE-1 will block the spoke-SDP to MTU-5.

Configuration

This section describes all the relevant configuration tasks for the setup shown in Figure 46: Example topology. The appropriate associated IP/MPLS configuration is out of the scope of this chapter. In this example, the following protocols will be configured beforehand:

- ISIS-TE as IGP with all the interfaces being level-2 (OSPF-TE could have been used instead).
- RSVP-TE as the MPLS protocol to signal the transport tunnels (LDP could have been used instead).
- LSPs between core PEs will be FRR protected (facility bypass tunnels) whereas LSP tunnels between MTUs and PEs will not be protected.

Note:
The designer can choose whether to protect access link failures by means of MPLS FRR or A/S pseudowire or BGP MH. Whereas FRR provides a faster convergence (around 50ms) and stability (it does not impact on the service layer, therefore, link failures do not trigger MAC flush and flooding), some interim inefficiencies can be introduced compared to A/S pseudowire or BGP MH.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

When the IP/MPLS infrastructure is up and running, the specific service configuration including the support for BGP MH can begin.

Global BGP configuration

BGP is used in this configuration guide for these purposes:

1. Exchange of multi-homing site NLRIs and redundancy handling from MTU-5 to the core.
2. Auto-discovery and signaling of the pseudowires in the core, as per RFC 4761.
3. Exchange of multi-homing site NLRIs and redundancy handling at the access for CE-7/CE-8.

A BGP route reflector (RR), PE-3, is used for the reflection of BGP updates corresponding to the preceding uses **1** and **2.**

A direct peering is established between MTU-4 and MTU-5 for use **3**. The same RR could have been used for the three cases, however, like in this example, the designer may choose to have a direct BGP peering between access devices. The reasons for this are:

- By having a direct BGP peering between MTU-4 and MTU-5, the BGP updates do not have to travel back and forth.
- On MTU-4 and MTU-5, BGP is exclusively used for multi-homing, therefore there will not be more BGP peers for either MTUs and a RR adds nothing in terms of control plane scalability.

On all nodes, the autonomous system number must be configured, as follows:

```
# on all nodes:
configure {
    router "Base" {
        autonomous-system 65000
    }
```

The following CLI output shows the global BGP configuration required on MTU-4. The 192.0.2.5 address will be replaced by the corresponding peer or the RR system address for PE-1 and PE-2.

```
# on MTU-4:
configure {
    router "Base" {
        autonomous-system 65000
        bgp {
            router-id 192.0.2.4
            rapid-withdrawal true
            family {
                ipv4 false
                l2-vpn true
            }
            rapid-update {
                l2-vpn true
            }
            group "Multi-Homing" {
            }
            neighbor "192.0.2.5" {
                group "Multi-Homing"
                peer-as 65000
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

In this example, PE-3 is the BGP RR with clients PE-1 and PE-2, as follows:

```
# on PE-3:
configure {
    router "Base" {
        autonomous-system 65000
        bgp {
            router-id 192.0.2.3
            rapid-withdrawal true
            family {
                ipv4 false
                l2-vpn true
            }
            rapid-update {
                l2-vpn true
            }
            group "internal" {
                cluster {
                    cluster-id 1.1.1.1
                }
            }
            neighbor "192.0.2.1" {
                group "internal"
                peer-as 65000
            }
            neighbor "192.0.2.2" {
                group "internal"
                peer-as 65000
            }
```

Some considerations about the relevant BGP commands for BGP-MH:

- It is required to specify **family l2-vpn** in the BGP configuration. That statement will allow the BGP peers to agree on the support for the family AFI=25 (Layer 2 VPN), SAFI=65 (VPLS). This family is used for BGP VPLS as well as for BGP MH and BGP AD.
- The **rapid-update l2-vpn** statement allows BGP MH to send BGP updates immediately after detecting link failures, without having to wait for the Minimum Route Advertisement Interval (MRAI) to send the updates in batches. This statement is required to guarantee a fast convergence for BGP MH.
- Optionally, **rapid-withdrawal** can also be added. In the context of BGP MH, this command is only useful if a particular multi-homing site is cleared. In that case, a BGP withdrawal is sent immediately without having to wait for the MRAI. A multi-homing site is cleared when the BGP MH site is removed or even the entire VPLS service.

Service level configuration

After the IP/MPLS infrastructure is configured, including BGP, this section shows the configuration required at service level (VPLS-500). The focus is on the nodes involved on BGP MH, that is, MTU-4, MTU-5, PE-1, and PE-2. These nodes are highlighted in Figure 47: Nodes involved in BGP MH.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 47: Nodes involved in BGP MH

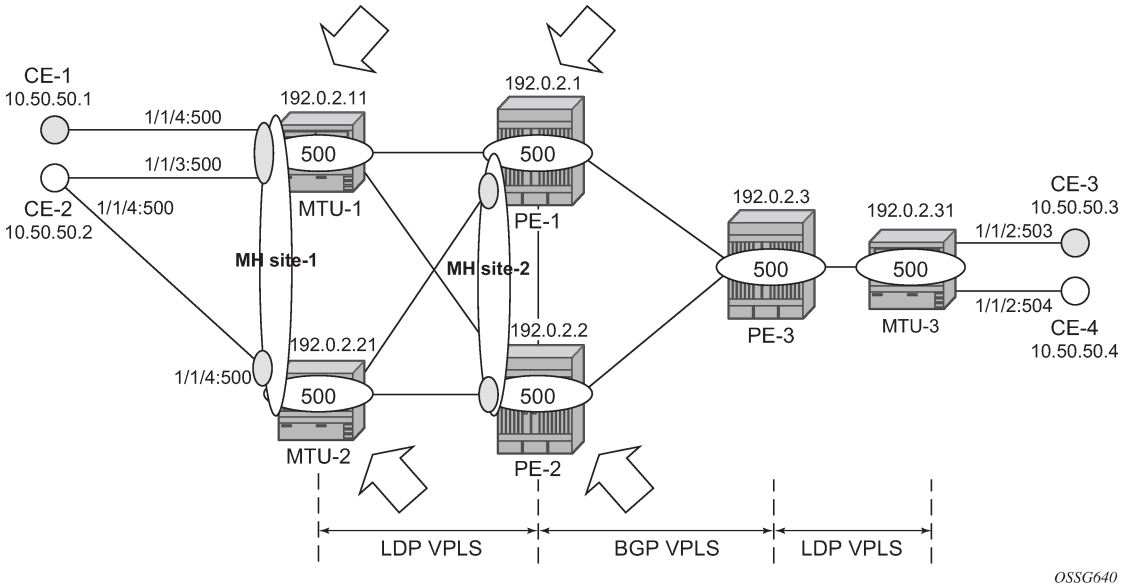

Core PE service configuration

The following CLI excerpt shows the service level configuration on PE-1. The import/export policies configured on the PE nodes are identical:

```
# on PE-1:
configure {
    policy-options {
        community "comm_core" {
            member "target:65000:500" { }
        }
        policy-statement "vsi500_export" {
            entry 10 {
                action {
                    action-type accept
                    community {
                        add ["comm_core"]
                    }
                }
            }
        }
        policy-statement "vsi500_import" {
            entry 10 {
                from {
                    family [l2-vpn]
                    community {
                        name "comm_core"
                    }
                }
                action {
                    action-type accept
                }
            }
            default-action {
                action-type reject
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
        }
```

The configuration of the SDPs, PW template, and VPLS on PE-1 is as follows:

```
# on PE-1:
configure {
    service {
        pw-template "PW500" {
            pw-template-id 500
            provisioned-sdp use
        }
        sdp 12 {
            admin-state enable
            description "SDP to transport BGP-signaled PWs"
            delivery-type mpls
            path-mtu 8000
            signaling bgp
            far-end {
                ip-address 192.0.2.2
            }
            lsp "LSP-PE-1-PE-2" { }
        }
        sdp 13 {
            admin-state enable
            description "SDP to transport BGP-signaled PWs"
            delivery-type mpls
            path-mtu 8000
            signaling bgp
            far-end {
                ip-address 192.0.2.3
            }
            lsp "LSP-PE-1-PE-3" { }
        }
        sdp 14 {
            admin-state enable
            delivery-type mpls
            path-mtu 8000
            far-end {
                ip-address 192.0.2.4
            }
            lsp "LSP-PE-1-MTU-4" { }
        }
        sdp 15 {
            admin-state enable
            delivery-type mpls
            path-mtu 8000
            far-end {
                ip-address 192.0.2.5
            }
            lsp "LSP-PE-1-MTU-5" { }
        }
        vpls "VPLS-500" {
            admin-state enable
            service-id 500
            customer "1"
            bgp 1 {
                route-distinguisher "65000:501"
                vsi-import ["vsi500_import"]
                vsi-export ["vsi500_export"]
                pw-template-binding "PW500" {
                    split-horizon-group "CORE"
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
            bgp-vpls {
                admin-state enable
                maximum-ve-id 65535
                ve {
                    name "501"
                    id 501
                }
            }
            spoke-sdp 14:500 {
            }
            spoke-sdp 15:500 {
            }
            bgp-mh-site "MH-site-2" {
                admin-state enable
                id 2
                spoke-sdp 15:500
            }
        }
```

The following are general comments about the configuration of VPLS-500:

- As seen in the preceding CLI output for PE-1, there are four provisioned SDPs that the service VPLS-500 will use in this example. SDP 14 and SDP 15 are tunnels over which the T-LDP FEC128 pseudowires for VPLS-500 will be carried (according to RFC 4762), whereas SDP 12 and SDP 13 are the tunnels for the core BGP pseudowires (based on RFC 4761).
- The BGP context provides the general service BGP configuration that will be used by BGP VPLS and BGP MH:
  - Route distinguisher (notation chosen is based on <AS_number:500 + node_id>)
  - VSI export policies are used to add the export route-targets included in all the BGP updates sent to the BGP peers.
  - VSI import policies are used to control the NLRIs accepted in the RIB, normally based on the route targets.
  - Both VSI-export and VSI-import policies can be used to modify attributes such as the Local Preference (LP) that will be used to influence the BGP MH Designated Forwarder (DF) election (LP is the second rule in the BGP MH election process, as previously discussed). The use of these policies will be described later in the chapter.
  - The **pw-template-binding** command maps the previously defined pw-template PW500 to the SHG "CORE". In this way, all the BGP-signaled pseudowires will be part of this SHG. Although not shown in this example, the **pw-template-binding** command can also be used to instantiate pseudowires within different SHGs, based on different import route targets:

Note:
Detailed BGP-VPLS configuration is out of the scope of this chapter. For more information, see chapter BGP VPLS.

```
[ex:configure service vpls "VPLS-500" bgp 1]
A:admin@PE-1# pw-template-binding "PW500" ?

 pw-template-binding

 apply-groups          - Apply a configuration group at this level
 apply-groups-exclude  - Exclude a configuration group at this level
 bfd-liveness          - Enable BFD
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
  bfd-template          - BFD template name for PW-Template binding
  import-rt             - Import route-target communities
  split-horizon-group   - Split horizon group

  Choice: oper-group-association
  monitor-oper-group   :- Operational group to monitor
  oper-group           :- Operational group
```

- The BGP-signaled pseudowires (from PE-1 to PE-2 and PE-3) are set up according to the configuration in the BGP context. Beside those pseudowires, the VPLS-500 also has two more pseudowires signaled by TLDP: spoke-SDP 14:500 (to MTU-4) and spoke-SDP 15:500 (to MTU-5).

The MH site name is defined by a string of up to 32 characters:

```
[ex:configure service vpls "VPLS-500"]
A:admin@PE-1# bgp-mh-site ?

 [site-name] <string>
 <string> - <1..32 characters>

    Name for the specific site
```

The general BGP MH configuration parameters for a particular multi-homing site are as follows:

```
[ex:configure service vpls "VPLS-500"]
A:admin@PE-1# bgp-mh-site "MH-site-2" ?

 bgp-mh-site

 activation-timer      - Time that the local sites are in standby status, waiting for
                         BGP updates
 admin-state           - Administrative state of the VPLS BGP multi-homing site
 apply-groups          - Apply a configuration group at this level
 apply-groups-exclude  - Exclude a configuration group at this level
 boot-timer            - Time that system waits after node reboot and before it runs
                         DF election algorithm
 failed-threshold      - Threshold for the site to be declared down
 id                    - ID for the site
 min-down-timer        - Minimum downtime for BGP multi-homing site after transition
                         from up to down
 monitor-oper-group    - Operational group to monitor

 Choice: site-object
 mesh-sdp-binds       :- Specify if a mesh-sdp-binding is associated with this site
 sap                  :- SAP to be associated with this site
 shg-name             :- Split horizon group to be associated with this site
 spoke-sdp            :- SDP to be associated with this site
```

Where:

- The site **name** is defined by a string of up to 64 characters.
- The **id** is an integer that identifies the multi-homing site and is encoded in the BGP MH NLRI. This ID must be the same one used on the peer node where the same multi-homing site is connected to. That is, MH-site-2 must use the same site-id in PE-1 and PE-2 (value = 2 in the PE-1 site configuration).
- Out of the four potential objects in a site—spoke SDP, SAP, SHG, and mesh SDP binding—only one can be used at the time on a particular site. To add more than just one SAP/spoke-SDP to the same site, an SHG composed of the SAP/spoke-SDP objects must be used in the site configuration. Otherwise, only one object—spoke SDP, SAP, SHG, or mesh SDP binding—is allowed per site. When a new object is configured in a site, it replaces the previous object in that site.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- The **failed-threshold** command defines how many objects should be down for the site to be declared down. This command is obviously only valid for multi-object sites (SHGs and mesh-SDP bindings). By default, all the objects in a site must be down for the site to be declared as operationally down.

```
[ex:configure service vpls "VPLS-500" bgp-mh-site "MH-site-2"]
A:admin@PE-1# failed-threshold ?

 failed-threshold (<number> | <keyword>)
 <number>  - <1..1000>
 <keyword> - all
 Default   - all

    Threshold for the site to be declared down
```

- The **boot-timer** specifies for how long the service manager waits after a node reboot before running the MH procedures. The boot-timer value should be configured to allow for the BGP sessions to come up and for the NLRI information to be refreshed/exchanged. In environments with the default BGP MRAI (30 seconds), it is highly recommended to increase this value (for instance, 120 seconds for a normal configuration). The **boot-timer** is only important when a node comes back up and would become the DF. Default value: 10 seconds.

```
[ex:configure service vpls "VPLS-500" bgp-mh-site "MH-site-2"]
A:admin@PE-1# boot-timer ?

 boot-timer <number>
 <number> - <0..600> - seconds

    Time that system waits after node reboot and before it runs DF election algorithm
```

- The **activation-timer** command defines the amount of time the service manager will keep the local objects in standby (in the absence of BGP updates from remote PEs) before running the DF election algorithm to decide whether the site should be unblocked. The timer is started when one of the following events occurs only if the site is operationally up:
  - Manual site activation by enabling the admin-state at the **id** level or at member objects level (SAPs or pseudowires)
  - Site activation after a failure
  - The BGP MH election procedures will be resumed upon expiration of this timer or the arrival of a BGP MH update for the multi-homing site. Default value: 2 seconds.

```
[ex:configure service vpls "VPLS-500" bgp-mh-site "MH-site-2"]
A:admin@PE-1# activation-timer ?

 activation-timer <number>
 <number> - <0..100> - seconds

    Time that the local sites are in standby status, waiting for BGP updates
```

- When a BGP MH site goes down, it may be preferred that it stays down for a minimum time. This is configurable by the **min-down-timer**. When set to zero, this timer is disabled.

```
[ex:configure service vpls "VPLS-500" bgp-mh-site "MH-site-2"]
A:admin@PE-1# min-down-timer ?

 min-down-timer <number>
 <number> - <0..100> - seconds
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    Minimum downtime for BGP multi-homing site after transition from up to down
```

- The **boot-timer**, **activation-timer**, and **min-down-timer** commands can be provisioned at service level or at global level. The service level settings have precedence and override the global configuration. When no timer values are provisioned at global level, the default values apply; when no timer values are provisioned at service level, the timers inherit the global values.

```
[ex:configure redundancy bgp-mh]
A:admin@PE-1# site ?

 site

 activation-timer     - Time to keep local sites in standby status before running DF
                        election algorithm
 boot-timer           - Time that system waits after node reboot and before it runs
                        DF election algorithm
 min-down-timer       - Minimum downtime for BGP multi-homing site after transition
                        from up to down
```

- Each site has three possible states:
  - Admin state — controlled by the admin-state command.
  - Operational state — controlled by the operational status of the individual site objects.
  - Designated Forwarder (DF) state — controlled by the BGP MH election algorithm.

The following CLI output shows the three states for BGP MH site "MH-site-1" on MTU-5:

```
[]
A:admin@MTU-5# show service id 500 site MH-site-1

===============================================================================
Site Information
===============================================================================
Site Name            : MH-site-1
-------------------------------------------------------------------------------
Site Id              : 1
Dest                 : sap:1/1/1:8          Mesh-SDP Bind    : no
Admin Status         : Enabled              Oper Status      : up
Designated Fwdr      : No
DF UpTime            : 0d 00:00:00          DF Chg Cnt       : 0
Boot Timer           : default              Timer Remaining  : 0d 00:00:00
Site Activation Timer: default              Timer Remaining  : 0d 00:00:00
Min Down Timer       : default              Timer Remaining  : 0d 00:00:00
Failed Threshold     : default(all)
Monitor Oper Grp     : (none)
===============================================================================
```

On PE-1, MH-site " MH-site-2" is configured with site ID 2 and object spoke-SDP 15:500 (pseudowire established from PE-1 to MTU-5).

The following CLI shows the service configuration for PE-2. The site ID is 2, that is, the same value configured in PE-1. The object defined in PE-2's site is spoke-SDP 25:500 (pseudowire established from PE-2 to MTU-5).

```
# on PE-2:
    service {
        pw-template "PW500" {
            pw-template-id 500
            provisioned-sdp use
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        sdp 21 {
            admin-state enable
            description "SDP to transport BGP-signaled PWs"
            delivery-type mpls
            path-mtu 8000
            signaling bgp
            far-end {
                ip-address 192.0.2.1
            }
            lsp "LSP-PE-2-PE-1" { }
        }
        sdp 23 {
            admin-state enable
            description "SDP to transport BGP-signaled PWs"
            delivery-type mpls
            path-mtu 8000
            signaling bgp
            far-end {
                ip-address 192.0.2.3
            }
            lsp "LSP-PE-2-PE-3" { }
        }
        sdp 24 {
            admin-state enable
            delivery-type mpls
            path-mtu 8000
            far-end {
                ip-address 192.0.2.4
            }
            lsp "LSP-PE-2-MTU-4" { }
        }
        sdp 25 {
            admin-state enable
            delivery-type mpls
            path-mtu 8000
            far-end {
                ip-address 192.0.2.5
            }
            lsp "LSP-PE-2-MTU-5" { }
        }
        vpls "VPLS-500" {
            admin-state enable
            service-id 500
            customer "1"
            bgp 1 {
                route-distinguisher "65000:502"
                vsi-import ["vsi500_import"]
                vsi-export ["vsi500_export"]
                pw-template-binding "PW500" {
                    split-horizon-group "CORE"
                }
            }
            bgp-vpls {
                admin-state enable
                maximum-ve-id 65535
                ve {
                    name "502"
                    id 502
                }
            }
            spoke-sdp 24:500 {
            }
            spoke-sdp 25:500 {
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            bgp-mh-site "MH-site-2" {
                admin-state enable
                id 2
                spoke-sdp 25:500
            }
        }
```

MTU service configuration

The following CLI output shows the service level configuration on MTU-4.

```
# on MTU-4:
configure {
    service {
        sdp 41 {
            admin-state enable
            delivery-type mpls
            path-mtu 8000
            far-end {
                ip-address 192.0.2.1
            }
            lsp "LSP-MTU-4-PE-1" { }
        }
        sdp 42 {
            admin-state enable
            delivery-type mpls
            path-mtu 8000
            far-end {
                ip-address 192.0.2.2
            }
            lsp "LSP-MTU-4-PE-2" { }
        }
        vpls "VPLS-500" {
            admin-state enable
            service-id 500
            customer "1"
            bgp 1 {
                route-distinguisher "65000:504"
                route-target {
                    export "target:65000:500"
                    import "target:65000:500"
                }
            }
            endpoint "CORE" {
                suppress-standby-signaling false
            }
            split-horizon-group "site-1" {
            }
            spoke-sdp 41:500 {
                endpoint {
                    name "CORE"
                    precedence primary
                }
                stp {
                    admin-state disable
                }
            }
            spoke-sdp 42:500 {
                endpoint {
                    name "CORE"
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                stp {
                    admin-state disable
                }
            }
            bgp-mh-site "MH-site-1" {
                admin-state enable
                id 1
                shg-name "site-1"
            }
            sap 1/1/1:7 {
                split-horizon-group "site-1"
            }
            sap 1/1/2:8 {
                split-horizon-group "site-1"
                eth-cfm {
                    mep md-admin-name "domain-1" ma-admin-name "assoc-1" mep-id 48 {
                        admin-state enable
                        direction down
                        fault-propagation use-if-status-tlv
                        ccm true
                    }
                }
            }
        }
```

MTU-4 is configured with the following characteristics:

- The BGP context provides the general BGP parameters for service 500 in MTU-4. The **route-target** command is now used instead of the vsi-import and vsi-export commands. The intent in this example is to configure only the export and import route-targets. There is no need to modify any other attribute. If the local preference is to be modified (to influence the DF election), a **vsi-policy** must be configured.
- An A/S pseudowire configuration is used to control the pseudowire redundancy towards the core.
- The multi-homing site, MH-site-1 has a site-id = 1 and an SHG as an object. The SHG site-1 is composed of SAP 1/1/1:7 and SAP 1/1/2:8. As previously discussed, the site will not be declared operationally down until the two SAPs belonging to the site are down. This behavior can be changed by the **failed-threshold** command (for instance, in order to bring the site down when only one object has failed even though the second SAP is still up).
- As an example, a Y.1731 MEP with fault-propagation has been defined in SAP 1/1/2:8. As discussed later in the chapter, this MEP will signal the status of the SAP (as a result of the BGP MH process) to CE-8.

The service configuration in MTU-5 is as follows:

```
# on MTU-5:
configure {
    service {
        sdp 51 {
            admin-state enable
            delivery-type mpls
            path-mtu 8000
            far-end {
                ip-address 192.0.2.1
            }
            lsp "LSP-MTU-5-PE-1" { }
        }
        sdp 52 {
            admin-state enable
            delivery-type mpls
            path-mtu 8000
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            far-end {
                ip-address 192.0.2.2
            }
            lsp "LSP-MTU-5-PE-2" { }
        }
        vpls "VPLS-500" {
            admin-state enable
            service-id 500
            customer "1"
            bgp 1 {
                route-distinguisher "65000:505"
                route-target {
                    export "target:65000:500"
                    import "target:65000:500"
                }
            }
            spoke-sdp 51:500 {
            }
            spoke-sdp 52:500 {
            }
            bgp-mh-site "MH-site-1" {
                admin-state enable
                id 1
                sap 1/1/1:8
            }
            sap 1/1/1:8 {
            }
        }
```

Influencing the DF election

As previously described, assuming that the sites on the two nodes taking part of the same multi-homing site are both up, the two tie-breakers for electing the DF are (in this order):

1. Highest LP
2. Lowest PE ID

The LP by default is 100 in all the routers. Under normal circumstances, if the LP in any router is not changed, MTU-4 will be elected the DF for MH-site-1, whereas PE-1 will be the DF for MH-site-2. Assume in this section that this behavior is changed for MH-site-2 to make PE-2 the DF. Because changing the system address (to make PE-2's ID the lower of the two IDs) is usually not an easy task to accomplish, the vsi-export policy on PE-2 is modified with an LP of 150 with which the MH-site-2 NLRI is announced to PE-1. Because LP 150 is greater than the default 100 in PE-1, PE-2 will be elected as the DF for MH-site-2. The vsi-import policy remains unchanged and the vsi-export policy is modified as follows:

```
# on PE-2:
configure {
    policy-options {
        community "comm_core" {
            member "target:65000:500" { }
        }
        policy-statement "vsi500_export" {
            entry 10 {
                action {
                    action-type accept
                    local-preference 150
                    community {
                        add ["comm_core"]
                    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
            }
        }
```

On PE-1, the import and export policies are not modified. The policies were already applied in the **bgp** context of VPLS-500, as follows:

```
# on PE-2:
configure {
    service {
        vpls "VPLS-500" {
            admin-state enable
            service-id 500
            customer "1"
            bgp 1 {
                route-distinguisher "65000:502"
                vsi-import ["vsi500_import"]
                vsi-export ["vsi500_export"]
                pw-template-binding "PW500" {
                    split-horizon-group "CORE"
                }
            }
---snip---
```

The DF state of PE-2 can be verified as follows:

```
[]
A:admin@PE-2# show service id 500 site MH-site-2

===============================================================================
Site Information
===============================================================================
Site Name            : MH-site-2
-------------------------------------------------------------------------------
Site Id              : 2
Dest                 : sdp:25:500          Mesh-SDP Bind    : no
Admin Status         : Enabled             Oper Status      : up
Designated Fwdr      : Yes
DF UpTime            : 0d 00:00:10         DF Chg Cnt       : 2
Boot Timer           : default             Timer Remaining  : 0d 00:00:00
Site Activation Timer: default             Timer Remaining  : 0d 00:00:00
Min Down Timer       : default             Timer Remaining  : 0d 00:00:00
Failed Threshold     : default(all)
Monitor Oper Grp     : (none)
===============================================================================
```

The import and export policies are applied at service 500 level, which means that the LP changes for all the potential multi-homing sites configured under service 500. Therefore, load balancing can be achieved on a per-service basis, but not within the same service.

These policies are applied on VPLS-500 for all the potential BGP applications: BGP VPLS, BGP MH, and BGP AD. In the example, the LP for the PE-2 BGP updates for BGP MH and BGP VPLS will be set to 150. However, this has no impact on BGP VPLS because a PE cannot receive two BGP VPLS NLRIs with the same VE-ID, which implies that a different VE-ID per PE within the same VPLS is required.

The vsi-export policy is restored to its original settings on PE-2, as follows:

```
# on PE-2:
configure {
    policy-options {
        community "comm_core" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            member "target:65000:500" { }
        }
        policy-statement "vsi500_export" {
            entry 10 {
                action {
                    action-type accept
                    delete local-preference
                    community {
                        add ["comm_core"]
                    }
                }
            }
        }
```

In all the PE nodes, the import and export policies applied in the **bgp** context of VPLS-500 have identical settings again, and PE-1 is the DF.

Black-hole avoidance

SR OS supports the appropriate MAC flush mechanisms for BGP MH, regardless of the protocol being used for the pseudowire signaling:

- LDP VPLS — The PE that contains the old DF site (the site that just experienced a DF to non-DF transition) always sends an LDP MAC flush-all-from-me to all LDP pseudowires in the VPLS, including the LDP pseudowires associated with the new DF site. No specific configuration is required.
- BGP VPLS — The remote BGP VPLS PEs interpret the F bit transitions from 1 to 0 as an implicit MAC flush-all-from-me indication. If a BGP update with the flag F=0 is received from the previous DF PE, the remote PEs perform MAC flush-all-from-me, flushing all the MACs associated with the pseudowire to the old DF PE. No specific configuration is required.

Double flushing will not happen because it is expected that between any pair of PEs there will exist only one type of pseudowires—either BGP or LDP pseudowire—, but not both types.

In the example, assuming MTU-4 and PE-1 are the DF nodes:

- When MH-site-1 is brought operationally down on MTU-4 (so by default, the two SAPs must go down unless the **failed-threshold** parameter is changed so that the site is down when only one SAP is brought down), MTU-4 will issue a flush-all-from-me message.
- When MH-site-2 is brought operationally down on PE-1, a BGP update with F=0 and D=1 is issued by PE-1. PE-2 and PE-3 will receive the update and will flush the MAC addresses learned on the pseudowire to PE-1.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 48: MAC flush for BGP MH

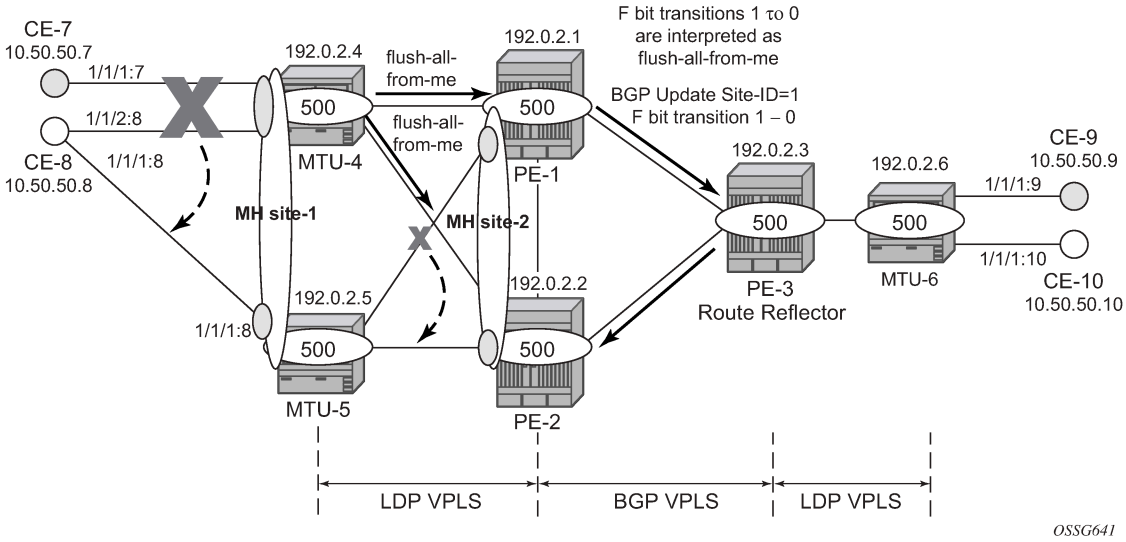

Node failures implicitly trigger a MAC flush on the remote nodes, because the TLDP/BGP session to the failed node goes down.

Access CE/PE signaling

BGP MH works at service level, therefore no physical ports are torn down on the non-DF, but rather the objects are brought down operationally, while the physical port will stay up and used for any other services existing on that port. Because of this reason, there is a need for signaling the standby status of an object to the remote PE or CE.

- Access PEs running BGP MH on spoke SDPs and elected non-DF, will signal pseudowire standby status (0x20) to the other end. If no pseudowire status is supported on the remote MTU, a label withdrawal is performed. If there is more than one spoke SDP on the site (part of the same SHG), the signaling is sent for all the pseudowires of the site.

  Note:
  The **configure service vpls x spoke-sdp y:z pw-status signaling false** parameter allows to send a TLDP label-withdrawal instead of pseudowire status bits, even though the peer supports pseudowire status.

- Multi-homed CEs connected through SAPs to the PEs running BGP MH, are signaled by the PEs using Y.1731 CFM, either by stopping the transmission of CCMs or by sending CCMs with isDown (interface status down encoding in the interface status TLV).

In this example, down MEPs on MTU-4 SAP 1/1/2:8 and CE-8 SAP 1/1/2:8 are configured. In a similar way, other MEPs can be configured on MTU-4 SAP 1/1/1:7, MTU-5 SAP 1/1/1:8, and CE-8 SAP 1/1/1:7 and SAP 1/1/1:8. Figure 49: Access PE/CE signaling shows the MEPs on MTU-4 SAP 1/1/2:8 and CE-8. Upon failure on the MTU-4 site MH-site-1, the MEP 48 will start sending CCMs with interface status down.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 49: Access PE/CE signaling

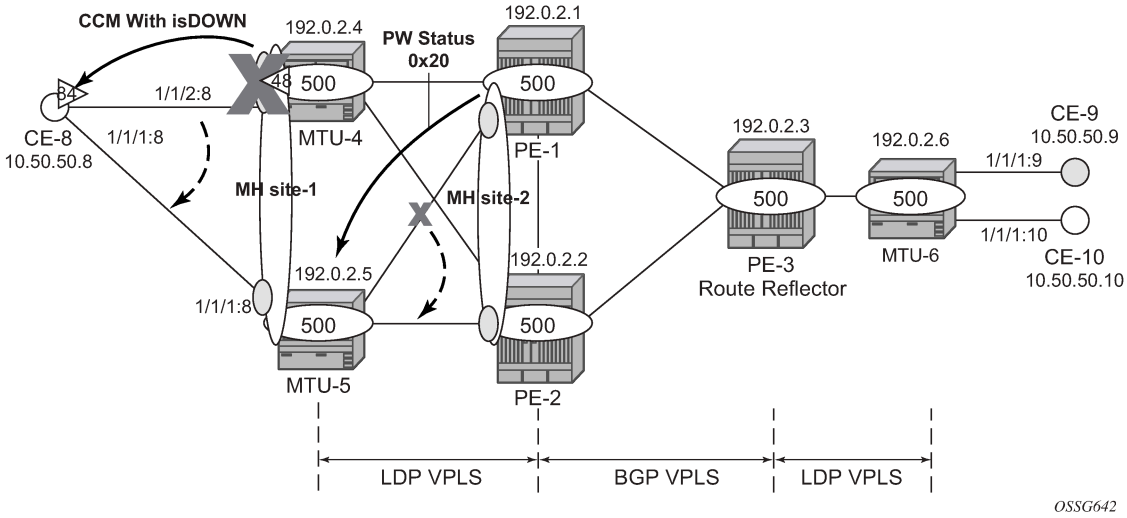

The CFM configuration required at SAP 1/1/2:8 is as follows. Down MEPs will be configured on CE-8 and MTU-5 SAPs in the same way, but in a different association. The option **fault-propagation use-if-status-tlv** must be added. In case the CE does not understand the CCM interface status TLV, the **fault-propagation suspend-ccm** option can be enabled instead. This will stop the transmission of CCMs upon site failures. Detailed configuration guidelines for Y.1731 are beyond the scope of this chapter.

```
# on MTU-4:
configure {
    eth-cfm {
        domain "domain-1" {
            level 3
            name "domain-1"
            md-index 1
            association "assoc-1" {
                icc-based "Association48"
                ma-index 1
                ccm-interval 1s
                bridge-identifier "VPLS-500" {
                }
                remote-mep 84 {
                }
            }
        }
```

```
# on MTU-4:
configure {
    service {
        vpls "VPLS-500" {
            sap 1/1/2:8 {
                split-horizon-group "site-1"
                eth-cfm {
                    mep md-admin-name "domain-1" ma-admin-name "assoc-1" mep-id 48 {
                        admin-state enable
                        direction down
                        fault-propagation use-if-status-tlv
                        ccm true
                    }
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
        }
```

If CE-8 is a service router, upon receiving a CCM with isDown, an alarm will be triggered and the SAP will be brought down:

```
# on CE-8:
71 2021/01/20 16:09:02.701 CET WARNING: OSPF #2047 vprn8 VR:  2 OSPFv2 (0)
"LCL_RTR_ID 10.50.50.8: Interface int-CE-8-MTU-4 state changed to down (event
 IF_DOWN)"

70 2021/01/20 16:09:02.701 CET WARNING: SNMP #2004 vprn8 int-CE-8-MTU-4
"Interface int-CE-8-MTU-4 is not operational"

69 2021/01/20 16:09:02.700 CET MINOR: SVCMGR #2203 vprn8
"Status of SAP 1/1/2:8 in service 8 (customer 1) changed to admin=up oper=down
 flags=OamDownMEPFault "

68 2021/01/20 16:09:02.700 CET MINOR: SVCMGR #2108 vprn8
"Status of interface int-CE-8-MTU-4 in service 8 (customer 1) changed to admin=up
 oper=down"

67 2021/01/20 16:09:02.700 CET MINOR: ETH_CFM #2001 Base
"MEP 1/1/84 highest defect is now defRemoteCCM"
```

On CE-8, the status of the SAP can be verified as follows:

```
[]
A:admin@CE-8# show service id 8 sap 1/1/2:8

===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 8
SAP                : 1/1/2:8                 Encap            : q-tag
Description        : (Not Specified)
Admin State        : Up                      Oper State       : Down
Flags              : OamDownMEPFault
Multi Svc Site     : None
Last Status Change : 01/20/2021 16:09:03
Last Mgmt Change   : 01/20/2021 16:05:49
===============================================================================
```

As also depicted in Figure 49: Access PE/CE signaling, PE-1 will signal pseudowire status standby (code 0x20) when PE-1 goes to non-DF state for MH-site-2. MTU-5 will receive that signaling and, based on the **ignore-standby-signaling** parameter, will decide whether to send the broadcast, unknown unicast, and multicast (BUM) traffic to PE-1. In case MTU-5 uses in its configuration **ignore-standby-signaling**, it will be sending BUM traffic on both pseudowires at the same time (which is not normally desired), ignoring the pseudowire status bits. The following output shows the MTU-5 spoke-SDP receiving the pseudowire status signaling. Although the spoke SDP stays operationally up, the Peer Pw Bits field shows **pwFwdingStandby** and MTU-5 will not send any traffic if the **ignore-standby-signaling** parameter is disabled.

```
[]
A:admin@MTU-5# show service id 500 sdp 51:500 detail

===============================================================================
Service Destination Point (Sdp Id : 51:500) Details
===============================================================================
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
 Sdp Id 51:500  -(192.0.2.1)
------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id             : 51:500                  Type             : Spoke
Spoke Descr     : (Not Specified)
Split Horiz Grp    : (Not Specified)
Etree Root Leaf Tag: Disabled                Etree Leaf AC    : Disabled
VC Type            : Ether                   VC Tag           : n/a
Admin Path MTU     : 8000                    Oper Path MTU    : 8000
Delivery           : MPLS
Far End            : 192.0.2.1               Tunnel Far End   : n/a
Oper Tunnel Far End: 192.0.2.1
LSP Types          : RSVP
---snip---

Admin State        : Up                      Oper State       : Up
---snip---

Endpoint           : N/A                     Precedence       : 4
PW Status Sig      : Enabled
Force Vlan-Vc      : Disabled                Force Qinq-Vc    : none
Class Fwding State : Down
Flags              : None
Time to RetryReset : never                   Retries Left     : 3
Mac Move           : Blockable               Blockable Level  : Tertiary
Local Pw Bits      : None
Peer Pw Bits       : pwFwdingStandby
---snip---
```

Operational groups for BGP-MH

Operational groups (**oper-group**) introduce the capability of grouping objects into a generic group object and associating its status to other service endpoints (pseudowires, SAPs, IP interfaces) located in the same or in different service instances. The operational group status is derived from the status of the individual components using certain rules specific to the application using the concept. A number of other service entities—the monitoring objects—can be configured to monitor the operational group status and to drive their own status based on the **oper-group** status. In other words, if the operational group goes down, the monitoring objects will be brought down. When one of the objects included in the operational group comes up, the entire group will also come up, and therefore so will the monitoring objects.

This concept can be used to enhance the BGP-MH solution for avoiding black-holes on the PE selected as the DF if the rest of the VPLS endpoints fail (pseudowire spoke(s)/pseudowire mesh and/or SAP(s)). Figure 50: Oper-groups and BGP-MH illustrates the use of operational groups together with BGP-MH. On PE-1 (and PE-2) all of the BGP-VPLS pseudowires in the core are configured under the same **oper-group***group-1*. MH-site-2 is configured as a monitoring object. When the two BGP-VPLS pseudowires go down, **oper-group***group-1* will be brought down, therefore MH-site-2 on PE-1 will go down as well (PE-2 will become DF and PE-1 will signal standby to MTU-5).

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 50: Oper-groups and BGP-MH

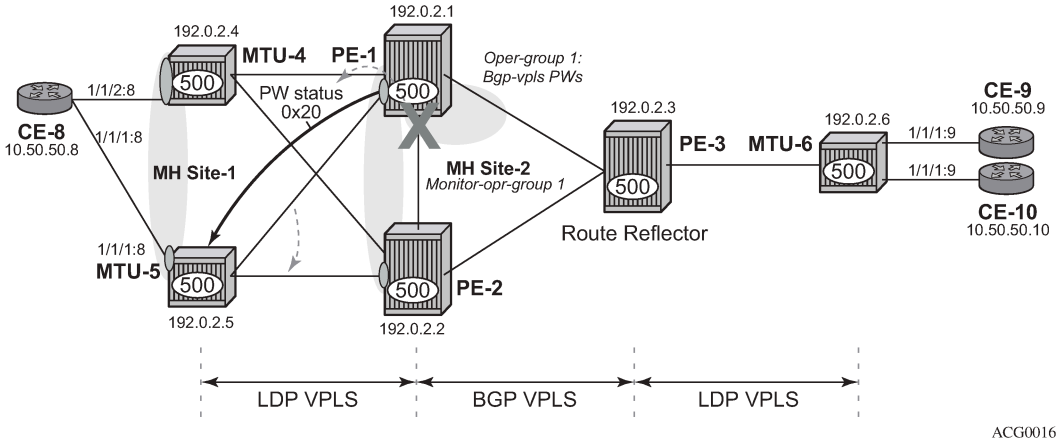

In the preceding example, this feature provides a solution to avoid a black-hole when PE-1 loses its connectivity to the core.

Operational groups are configured in two steps:

1. Identify a set of objects whose forwarding state should be considered as a whole group, then group them under an operational group (in this case **oper-group***group-1*, which is configured in the **bgp pw-template-binding** context).
2. Associate other existing objects (clients) with the oper-group using the **monitor-group** command (configured, in this case, in the **site** *MH-site-2*).

The following CLI excerpt shows the commands required (**oper-group**, **monitor-oper-group**).

```
# on PE-1:
configure {
    service {
    oper-group "group-1" {
        }
        vpls "VPLS-500"
            bgp 1 {
                pw-template-binding "PW500" {
                    split-horizon-group "CORE"
                    oper-group "group-1"
                }
            }
            bgp-mh-site "MH-site-2"
                monitor-oper-group "group-1"
            }
        }
```

When all the BGP-VPLS pseudowires go down, **oper-group***group-1* will go down and therefore the monitoring object, **site** *MH-site-2*, will also go down and PE-2 will then be elected as DF. The log 99 gives information about this sequence of events:

```
# on PE-1:
configure {
    service {
        sdp 12 {
            admin-state disable
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        sdp 13 {
            admin-state disable
        }
```

```
175 2021/01/20 16:15:32.377 CET WARNING: SVCMGR #2531 Base BGP-MH
"Service-id 500 site MH-site-2 is not the designated-forwarder"

174 2021/01/20 16:15:32.377 CET MAJOR: SVCMGR #2316 Base
"Processing of a SDP state change event is finished and the status of all affected SDP
 Bindings on SDP 12 has been updated."

173 2021/01/20 16:15:32.377 CET MAJOR: SVCMGR #2316 Base
"Processing of a SDP state change event is finished and the status of all affected SDP
 Bindings on SDP 13 has been updated."

172 2021/01/20 16:15:32.377 CET MINOR: SVCMGR #2306 Base
"Status of SDP Bind 15:500 in service 500 (customer 1) changed to admin=up oper=down
 flags="

171 2021/01/20 16:15:32.376 CET MINOR: SVCMGR #2326 Base
"Status of SDP Bind 15:500 in service 500 (customer 1) local PW status bits changed
 to pwFwdingStandby "

170 2021/01/20 16:15:32.376 CET MINOR: SVCMGR #2542 Base
"Oper-group group-1 changed status to down"

169 2021/01/20 16:15:32.376 CET MINOR: SVCMGR #2303 Base
"Status of SDP 13 changed to admin=down oper=down"

168 2021/01/20 16:15:32.376 CET MINOR: SVCMGR #2303 Base
"Status of SDP 12 changed to admin=down oper=down"
```

PE-1 is no longer the DF, as follows:

```
[]
A:admin@PE-1# show service id 500 site

===============================================================================
VPLS Sites
===============================================================================
Site                Site-Id   Dest                Mesh-SDP  Admin   Oper  Fwdr
-------------------------------------------------------------------------------
MH-site-2           2         sdp:15:500          no        Enabled down  No
-------------------------------------------------------------------------------
Number of Sites : 1
-------------------------------------------------------------------------------
===============================================================================
```

PE-2 becomes the DF.

```
[]
A:admin@PE-2# show service id 500 site

===============================================================================
VPLS Sites
===============================================================================
Site                Site-Id   Dest                Mesh-SDP  Admin   Oper  Fwdr
-------------------------------------------------------------------------------
MH-site-2           2         sdp:25:500          no        Enabled up    Yes
-------------------------------------------------------------------------------
Number of Sites : 1
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
```

The process reverts when at least one BGP-VPLS pseudowire comes back up.

Show commands and debugging options

The main command to find out the status of a site is the **show service id x site** command.

```
[]
A:admin@MTU-5# show service id 500 site

===============================================================================
VPLS Sites
===============================================================================
Site                Site-Id   Dest               Mesh-SDP  Admin   Oper  Fwdr
-------------------------------------------------------------------------------
MH-site-1           1         sap:1/1/1:8        no        Enabled up    No
-------------------------------------------------------------------------------
Number of Sites : 1
-------------------------------------------------------------------------------
===============================================================================
```

A **detail** modifier is available:

```
[]
A:admin@MTU-5# show service id 500 site detail

===============================================================================
Site Information
===============================================================================
Site Name           : MH-site-1
-------------------------------------------------------------------------------
Site Id             : 1
Dest                : sap:1/1/1:8          Mesh-SDP Bind    : no
Admin Status        : Enabled              Oper Status      : up
Designated Fwdr     : No
DF UpTime           : 0d 00:00:00          DF Chg Cnt       : 0
Boot Timer          : default              Timer Remaining  : 0d 00:00:00
Site Activation Timer: default             Timer Remaining  : 0d 00:00:00
Min Down Timer      : default              Timer Remaining  : 0d 00:00:00
Failed Threshold    : default(all)
Monitor Oper Grp    : (none)
-------------------------------------------------------------------------------
Number of Sites : 1
===============================================================================
```

The **detail**view of the command displays information about the BGP MH timers. The values are only shown if the global values are overridden by specific ones at service level (and will be tagged with **Ovr** if they have been configured at service level). The **Timer Remaining** field reflects the count down from the boot timer and activation timer down to the moment when this router tries to become DF again. Again, this is only shown when the global timers have been overridden by the ones at service level.

The objects on the non-DF site will be brought down operationally and flagged with **StandByForMHProtocol**, for example, for SAP 1/1/1:8 on non-DF MTU-5:

```
[]
A:admin@MTU-5# show service id 500 sap 1/1/1:8

===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Service Access Points(SAP)
===============================================================================
Service Id         : 500
SAP                : 1/1/1:8                 Encap            : q-tag
Description        : (Not Specified)
Admin State        : Up                      Oper State       : Down
Flags              : StandByForMHProtocol
Multi Svc Site     : None
Last Status Change : 01/20/2021 15:11:14
Last Mgmt Change   : 01/20/2021 15:44:01
===============================================================================
```

For spoke SDP 25:500 on non-DF PE-2:

```
[]
A:admin@PE-2# show service id 500 sdp 25:500 detail

===============================================================================
Service Destination Point (Sdp Id : 25:500) Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 25:500  -(192.0.2.5)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id             : 25:500                  Type             : Spoke
---snip---

Admin State        : Up                      Oper State       : Down
---snip---

Flags              : StandbyForMHProtocol
---snip---
```

The BGP MH routes in the RIB, RIB-In and RIB-Out can be shown by using the corresponding **show router bgp routes l2-vpn** and **show router bgp neighbor x.x.x.x filter1 received-routes|advertised-routes family l2-vpn** commands. The BGP MH routes are only shown when the operator uses the **l2-vpn** family modifier. Should the operator want to filter only the BGP MH routes out of the l2-vpn routes, the **l2vpn-type multi-homing** filter has to be added to the **show router bgp routes** commands.

```
[]
A:admin@PE-3# show router bgp routes l2-vpn
===============================================================================
 BGP Router ID:192.0.2.3        AS:65000       Local AS:65000
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP L2VPN Routes
===============================================================================
Flag  RouteType                  Prefix                            MED
      RD                         SiteId                            Label
      Nexthop                    VeId                  BlockSize   LocalPref
      As-Path                    BaseOffset            vplsLabelBa
                                                       se
-------------------------------------------------------------------------------
u*>i  VPLS                       -                     -           0
      65000:501                  -                                 -
      192.0.2.1                  501                   8           100
      No As-Path                 497                   524271
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
u*>i  MultiHome                   -                       -           0
      65000:501                   2                                   -
      192.0.2.1                   -                       -           100
      No As-Path                  -                       -
u*>i  VPLS                        -                       -           0
      65000:502                   -                                   -
      192.0.2.2                   502                     8           100
      No As-Path                  497                     524271
u*>i  MultiHome                   -                       -           0
      65000:502                   2                                   -
      192.0.2.2                   -                       -           100
      No As-Path                  -                       -
-------------------------------------------------------------------------------
Routes : 4
===============================================================================
```

The following output shows the L2-VPN BGP-MH routes from site 2 (PE-1 and PE-2) in detail:

```
[]
A:admin@PE-3# show router bgp routes l2-vpn l2vpn-type multi-homing siteid 2 hunt
===============================================================================
 BGP Router ID:192.0.2.3        AS:65000       Local AS:65000
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP L2VPN-MULTIHOME Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Route Type     : MultiHome
Route Dist.    : 65000:501
Site Id        : 2
Nexthop        : 192.0.2.1
From           : 192.0.2.1
Res. Nexthop   : n/a
Local Pref.    : 100                    Interface Name : NotAvailable
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : 0
AIGP Metric    : None                   IGP Cost       : n/a
Connector      : None
Community      : target:65000:500
                 l2-vpn/vrf-imp:Encap=19: Flags=-DF: MTU=0: PREF=0
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.1
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
Add Paths Send : Default
Last Modified  : 00h07m03s

Route Type     : MultiHome
Route Dist.    : 65000:502
Site Id        : 2
Nexthop        : 192.0.2.2
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
From           : 192.0.2.2
Res. Nexthop   : n/a
Local Pref.    : 100                    Interface Name : NotAvailable
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : 0
AIGP Metric    : None                   IGP Cost       : n/a
Connector      : None
Community      : target:65000:500
                 l2-vpn/vrf-imp:Encap=19: Flags=none: MTU=0: PREF=0
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.2
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
Add Paths Send : Default
Last Modified  : 00h07m03s

 ---snip---
```

The following shows the Layer 2 BGP routes on PE-1:

```
[]
A:admin@PE-1# show service l2-route-table ?

 l2-route-table [detail] [bgp-ad] [multi-homing] [bgp-vpls] [bgp-vpws] [all-routes]

 all-routes            - <keyword>
 bgp-ad                - <keyword>
 bgp-vpls              - <keyword>
 bgp-vpws              - <keyword>
 detail                - keyword - display detailed information
 multi-homing          - <keyword>
```

```
[]
A:admin@PE-1# show service l2-route-table multi-homing

===============================================================================
Services: L2 Multi-Homing Route Information - Summary
===============================================================================
Svc Id     L2-Routes (RD-Prefix)        Next Hop        SiteId     State  DF
-------------------------------------------------------------------------------
500        65000:502                    192.0.2.2       2          up(0)  clear
-------------------------------------------------------------------------------
No. of L2 Multi-Homing Route Entries: 1
===============================================================================
```

In case PE-3 were the RR for MTU-4 and MTU-5 as well as for PE-1 and PE-2, PE-1 would have two more L2-routes for multi-homing in this table, as follows:

```
[]
A:admin@PE-1# show service l2-route-table multi-homing

===============================================================================
Services: L2 Multi-Homing Route Information - Summary
===============================================================================
Svc Id     L2-Routes (RD-Prefix)        Next Hop        SiteId     State  DF
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
500        65000:504                   192.0.2.4       1          up(0)  set
500        65000:505                   192.0.2.5       1          up(0)  clear
500        65000:502                   192.0.2.2       2          up(0)  clear
-------------------------------------------------------------------------------
No. of L2 Multi-Homing Route Entries: 3
===============================================================================
```

When operational groups are configured (as previously shown), the following **show** command helps to find the operational dependencies between monitoring objects and group objects.

```
[]
A:admin@PE-1# show service oper-group "group-1" detail

===============================================================================
Service Oper Group Information
===============================================================================
Oper Group       : group-1
Creation Origin  : manual                          Oper Status: up
Hold DownTime    : 0 secs                          Hold UpTime: 4 secs
Members          : 2                               Monitoring : 1
===============================================================================

=========================================================================
Member SDP-Binds for OperGroup: group-1
=========================================================================
SdpId            SvcId      Type     IP address      Adm     Opr
-------------------------------------------------------------------------
12:4294967295    500        BgpVpls  192.0.2.2       Up      Up
13:4294967294    500        BgpVpls  192.0.2.3       Up      Up
-------------------------------------------------------------------------
SDP Entries found: 2
=========================================================================

===============================================================================
Monitoring Sites for OperGroup: group-1
===============================================================================
SvcId     Site               Site-Id   Dest               Admin   Oper  Fwdr
-------------------------------------------------------------------------------
500       MH-site-2          2         sdp:15:500         Enabled up    Yes
-------------------------------------------------------------------------------
Site Entries found: 1
===============================================================================
```

For debugging, the following CLI sources can be used:

- **log-id 99** — Provides information about the site object changes and DF changes.
- **debug router bgp update** (in classic CLI) — Shows the BGP updates for BGP MH, including the sent and received BGP MH NLRIs and flags.

  ```
  # on MTU-4 (classic CLI):
  debug
      router "Base"
          bgp
              update
  ```

- **debug router ldp** commands (in classic CLI) — Provides information about the pseudowire status bits being signaled as well as the MAC flush messages.

  ```
  # on MTU-4 (classic CLI):
  debug
      router "Base"
  ```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        ldp
            peer 192.0.2.1
                packet
                    init detail
                    label detail
```

As an example, log-id 99 shows the following debug output after disabling MH-site-1 on MTU-4:

```
# on MTU-4:
configure {
    service {
        vpls "VPLS-500"
            sap 1/1/1:7 {
                admin-state disable
            }
            sap 1/1/2:8 {
                admin-state disable
            }
```

```
120 2021/01/20 16:38:54.685 CET WARNING: SVCMGR #2531 Base BGP-MH
"Service-id 500 site MH-site-1 is not the designated-forwarder"

119 2021/01/20 16:38:54.685 CET MINOR: SVCMGR #2203 Base
"Status of SAP 1/1/2:8 in service 500 (customer 1) changed to admin=down oper=down
 flags=SapAdminDown MhStandby"

---snip---
```

On MTU-4, debugging is enabled for BGP updates and the following BGP-MH updates are logged:

```
4 2021/01/20 16:38:54.692 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 86
    Flag: 0x90 Type: 14 Len: 28 Multiprotocol Reachable NLRI:
        Address Family L2VPN
        NextHop len 4 NextHop 192.0.2.5
        [MH] site-id: 1, RD 65000:505
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x80 Type: 4 Len: 4 MED: 0
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0x80 Type: 9 Len: 4 Originator ID: 192.0.2.5
    Flag: 0x80 Type: 10 Len: 4 Cluster ID:
        1.1.1.1
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:65000:500
        l2-vpn/vrf-imp:Encap=19: Flags=-DF: MTU=0: PREF=0
"

---snip---

2 2021/01/20 16:38:54.686 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 72
    Flag: 0x90 Type: 14 Len: 28 Multiprotocol Reachable NLRI:
        Address Family L2VPN
        NextHop len 4 NextHop 192.0.2.4
        [MH] site-id: 1, RD 65000:504
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x80 Type: 4 Len: 4 MED: 0
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:65000:500
        l2-vpn/vrf-imp:Encap=19: Flags=D: MTU=0: PREF=0
"
```

As described earlier, debugging is enabled on MTU-4 for LDP messages between MTU-4 and PE-1. The following MAC flush-all-from-me message is sent by MTU-4 to PE-1.

```
1 2021/01/20 16:38:54.686 CET MINOR: DEBUG #2001 Base LDP
"LDP: LDP
Send Address Withdraw packet (msgId 383) to 192.0.2.1:0
Protocol version = 1
MAC Flush (All MACs learned from me)
Service FEC PWE3: ENET(5)/500 Group ID = 0 cBit = 0
"
```

Assuming all the recommended tools are enabled, a DF to non-DF transition can be shown as well as the corresponding MAC flush messages and related BGP processing.

On PE-1, MH-site-2 is brought down by disabling the spoke-SDP 15:500 object. A BGP-MH update will be sent when the MH site goes down. When all objects on the VPLS are disabled as in the following configuration, a BGP VPLS update will be sent as well.

```
# on PE-1:
configure {
    service {
        vpls "VPLS-500" {
            spoke-sdp 14:500 {
                admin-state disable
            }
            spoke-sdp 15:500 {
                admin-state disable
            }
```

When MH-site-2 is torn down on PE-1, the **debug router bgp update** command allows us to see two BGP updates from PE-1:

- A BGP MH update for site ID 2 with flag D set (because the site is down).
- A BGP VPLS update for veid=501 and flag D set. This is due to the fact that there are no more active objects on the VPLS, besides the BGP pseudowires.

```
4 2021/01/20 16:43:15.326 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 72
    Flag: 0x90 Type: 14 Len: 28 Multiprotocol Reachable NLRI:
        Address Family L2VPN
        NextHop len 4 NextHop 192.0.2.1
        [VPLS/VPWS] preflen 17, veid: 501, vbo: 497, vbs: 8, label-base: 524271,
                    RD 65000:501
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x80 Type: 4 Len: 4 MED: 0
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        target:65000:500
        l2-vpn/vrf-imp:Encap=19: Flags=D: MTU=1514: PREF=0
"

3 2021/01/20 16:43:15.326 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 72
    Flag: 0x90 Type: 14 Len: 28 Multiprotocol Reachable NLRI:
        Address Family L2VPN
        NextHop len 4 NextHop 192.0.2.1
        [MH] site-id: 2, RD 65000:501
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x80 Type: 4 Len: 4 MED: 0
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:65000:500
        l2-vpn/vrf-imp:Encap=19: Flags=D: MTU=0: PREF=0
"
```

The D flag, sent along with the BGP VPLS update for veid 501, would be seen on the remote core PEs as though it was a pseudowire status fault (although there is no TLDP running in the core).

```
[]
A:admin@PE-2# show service id 500 all | match Flag
Flags              : PWPeerFaultStatusBits
Flags              : None
Flags              : None
Flags              : None
```

Conclusion

SR OS supports a wide range of service resiliency options as well as the best-of-breed system level HA and MPLS mechanisms for the access and the core. BGP MH for VPLS completes the service resiliency tool set by adding a mechanism that has some good advantages over the alternative solutions:

- BGP MH provides a common resiliency mechanism for attachment circuits (SAPs), pseudowires (spoke SDPs), split horizon groups and mesh bindings
- BGP MH is a network-based technique which does not need interaction to the CE or MTU to which it is providing redundancy to.

The examples used in this chapter illustrate the configuration of BGP MH for access CEs and MTUs. Show and debug commands have also been suggested so that the operator can verify and troubleshoot the BGP MH procedures.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

BGP Virtual Private Wire Services

This chapter describes BGP Virtual Private Wire Service (VPWS) configurations.

Topics in this chapter include:

Applicability

This chapter is applicable to SR OS and was initially written for SR OS Release 11.0.R4. The MD-CLI in the current edition is based on SR OS Release 21.2.R1. There are no prerequisites for this configuration.

Overview

The following two IETF standards describe the provisioning of Virtual Private Wire Services (VPWS):

- RFC 4447, *Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)*, describes Label Distribution Protocol (LDP) VPWS, where VPWS pseudowires are signaled using LDP between Provider Edge (PE) Routers.
- RFC 6624, *Layer 2 Virtual Private Networks Using BGP for Auto-Discovery and Signaling*, describes the use of Border Gateway Protocol (BGP) for signaling of pseudowires between such PEs.

Figure 51: Example topology shows the example topology with five SR OS routers located in the same Autonomous System (AS). There are three PE routers connected to a single P router and a route reflector (RR) for the AS. The PE routers are all BGP VPWS-aware. The Provider (P) router is BGP VPWS-unaware and does not take part in the BGP process.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 51: Example topology

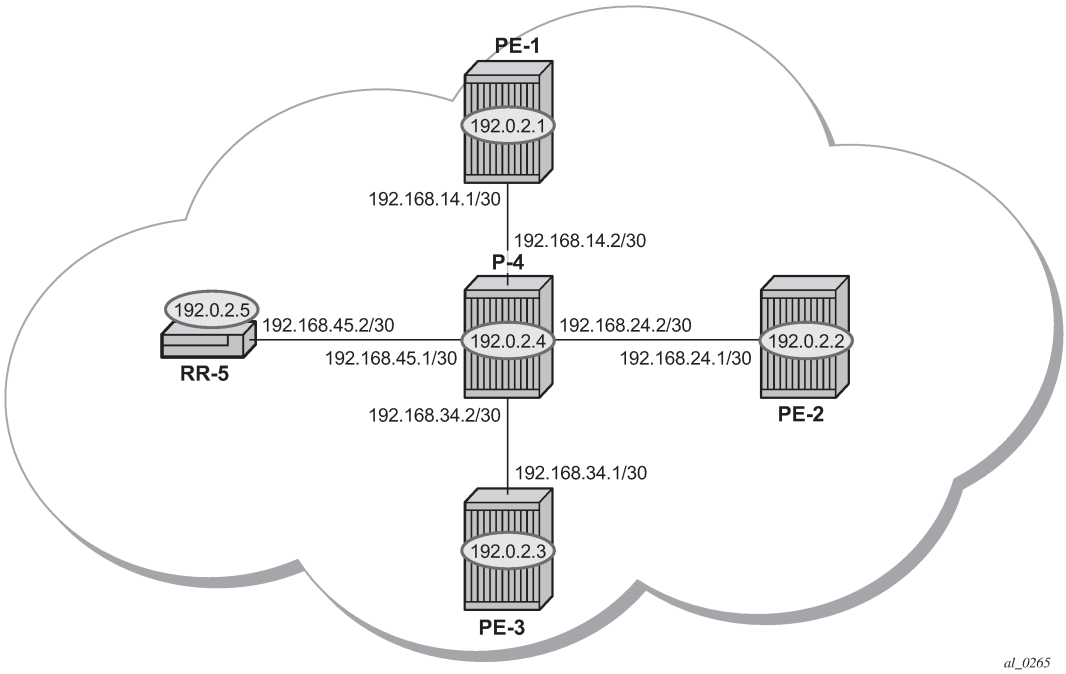

The following configuration tasks are completed as a prerequisite:

- IS-IS or OSPF is configured on each of the network interfaces between the PE/P routers and route reflector.
- MPLS is configured on all interfaces between PE routers and P routers. It is not required between P-4 and RR-5.
- LDP is configured on interfaces between PE and P routers. It is not required between P-4 and the RR-5.
- RSVP is configured on interfaces between PE and P routers. It is not required between P-4 and the RR-5.

BGP VPWS

In this architecture, a VPWS is a collection of two (or three in case of redundancy) BGP VPWS service instances present on different PEs in a provider network.

The PEs communicate with each other at the control plane level by means of BGP updates containing BGP VPWS Network Layer Reachability Information (NLRI). Each update contains enough information for a PE to determine the presence of other BGP VPWS instances on peering PEs and to set up pseudowire connectivity for data flow between peers containing the same BGP VPWS service. Therefore, auto-discovery and pseudowire signaling is achieved using a single BGP update message.

Each PE with a BGP VPWS instance is identified by a VPWS edge identifier (VE-ID) and the presence of other BGP VPWS instances is determined using the exchange of standard BGP extended community route targets (RTs) between PEs.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Each PE will advertise, via the RR, the presence of its BGP VPWS instance to all other PEs, along with a block of multiplexer labels (for BGP VPWS, one label per block) that can be used to communicate between each instance, plus a BGP next-hop that determines a labeled transport tunnel to be used between PEs.

Each BGP VPWS instance is configured with import and export route target extended communities for topology control, along with VE identification.

Configuration

The following examples show the configuration of four BGP VPWS scenarios:

- Single homed BGP VPWS
  - using auto-provisioned SDPs
  - using pre-provisioned SDPs
- Dual homed BGP VPWS
  - with single pseudowire
  - with active/standby pseudowire

Configure MP-iBGP

The first step is to configure an MP-iBGP session between each of the PEs and the RR. The configuration for all PEs is as follows:

```
# on PE-1, PE-2, and PE-3:
configure {
    router "Base"{
        autonomous-system 65536
        bgp {
            group "INTERNAL" {
                peer-as 65536
                family {
                    l2-vpn true
                }
            }
            neighbor "192.0.2.5" {
                group "INTERNAL"
            }
```

The IP addresses can be derived from Figure 51: Example topology.

On RR-5, the BGP configuration is as follows:

```
# on RR-5:
configure {
    router "Base"{
        autonomous-system 65536
         bgp {
            group "INTERNAL" {
                peer-as 65536
                family {
                    l2-vpn true
                }
                cluster {
                    cluster-id 1.1.1.1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
        }
        neighbor "192.0.2.1" {
            group "INTERNAL"
        }
        neighbor "192.0.2.2" {
            group "INTERNAL"
        }
        neighbor "192.0.2.3" {
            group "INTERNAL"
        }
```

The following command on RR-5 shows that BGP sessions with each PE are established and have a negotiated L2 VPN address family capability.

```
[/]
A:admin@RR-5# show router bgp summary all

===============================================================================
BGP Summary
===============================================================================
Legend : D - Dynamic Neighbor
===============================================================================
Neighbor
Description
ServiceId          AS PktRcvd InQ  Up/Down   State|Rcv/Act/Sent (Addr Family)
                      PktSent OutQ
-------------------------------------------------------------------------------
192.0.2.1
Def. Instance  65536        5    0 00h01m24s 0/0/0 (L2VPN)
                            5    0
192.0.2.2
Def. Instance  65536        5    0 00h01m24s 0/0/0 (L2VPN)
                            5    0
192.0.2.3
Def. Instance  65536        5    0 00h01m24s 0/0/0 (L2VPN)
                            5    0

-------------------------------------------------------------------------------
```

Pseudowire templates

BGP VPWS utilizes pseudowire (PW) templates to dynamically instantiate SDP bindings for a service to signal the egress service de-multiplexer labels used by remote PEs to reach the local PE. The template determines the signaling parameters of the pseudowire, such as vc-type, vlan-vc-tag, hash-label, filters, and so on.

- The encapsulation type in the Layer-2 extended community is either 4 (Ethernet VLAN tagged mode) or 5 (Ethernet raw mode), depending on the **vc-type** parameter.
- The **force-vc-forwarding** function will add a tag (equivalent to vc-type vlan) and will allow for customer QoS transparency (dot1p + Drop Eligibility (DE) bits).

The MPLS transport tunnel between PEs can be signaled using LDP or RSVP-TE.

LDP-based SDPs can be automatically instantiated or pre-provisioned. RSVP-TE-based SDPs have to be pre-provisioned. If pre-provisioned pseudowires are used, the PW template must be created with the **provisioned-sdp use** parameter. Alternatively, the **provisioned-sdp prefer** parameter can be used, in

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

which case a pre-provisioned SDP will be used if available; if not, LDP-based SDPs can be automatically instantiated, see chapter LDP VPLS Using BGP Auto-Discovery — Prefer Provisioned SDP.

Pseudowire templates for auto-SDP creation using LDP

In order to use an LDP transport tunnel for data flow between PEs, link layer LDP needs to be configured between all PEs/Ps so that a transport label for each PE system interface is available. For example, on PE-1:

```
# on PE-1:
configure {
    router "Base" {
        ldp {
            interface-parameters {
                interface "int-PE-1-P-4" {
                    ipv4 {
                    }
                }
            }
        }
```

Using this mechanism, SDPs can be auto-instantiated with SDP-IDs starting at the higher end of the SDP numbering range, such as 32767. Any subsequent SDPs created use SDP-IDs decrementing from this value.

A pseudowire template is required. The following example is created using the default values:

```
# on PE-1, PE-2, and PE-3:
configure {
    service {
        pw-template "PW3" {
            pw-template-id 3
        }
```

Pseudowire templates for provisioned SDPs using RSVP-TE

RSVP-TE LSPs need to be created between the PE routers on which provisioned SDPs will be used as prerequisite.

The MPLS interface and LSP configuration for PE-1 are:

```
# on PE-1:
configure {
    router "Base"
        mpls {
            admin-state enable
            interface "int-PE-1-P-4" {
            }
            path "dyn" {
                admin-state enable
            }
            lsp "LSP-PE-1-PE-2" {
                admin-state enable
                type p2p-rsvp
                to 192.0.2.2
                primary "dyn" {
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
            lsp "LSP-PE-1-PE-3" {
                admin-state enable
                type p2p-rsvp
                to 192.0.2.3
                primary "dyn" {
                }
            }
```

The MPLS and LSP configuration for PE-2 are similar to that of PE-1 with the appropriate interfaces and LSP names configured.

To use an RSVP-TE tunnel as transport between PEs, it is necessary to bind the RSVP-TE LSP between PEs to an SDP.

On PE-1, the SDP toward PE-2 is configured as follows. Similar SDPs are required on each PE to the remote PEs in the service where provisioned SDPs are to be used.

```
# on PE-1:
configure {
    service {
        sdp 12 {
            admin-state enable
            description "SDP-PE-1-PE-2_RSVP_BGP"
            delivery-type mpls
            signaling bgp
            far-end {
                ip-address 192.0.2.2
            }
            lsp "LSP-PE-1-PE-2" { }
        }
```

The **signaling bgp** parameter is required. BGP VPWS instances using BGP VPWS signaling can use BGP-signaled SDPs. However, TLDP-signaled (default) SDPs that are bound to RSVP-based LSPs will not be used as SDPs within BGP VPWS.

Single-homed BGP VPWS using auto-provisioned SDPs

Figure 52: Single-homed BGP VPWS using auto-provisioned SDPs shows a schematic of a single homed BGP VPWS between PE-1 and PE-3 where SDPs are auto-provisioned. In this case, the transport tunnels are LDP-signaled.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 52: Single-homed BGP VPWS using auto-provisioned SDPs

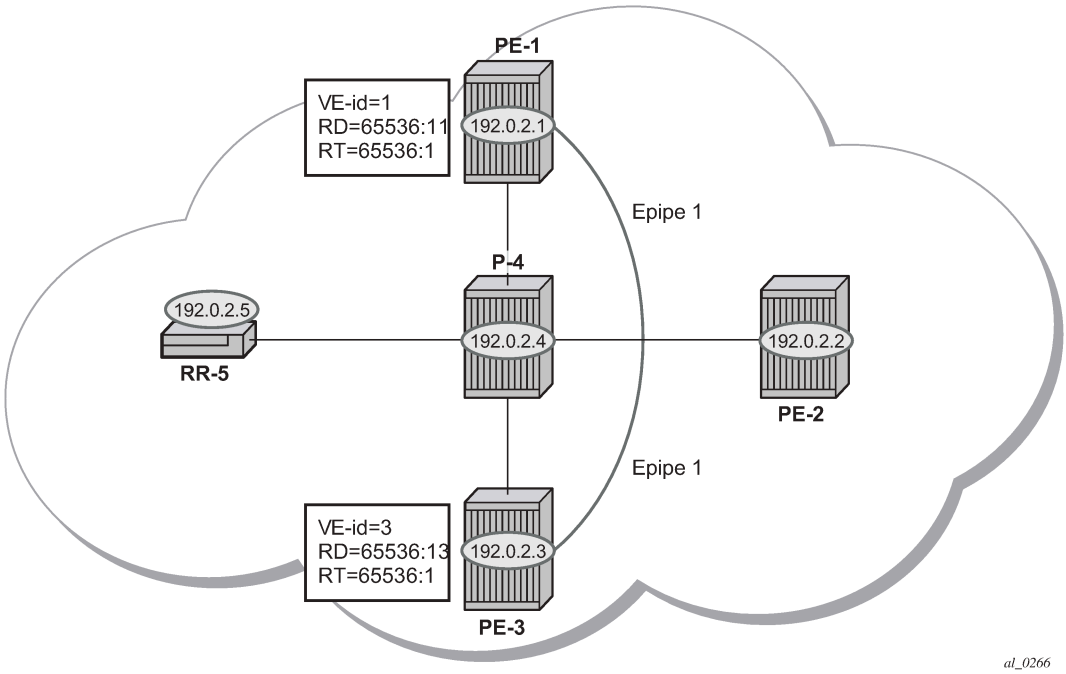

The following shows the configuration required on PE-1 for a BGP VPWS service using a pseudowire template configured for auto-provisioning of SDPs.

```
# on PE-1:
configure {
    service {
        pw-template "PW1" {
            pw-template-id 1
            vc-type vlan
        }
        epipe "Epipe-1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
                route-distinguisher "65536:11"
                route-target {
                    export "target:65536:1"
                    import "target:65536:1"
                }
                pw-template-binding "PW1" {
                }
            }
            bgp-vpws {
                admin-state enable
                local-ve {
                    name "PE-1"
                    id 1
                }
                remote-ve "PE-3" {
                    id 3
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            sap 1/1/4:1 {
                admin-state enable
            }
        }
```

The **bgp** context specifies parameters that are required for BGP VPWS.

Within the **bgp** context, parameters are configured that are used by the neighboring PEs to determine the membership of a BGP VPWS, in other words, the auto-discovery of PEs in the same BGP VPWS. Within the **bgp** context, the RD is configured, along with the route target extended communities. Route target communities are used to determine membership of a BGP VPWS. The import and export route targets at the BGP level are mandatory. The PW template binding is then applied and its parameters are used for both the routes sent by this PE and the received routes matching the route target value.

Within the **bgp-vpws** context, the signaling parameters are configured. These determine the service labels required for the data plane of the VPWS instance.

The VPWS Edge ID (VE-ID) is a numerical value assigned to each PE within a BGP VPWS. This value must be unique for a BGP VPWS, with the exception of multi-homed scenarios, where two dual-homed PEs can have the same VE-ID and are distinguishable by the site preference (or by the tie breaking rules from the *draft-ietf-bess-vpls-multihoming-03*).

Changes to the pseudowire template are not taken into account once the pseudowire has been set up (changes of RT are refreshed though). PW-templates can be re-evaluated with the **tools perform service eval-pw-template** command. The **eval-pw-template** checks if all of the bindings using this PW template policy are still meant to be using this policy. If the template has changed and **allow-service-impact** is true, then the old binding is removed and it is re-added using the new template.

```
[/]
A:admin@PE-1# tools perform service eval-pw-template 1
eval-pw-template succeeded for Svc 1 Tx L2 ExtComm, Policy 1
eval-pw-template succeeded for Svc 1 32767:4294967295 Policy 1
```

VE-ID and BGP label allocations

For a point-to-point VPWS, there are only two members within the BGP VPWS service, so only one label entry is required by each remote service. For dual-homed scenarios, there are two labels for the redundant site, one from each dual-homed PE.

Each PE allocates a label per BGP VPWS instance for the remote PEs, so it signals blocks with one label. It achieves this by advertising three parameters in a BGP update message. For more information about these parameters, see chapter BGP VPLS.

- A Label Base (LB) which is the lowest label in the block.
- A VE Block size (VBS) which is always 1 and cannot be changed.
- A VE Base Offset (VBO) corresponding to the first label in the label block.

PE-3 service creation

On PE-3, Epipe 1 is configured using PW template 1, as follows. PE-3 has been allocated a VE-ID of 3. For completeness, the PW template is also shown.

```
# on PE-3:
configure {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    service {
        pw-template "PW1" {
            pw-template-id 1
            vc-type vlan
        }
        epipe "Epipe-1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
                route-distinguisher "65536:13"
                route-target {
                    export "target:65536:1"
                    import "target:65536:1"
                }
                pw-template-binding "PW1" {
                }
            }
            bgp-vpws {
                admin-state enable
                local-ve {
                    name "PE-3"
                    id 3
                }
                remote-ve "PE-1" {
                    id 1
                }
            }
            sap 1/1/4:1 {
            }
        }
```

PE-1 service operation verification

The following command shows that the BGP VPWS service is enabled on PE-1:

```
[/]
A:admin@PE-1# show service id 1 bgp-vpws

===============================================================================
BGP VPWS Information
===============================================================================
Admin State          : Enabled
VE Name              : PE-1                VE Id           : 1
PW Tmpl used         : 1

Remote-Ve Information
-------------------------------------------------------------------------------
Remote VE Name       : PE-3                Remote VE Id    : 3
===============================================================================
```

The following shows the BGP information used by the BGP VPWS service on PE-1:

```
[/]
A:admin@PE-1# show service id 1 bgp

===============================================================================
BGP Information
===============================================================================
Vsi-Import           : None
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Vsi-Export           : None
Route Dist           : 65536:11
Oper Route Dist      : 65536:11
Oper RD Type         : configured
Rte-Target Import    : 65536:1             Rte-Target Export: 65536:1
Oper RT Imp Origin   : configured          Oper RT Import   : 65536:1
Oper RT Exp Origin   : configured          Oper RT Export   : 65536:1

PW-Template Id       : 1
Endpoint             : <none>
BFD Template         : None
BFD-Enabled          : no                  BFD-Encap        : ipv4
Import Rte-Tgt       : None
-------------------------------------------------------------------------------
===============================================================================
```

Epipe 1 is operationally up on PE-1, as follows:

```
[/]
A:admin@PE-1# show service id 1 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 1                   Vpn Id            : 0
Service Type      : Epipe
MACSec enabled    : no
Name              : Epipe1
Description       : (Not Specified)
Customer Id       : 1                   Creation Origin   : manual
Last Status Change: 03/04/2021 15:25:11
Last Mgmt Change  : 03/04/2021 15:25:11
Test Service      : No
Admin State       : Up                  Oper State        : Up
MTU               : 1514
Vc Switching      : False
SAP Count         : 1                   SDP Bind Count    : 1
Per Svc Hashing   : Disabled
Vxlan Src Tep Ip  : N/A
Force QTag Fwd    : Disabled
Oper Group        : <none>

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                             Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/4:1                            q-tag        1578    1578    Up   Up
sdp:32767:4294967295 SB(192.0.2.3)     BgpVpws      0       1552    Up   Up
===============================================================================
```

The SAP and SDP are all operationally up. The indication "**SB**" next to the SDP-ID signifies "Spoke" and "BGP".

The following output shows the ingress and egress labels for PE-1.

```
[/]
A:admin@PE-1# show service id 1 sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
32767:4294967295 BgpVpws  192.0.2.3       Up      Up        524281    524281
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
```

The following debug output from PE-1 shows the BGP VPWS NLRI update for Epipe 1 sent by PE-1 to RR-5. This update will then be received by the other PEs.

```
# debugging is enabled in classic CLI on PE-1:
debug
    router "Base"
        bgp
            update
```

```
3 2021/03/04 15:25:41.024 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.5
"Peer 1: 192.0.2.5: UPDATE
Peer 1: 192.0.2.5 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 76
    Flag: 0x90 Type: 14 Len: 32 Multiprotocol Reachable NLRI:
        Address Family L2VPN
        NextHop len 4 NextHop 192.0.2.1
        [VPLS/VPWS] preflen 21, veid: 1, vbo: 3, vbs: 1, label-base: 524281,
                    RD 65536:11, csv: 0x00000000, type 1, len 1,
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x80 Type: 4 Len: 4 MED: 0
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:65536:1
        l2-vpn/vrf-imp:Encap=4: Flags=none: MTU=1514: PREF=0
"
```

The control flags within the extended community indicate the status of the BGP VPWS instance.

The control flags are the following:

```
0 1 2 3 4 5 6 7
+-+-+-+-+-+-+-+-+
|D|A|F|Z|Z|Z|C|S| (Z = MUST Be Zero)
+-+-+-+-+-+-+-+-+
```

- D: access circuit down indicator. D is 1 if all access circuits are down, otherwise D is 0.
- A: automatic site ID allocation, which is not supported. This is ignored on receipt and set to 0 on sending.
- F: MAC flush indicator, this relates to VPLS. This is set to 0 and ignored on receipt.
- C: presence of a control word. Control word usage is not supported. This is set to 0 on sending (control word not present) and if a non-zero value is received (indicating a control word is required), the pseudowire will not be created.
- S: sequenced delivery. Sequenced delivery is not supported. This is set to 0 on sending (no sequenced delivery) and if a non-zero value is received (indicating sequenced delivery required), the pseudowire will not be created.

The BGP VPWS NLRI is based on the BGP VPLS NLRI, but is extended with a Circuit Status Vector (CSV). The circuit status vector is used to indicate the status of both the SAP and the spoke-SDP within

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

the local service. Because the VE block size used is 1, the most significant bit in the circuit status vector TLV value will be set to 1 if either the SAP or spoke-SDP is down; otherwise, it will be set to 0.

```
# on PE-1:
configure {
    service {
        epipe "Epipe-1"
            sap 1/1/4:1 {
                admin-state disable
```

```
6 2021/03/04 15:31:59.024 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.5
"Peer 1: 192.0.2.5: UPDATE
Peer 1: 192.0.2.5 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 76
    Flag: 0x90 Type: 14 Len: 32 Multiprotocol Reachable NLRI:
        Address Family L2VPN
        NextHop len 4 NextHop 192.0.2.1
        [VPLS/VPWS] preflen 21, veid: 1, vbo: 3, vbs: 1, label-base: 524281,
                    RD 65536:11, csv: 0x00000080, type 1, len 1,
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x80 Type: 4 Len: 4 MED: 0
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:65536:1
        l2-vpn/vrf-imp:Encap=4: Flags=D: MTU=1514: PREF=0
"
```

After disabling the local SAP, the CSV has the most significant bit set to 1 (0x80). The following command shows the BGP VPWS update received on PE-3:

```
[/]
A:admin@PE-3# show service l2-route-table bgp-vpws detail

===============================================================================
Services: L2 Bgp-Vpws Route Information - Summary
===============================================================================

Svc Id         : 1
VeId           : 1
PW Temp Id     : 1
RD             : *65536:11
Next Hop       : 192.0.2.1
State (D-Bit)  : down(1)
Path MTU       : 1514
Control Word   : 0
Seq Delivery   : 0
Status         : active
Tx Status      : active
CSV            : 80
Preference     : 0
Sdp Bind Id    : 32767:4294967295
===============================================================================
```

On PE-1, SAP 1/1/4:1 is re-enabled as follows:

```
# on PE-1:
configure {
    service {
        epipe "Epipe-1"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            sap 1/1/4:1 {
                admin-state enable
            }
```

PE-3 service operation verification

Similar to PE-1, the service operation should be validated on PE-3.

Single-homed BGP VPWS using pre-provisioned SDP

It is possible to configure BGP VPWS instances that use RSVP-TE transport tunnels. In this case, the SDPs must be created with the MPLS LSPs mapped and with the signaling set to BGP, because the service labels are signaled using BGP. The PW template configured within the BGP VPWS instance must use the keyword provisioned-sdp use (or provisioned-sdp prefer).

Figure 53: Single-homed BGP VPWS using pre-provisioned SDP

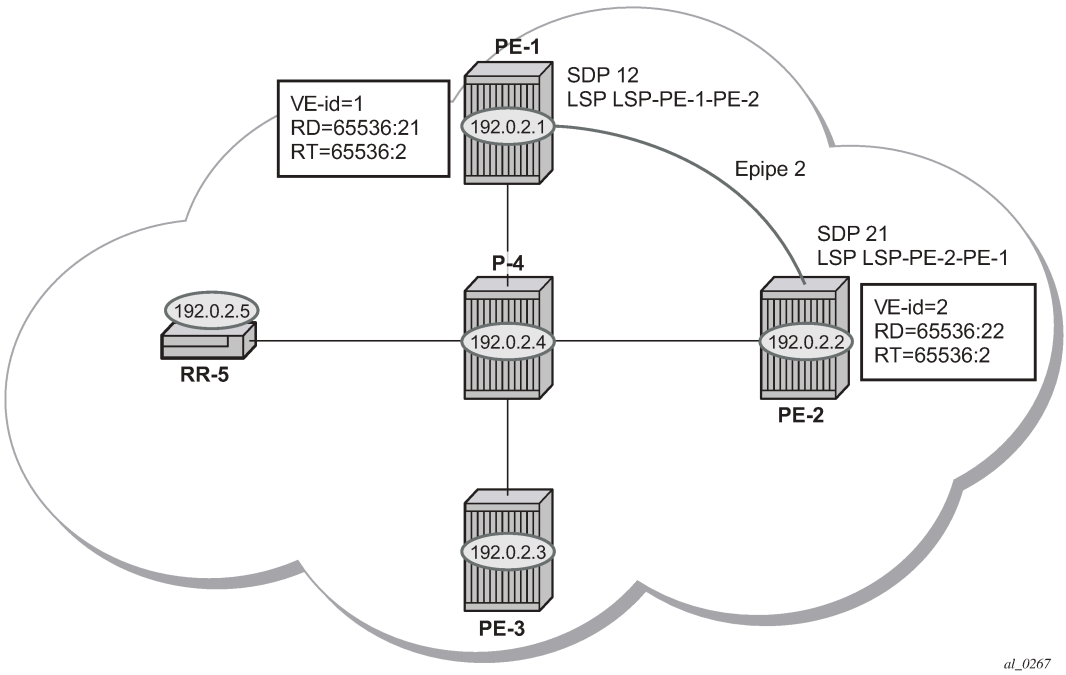

Figure 53: Single-homed BGP VPWS using pre-provisioned SDP shows a schematic of a BGP VPWS where SDPs are pre-provisioned with RSVP-TE signaled transport tunnels.

On PE-1, SDP 12 toward PE-2 is configured as follows:

```
# on PE-1:
configure {
    service {
        sdp 12 {
            admin-state enable
            description "SDP-PE-1-PE-2_RSVP_BGP"
            delivery-type mpls
            signaling bgp
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            far-end {
                ip-address 192.0.2.2
            }
            lsp "LSP-PE-1-PE-2" { }
        }
```

On PE-2, SDP 21 toward PE-1 is configured as follows:

```
# on PE-2:
configure {
    service {
        sdp 21 {
            admin-state enable
            description "SDP-PE-2-PE-1_RSVP_BGP"
            delivery-type mpls
            signaling bgp
            far-end {
                ip-address 192.0.2.1
            }
            lsp "LSP-PE-2-PE-1" { }
        }
```

To create a spoke SDP within a service that uses the RSVP-TE transport tunnel, a pseudowire template is required that has the **provisioned-sdp use** parameter set.

The PW template is provisioned on both PEs as follows:

```
# on PE-1 and PE-2:
configure {
    service {
        pw-template "PW2" {
            pw-template-id 2
            provisioned-sdp use
        }
```

The following output shows the configuration required for a BGP VPWS service using a PW template configured for using pre-provisioned RSVP-TE SDPs.

```
# on PE-1:
configure {
    service {
        epipe "Epipe-2" {
            admin-state enable
            service-id 2
            customer "1"
            bgp 1 {
                route-distinguisher "65536:21"
                route-target {
                    export "target:65536:2"
                    import "target:65536:2"
                }
                pw-template-binding "PW2" {
                }
            }
            bgp-vpws {
                admin-state enable
                local-ve {
                    name "PE-1"
                    id 1
                }
                remote-ve "PE-2" {
                    id 2
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
            }
            sap 1/1/4:2 {
            }
        }
```

The route distinguisher and route target extended community values for Epipe 2 are different from those in Epipe 1. This is to differentiate between the two as their visibility is global within the BGP domain. The VE-ID values can be reused in each Epipe instance, as long as they are unique within the instance.

Similarly, the configuration is as follows on PE-2, where the VE-ID is 2:

```
# on PE-2:
configure {
    service {
        epipe "Epipe-2" {
            admin-state enable
            service-id 2
            customer "1"
            bgp 1 {
                route-distinguisher "65536:22"
                route-target {
                    export "target:65536:2"
                    import "target:65536:2"
                }
                pw-template-binding "PW2" {
                }
            }
            bgp-vpws {
                admin-state enable
                local-ve {
                    name "PE-2"
                    id 2
                }
                remote-ve "PE-1" {
                    id 1
                }
            }
            sap 1/1/4:2 {
            }
        }
```

The service Epipe 2 is operationally up on PE-1, as follows:

```
[/]
A:admin@PE-1# show service id 2 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 2                   Vpn Id            : 0
Service Type      : Epipe
---snip---

Admin State       : Up                  Oper State        : Up
---snip---

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                             Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
sap:1/1/4:2                              q-tag        1578    1578    Up   Up
sdp:12:4294967294 S(192.0.2.2)           BgpVpws      0       1552    Up   Up
===============================================================================
```

The SDP-ID is the pre-provisioned SDP 12.

For completeness, the following command shows that the service is operationally up on PE-2.

```
[/]
A:admin@PE-2# show service id 2 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 2                    Vpn Id            : 0
Service Type      : Epipe
---snip---

Admin State       : Up                   Oper State        : Up
---snip---

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                                Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/4:2                               q-tag        1578    1578    Up   Up
sdp:21:4294967295 S(192.0.2.1)            BgpVpws      0       1552    Up   Up
===============================================================================
```

The SDP-ID used is the pre-provisioned SDP 21.

Dual-homed BGP VPWS with single pseudowire

For access redundancy, an Epipe using a BGP VPWS service can be configured as dual-homed, as described in *draft-ietf-bess-vpls-multihoming-03*. It can be configured with a single pseudowire setup, where the redundant pseudowire is not created until the initially active pseudowire is removed.

The following diagram shows a setup where an Epipe is configured on each PE. Site B is dual-homed to PE-1 and PE-3 with the remote PE-2 connected to site A; each site connection uses a SAP. A single pseudowire using Ethernet Raw Mode encapsulation connects PE-2 to PE-1 or PE-3 (but not both at the same time). The pseudowire is signaled using BGP VPWS over a tunnel LSP between the PEs.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 54: Dual-homed BGP VPWS with single pseudowire

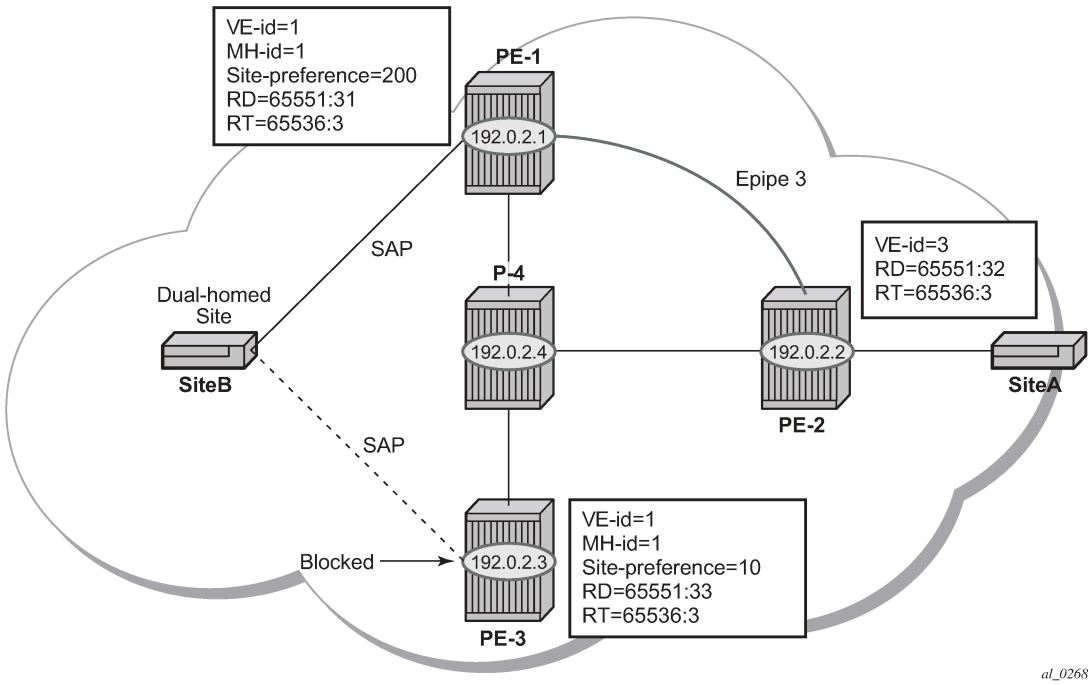

BGP multi-homing is configured for the dual-homed site B using a site-ID=1. The site-preference on PE-1 is set to 200 and to 10 on PE-3, this ensures that PE-1 will be the site's Designated Forwarder (DF) and the pseudowire from PE-2 will be created to PE-1 when PE-1 is fully operational (no pseudowire is created on PE-2 to PE-3). If PE-1 fails, or the multi-homing site fails over to PE-3, then the pseudowire from PE-2 to PE-1 will be removed and a new pseudowire will be created from PE-2 to PE-3.

On PE-1, Epipe 3 is configured as follows:

```
# on PE-1:
configure {
    service {
        pw-template "PW3" {
            pw-template-id 3
        }
        epipe "Epipe-3" {
            admin-state enable
            service-id 3
            customer "1"
            bgp 1 {
                route-distinguisher "65536:31"
                route-target {
                    export "target:65536:3"
                    import "target:65536:3"
                }
                pw-template-binding "PW3" {
                }
            }
            bgp-vpws {
                admin-state enable
                local-ve {
                    name "PE-1"
                    id 1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
                remote-ve "PE-2" {
                    id 2
                }
            }
            bgp-mh-site "SITEB" {
                admin-state enable
                id 1
                sap 1/1/4:3
                preference 200
            }
            sap 1/1/4:3 {
                admin-state enable
            }
        }
```

Epipe 3 is configured on PE-3 with the same local VE-ID as on PE-1, as follows:

```
# on PE-3:
configure {
    service {
        pw-template "PW3" {
            pw-template-id 3
        }
        epipe "Epipe-3" {
            admin-state enable
            service-id 3
            customer "1"
            bgp 1 {
                route-distinguisher "65536:33"
                route-target {
                    export "target:65536:3"
                    import "target:65536:3"
                }
                pw-template-binding "PW3" {
                }
            }
            bgp-vpws {
                admin-state enable
                local-ve {
                    name "PE-3"
                    id 1
                }
                remote-ve "PE-2" {
                    id 2
                }
            }
            bgp-mh-site "SITEB" {
                admin-state enable
                id 1
                sap 1/1/4:3
                preference 10
            }
            sap 1/1/4:3 {
            }
        }
```

In the preceding configurations, the **remote-ve** for PE-2 uses VE-ID 2 on both PE-1 and PE-3.

Epipe 3 is configured on PE-2 as follows:

```
# on PE-2:
configure {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    service {
        pw-template "PW3" {
            pw-template-id 3
        }
        epipe "Epipe-3" {
            admin-state enable
            service-id 3
            customer "1"
            bgp 1 {
                route-distinguisher "65536:32"
                route-target {
                    export "target:65536:3"
                    import "target:65536:3"
                }
                pw-template-binding "PW3" {
                }
            }
            bgp-vpws {
                admin-state enable
                local-ve {
                    name "PE-2"
                    id 2
                }
                remote-ve "PE-1 or PE-3" {
                    id 1
                }
            }
            sap 1/1/4:3 {
            }
        }
```

On PE-2, the **remote-ve** is configured as "PE-1 or PE-3"; this is because both of these PEs are configured with VE-ID 1.

As a result of this configuration, there are multiple route entries for RD 65536:31 on PE-2. In the BGP routing table, there are two entries per partner PE, one for the BGP-MH update (with site-ID=1) and the other for the BGP-VPWS update (with VE-ID=1).

```
[/]
A:admin@PE-2# show router bgp routes l2-vpn rd 65536:31
===============================================================================
 BGP Router ID:192.0.2.2        AS:65536       Local AS:65536
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP L2VPN Routes
===============================================================================
Flag  RouteType                  Prefix                            MED
      RD                         SiteId                            Label
      Nexthop                    VeId                   BlockSize  LocalPref
      As-Path                    BaseOffset             vplsLabelBa
                                                        se
-------------------------------------------------------------------------------
u*>i  MultiHome                  -                      -          0
      65536:31                   1                                 -
      192.0.2.1                  -                      -          200
      No As-Path                 -                      -
u*>i  VPWS                       -                      -          0
      65536:31                   -                                 -
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
      192.0.2.1                   1                     1          200
      No As-Path                  2                     524279
-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

```
[/]
A:admin@PE-2# show router bgp routes l2-vpn rd 65536:33
===============================================================================
 BGP Router ID:192.0.2.2        AS:65536       Local AS:65536
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP L2VPN Routes
===============================================================================
Flag  RouteType                   Prefix                            MED
      RD                          SiteId                            Label
      Nexthop                     VeId                  BlockSize   LocalPref
      As-Path                     BaseOffset            vplsLabelBa
                                                        se
-------------------------------------------------------------------------------
u*>i  MultiHome                   -                     -           0
      65536:33                    1                                 -
      192.0.2.3                   -                     -           10
      No As-Path                  -                     -
u*>i  VPWS                        -                     -           0
      65536:33                    -                                 -
      192.0.2.3                   1                     1           10
      No As-Path                  2                     524280
-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

The route to PE-1 has the higher site preference, so it is selected as the target for the pseudowire.

```
[/]
A:admin@PE-2# show service l2-route-table bgp-vpws detail

===============================================================================
Services: L2 Bgp-Vpws Route Information - Summary
===============================================================================

---snip---

Svc Id         : 3
VeId           : 1
PW Temp Id     : 3
RD             : *65536:31
Next Hop       : 192.0.2.1
State (D-Bit)  : up(0)
Path MTU       : 1514
Control Word   : 0
Seq Delivery   : 0
Status         : active
Tx Status      : active
CSV            : 0
Preference     : 200
Sdp Bind Id    : 32767:4294967292
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
```

After disabling the SAP in the service "Epipe3" on PE-1, BGP update messages are received. The VPLS/VPWS message received on PE-2 from PE-1 shows in the CSV that the access circuit is down (the CSV has the most-significant bit set to 1 (0x80)), so PE-2 selects the update from PE-3 to create the pseudowire. The BGP-MH update received by PE-2 from PE-1 also shows that the local site is down as indicated by the flags=D.

Note in the following debug output:

- BGP MH (multi-homing) entry uses encap-type=19.
- BGP VPWS entry uses encap-type=5 (Ethernet raw mode).

```
# Disable SAP in Epipe 3 on PE-1:
configure {
    service {
        epipe "Epipe-3"
            sap 1/1/4:3 {
                admin-state disable
            }
```

```
34 2021/03/04 15:56:35.904 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.5
"Peer 1: 192.0.2.5: UPDATE
Peer 1: 192.0.2.5 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 90
    Flag: 0x90 Type: 14 Len: 32 Multiprotocol Reachable NLRI:
        Address Family L2VPN
        NextHop len 4 NextHop 192.0.2.1
        [VPLS/VPWS] preflen 21, veid: 1, vbo: 2, vbs: 1, label-base: 524279,
                    RD 65536:31, csv: 0x00000080, type 1, len 1,
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x80 Type: 4 Len: 4 MED: 0
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 0
    Flag: 0x80 Type: 9 Len: 4 Originator ID: 192.0.2.1
    Flag: 0x80 Type: 10 Len: 4 Cluster ID:
        1.1.1.1
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:65536:3
        l2-vpn/vrf-imp:Encap=5: Flags=D: MTU=1514: PREF=200
"
```

```
35 2021/03/04 15:56:35.904 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.5
"Peer 1: 192.0.2.5: UPDATE
Peer 1: 192.0.2.5 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 86
    Flag: 0x90 Type: 14 Len: 28 Multiprotocol Reachable NLRI:
        Address Family L2VPN
        NextHop len 4 NextHop 192.0.2.1
        [MH] site-id: 1, RD 65536:31
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x80 Type: 4 Len: 4 MED: 0
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 0
    Flag: 0x80 Type: 9 Len: 4 Originator ID: 192.0.2.1
    Flag: 0x80 Type: 10 Len: 4 Cluster ID:
        1.1.1.1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:65536:3
        l2-vpn/vrf-imp:Encap=19: Flags=D: MTU=0: PREF=200
"
```

The result can be shown on PE-2 because the spoke SDP to PE-3 is now up (active).

```
[/]
A:admin@PE-2# show service l2-route-table bgp-vpws detail

===============================================================================
Services: L2 Bgp-Vpws Route Information - Summary
===============================================================================

---snip---

Svc Id         : 3
VeId           : 1
PW Temp Id     : 3
RD             : *65536:33
Next Hop       : 192.0.2.3
State (D-Bit)  : up(0)
Path MTU       : 1514
Control Word   : 0
Seq Delivery   : 0
Status         : active
Tx Status      : active
CSV            : 0
Preference     : 10
Sdp Bind Id    : 32767:4294967291
===============================================================================
```

Dual-homed BGP VPWS with active/standby pseudowire

The second method for BGP VPWS pseudowire redundancy is an active/standby configuration. Whereas in the solution with one pseudowire, the redundant nodes use the same VE-ID for the remote PE and different preferences; in the active/standby solution, the redundant nodes use different VE-IDs for the remote PE and different preferences. The node connecting to both pseudowires (PE-2 in this example) has both remote VE-IDs configured. This allows for faster failover because the standby pseudowire is instantiated in addition to the active pseudowire. If more than two applicable BGP updates are received, at most one standby pseudowire is created (based on the BGP VPWS tie breaking rules).

Figure 55: Dual-homed BGP VPWS with active/standby pseudowire shows a setup where an Epipe is configured on each PE. Site B is dual-homed to PE-1 and PE-3 with the remote PE-2 connected to site A; each site connection uses a SAP. The active/standby pseudowires using Ethernet raw mode encapsulation connect PE-2 to PE-1 and PE-3. The pseudowires are signaled using BGP VPWS over tunnel LSPs between the PEs.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 55: Dual-homed BGP VPWS with active/standby pseudowire

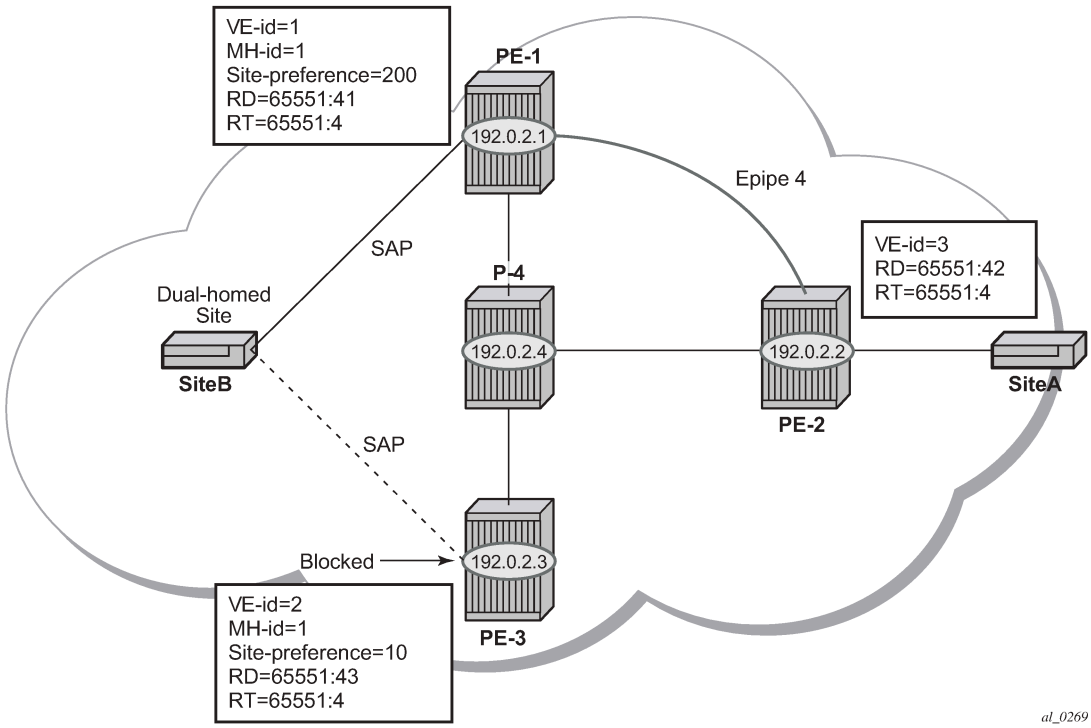

BGP Multi-Homing (MH) is configured for the dual-homed site B using site ID 1. The site preference on PE-1 is set to 200 and to 10 on PE-3; this ensures that PE-1 will be the site's DF for the MH site. The active pseudowire from PE-2 will be created to PE-1 with the standby pseudowire being created to PE-3. If PE-1 fails, or the multi-homing site fails over to PE-3, then the pseudowire from PE-2 to PE-3 will become active (used as the data path between site A and B).

Epipe 4 is configured on PE-1 as follows:

```
# on PE-1:
configure {
    service {
        pw-template "PW3" {
            pw-template-id 3
        }
        epipe "Epipe-4" {
            admin-state enable
            service-id 4
            customer "1"
            bgp 1 {
                route-distinguisher "65536:41"
                route-target {
                    export "target:65536:4"
                    import "target:65536:4"
                }
                pw-template-binding "PW3" {
                }
            }
            bgp-vpws {
                admin-state enable
                local-ve {
                    name "PE-1"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    id 1
                }
                remote-ve "PE-2" {
                    id 2
                }
            }
            bgp-mh-site "SITEB" {
                admin-state enable
                id 1
                sap 1/1/4:4
                preference 200
            }
            sap 1/1/4:4 {
            }
        }
```

Epipe 4 is configured on PE-3 with local VE-ID 3 (different from the previous example), as follows:

```
# on PE-3:
configure {
    service {
        pw-template "PW3" {
            pw-template-id 3
        }
        epipe "Epipe-4" {
            admin-state enable
            service-id 4
            customer "1"
            bgp 1 {
                route-distinguisher "65536:43"
                route-target {
                    export "target:65536:4"
                    import "target:65536:4"
                }
                pw-template-binding "PW3" {
                }
            }
            bgp-vpws {
                admin-state enable
                local-ve {
                    name "PE-3"
                    id 3
                }
                remote-ve "PE-2" {
                    id 2
                }
            }
            bgp-mh-site "SITEB" {
                admin-state enable
                id 1
                sap 1/1/4:4
                preference 10
            }
            sap 1/1/4:4 {
            }
        }
```

Epipe 4 is configured on PE-2 as follows. Two remote VE names are configured, PE-1 and PE-3 (this is the maximum number allowed).

```
# on PE-2:
configure {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    service {
        pw-template "PW3" {
            pw-template-id 3
        }
        epipe "Epipe-4" {
            admin-state enable
            service-id 4
            customer "1"
            bgp 1 {
                route-distinguisher "65536:42"
                route-target {
                    export "target:65536:4"
                    import "target:65536:4"
                }
                pw-template-binding "PW3" {
                }
            }
            bgp-vpws {
                admin-state enable
                local-ve {
                    name "PE-2"
                    id 2
                }
                remote-ve "PE-1" {
                    id 1
                }
                remote-ve "PE-3" {
                    id 3
                }
            }
            sap 1/1/4:4 {
            }
        }
```

Compared with the single pseudowire solution, both pseudowires are signaled and up on all PEs. The pseudowire with the higher preference is forwarding traffic (to PE-1), while the Tx status to the standby PE-3 is set to inactive, as follows:

```
[/]
A:admin@PE-2# show service l2-route-table bgp-vpws detail

===============================================================================
Services: L2 Bgp-Vpws Route Information - Summary
===============================================================================

---snip---

Svc Id         : 4
VeId           : 1
PW Temp Id     : 3
RD             : *65536:41
Next Hop       : 192.0.2.1
State (D-Bit)  : up(0)
Path MTU       : 1514
Control Word   : 0
Seq Delivery   : 0
Status         : active
Tx Status      : active
CSV            : 0
Preference     : 200
Sdp Bind Id    : 32767:4294967289

Svc Id         : 4
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
VeId           : 3
PW Temp Id     : 3
RD             : *65536:43
Next Hop       : 192.0.2.3
State (D-Bit)  : up(0)
Path MTU       : 1514
Control Word   : 0
Seq Delivery   : 0
Status         : active
Tx Status      : inactive
CSV            : 0
Preference     : 10
Sdp Bind Id    : 32766:4294967288
===============================================================================
```

The choice of pseudowire to be used to transmit traffic from PE-2 to PE-1 can also be seen in the endpoint created in the BGP VPWS service. Endpoints are automatically created for the pseudowires within a BGP VPWS service, regardless of whether active/standby pseudowires are used; these endpoints are created with a system generated name that ends with the BGP VPWS service id.

```
[/]
A:admin@PE-2# show service id 4 endpoint

===============================================================================
Service 4 endpoints
===============================================================================
Endpoint name                : _tmnx_BgpVpws-4
Description                  : Automatically created BGP-VPWS endpoint
Creation Origin              : bgpVpws
Revert time                  : 0
Act Hold Delay               : 0
Standby Signaling Master     : false
Standby Signaling Slave      : false
Tx Active (SDP)              : 32767:4294967289
Tx Active Up Time            : 0d 00:02:07
Revert Time Count Down       : never
Tx Active Change Count       : 3
Last Tx Active Change        : 03/04/2021 16:04:40
-------------------------------------------------------------------------------
Members
-------------------------------------------------------------------------------
Spoke-sdp: 32766:4294967288 Prec:4                 Oper Status: Up
Spoke-sdp: 32767:4294967289 Prec:4                 Oper Status: Up
===============================================================================
===============================================================================
```

The following command has no effect on an automatically created VPWS endpoint.

```
tools perform service id <service-id> endpoint <endpoint-name> force-switchover <..>
```

Conclusion

BGP VPWS allows the delivery of Layer 2 virtual private wire services to customers where BGP is commonly used. This chapter shows the configuration of single and dual-homed BGP VPWS services together with the associated show output, which can be used to verify and troubleshoot them.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

BGP VPLS

This chapter describes advanced BGP VPLS configurations.

Topics in this chapter include:

Applicability

This chapter was initially written for SR OS Release 9.0.R3. The MD-CLI in the current edition corresponds to SR OS Release 20.10.R2. There are no prerequisites for this configuration.

Overview

The following two IETF standards describe the provisioning of Virtual Private LAN Services (VPLS).

- RFC 4762, *Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling*, describes Label Distribution Protocol (LDP) VPLS, where VPLS pseudowires are signaled using LDP between VPLS Provider Edge (PE) routers, either configured manually or auto-discovered using BGP.
- RFC 4761, *Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling*, describes the use of Border Gateway Protocol (BGP) for both the auto-discovery of VPLS PEs and signaling of pseudowires between such PEs.

The purpose of this chapter is to describe the configuration and troubleshooting for BGP-VPLS.

Knowledge of BGP-VPLS RFC 4761 architecture and functionality is assumed throughout this chapter, as well as knowledge of Multi-Protocol BGP (MP-BGP).

Figure 56: Example topology shows the example topology with seven SR OS nodes located in the same Autonomous System (AS).

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 56: Example topology

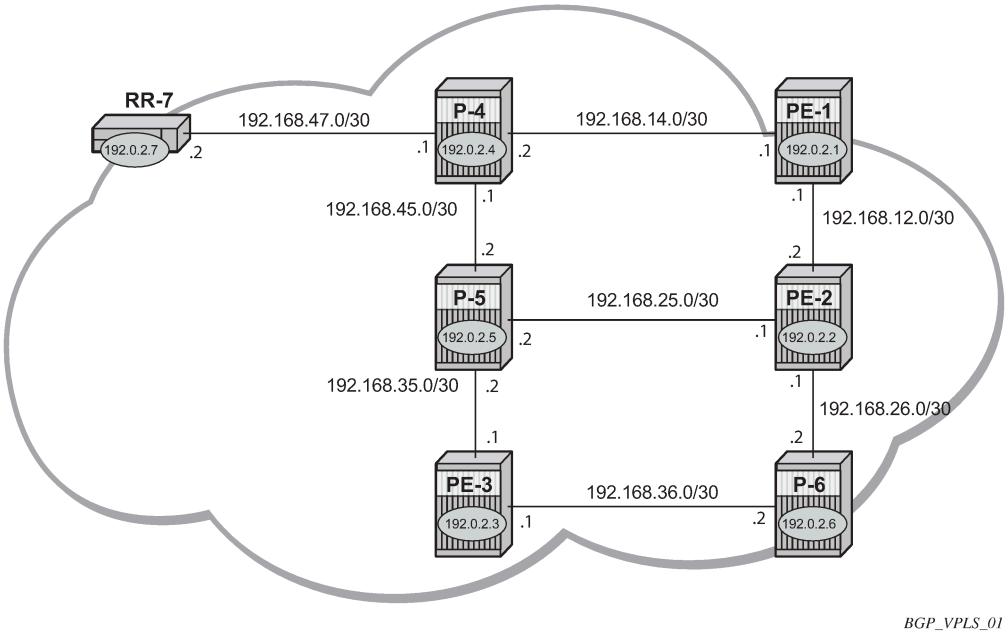

There are three Provider Edge (PE) routers, and RR-7 acts as a Route Reflector (RR) for the AS. The PE routers are all VPLS-aware, the Provider (P) routers are VPLS-unaware and do not take part in the BGP process.

The following configuration tasks are completed as a prerequisite:

- IS-IS or OSPF on each of the network interfaces between the PE/P routers and RR.
- MPLS is configured on all interfaces between PE routers and P routers. MPLS is not required between P-4 and RR-7.
- LDP is configured on interfaces between PE and P routers. It is not required between P-4 and the RR-7.
- The RSVP protocol is enabled.

BGP VPLS

In this architecture, a VPLS instance is a collection of local VPLS instances present on a number of PEs in a provider network. In this context, any VPLS-aware PE is also known as a VPLS Edge (VE) device.

The PEs communicate with each other at the control plane level by means of BGP updates containing BGP-VPLS Network Layer Reachability Information (NLRI). Each update contains enough information for a PE to determine the presence of other local VPLS instances on peering PEs and to set up pseudowire connectivity for data flow between peers containing a local VPLS within the same VPLS instance. Therefore, auto-discovery and pseudowire signaling are achieved using a single BGP update message.

Each PE within a VPLS instance is identified by a VPLS Edge identifier (VE-ID) and the presence of a VPLS instance is determined using the exchange of standard BGP extended community RTs between PEs.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Each PE will advertise, via the route reflectors, the presence of each VPLS instance to all other PEs, along with a block of multiplexer labels that can be used to communicate between such instances plus a BGP next hop that determines a labeled transport tunnel between PEs.

Each VPLS instance is configured with import and export RT extended communities for topology control, along with VE identification.

Configuration

The first step is to configure an MP-iBGP session between each of the PEs and the RR for the L2-VPN address family, as follows:

```
# on PE-1, PE-2, and PE-3:
configure {
    router "Base" {
        autonomous-system 65536
        bgp {
            group "INTERNAL" {
                peer-as 65536
                family {
                    l2-vpn true
                }
            }
            neighbor "192.0.2.7" {
                group "INTERNAL"
            }
```

The IP addresses can be derived from Figure 56: Example topology.

The BGP configuration for RR-7 is as follows:

```
# on RR-7:
configure {
    router "Base" {
        autonomous-system 65536
        bgp {
            cluster {
                cluster-id 1.1.1.1
            }
            group "RR-INTERNAL" {
                peer-as 65536
                family {
                    l2-vpn true
                }
            }
            neighbor "192.0.2.1" {
                group "RR-INTERNAL"
            }
            neighbor "192.0.2.2" {
                group "RR-INTERNAL"
            }
            neighbor "192.0.2.3" {
                group "RR-INTERNAL"
            }
```

On PE-1, the BGP session with RR-7 is established with the L2-VPN address family capability negotiated, as follows:

```
[]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
A:admin@PE-1# show router bgp neighbor 192.0.2.7

===============================================================================
BGP Neighbor
===============================================================================
-------------------------------------------------------------------------------
Peer                 : 192.0.2.7
Description          : (Not Specified)
Group                : INTERNAL
-------------------------------------------------------------------------------
Peer AS              : 65536            Peer Port            : 50198
Peer Address         : 192.0.2.7
Local AS             : 65536            Local Port           : 179
Local Address        : 192.0.2.1
Peer Type            : Internal         Dynamic Peer         : No
State                : Established      Last State           : Established
Last Event           : recvOpen
Last Error           : Cease (Connection Collision Resolution)
Local Family         : L2-VPN
Remote Family        : L2-VPN
Hold Time            : 90               Keep Alive           : 30
Min Hold Time        : 0
Active Hold Time     : 90               Active Keep Alive    : 30
Cluster Id           : None
Preference           : 170              Num of Update Flaps  : 0
---snip---

Local Capability     : RtRefresh MPBGP 4byte ASN
Remote Capability    : RtRefresh MPBGP 4byte ASN
---snip---
```

On RR-7, the BGP sessions with each PE are established, and have negotiated the L2-VPN address family capability, as follows:

```
[]
A:admin@RR-7# show router bgp summary all

===============================================================================
BGP Summary
===============================================================================
Legend : D - Dynamic Neighbor
===============================================================================
Neighbor
Description
ServiceId          AS PktRcvd InQ  Up/Down   State|Rcv/Act/Sent (Addr Family)
                      PktSent OutQ
-------------------------------------------------------------------------------
192.0.2.1
Def. Instance  65536       11    0 00h04m29s 0/0/0 (L2VPN)
                           11    0
192.0.2.2
Def. Instance  65536       11    0 00h04m29s 0/0/0 (L2VPN)
                           11    0
192.0.2.3
Def. Instance  65536       11    0 00h04m29s 0/0/0 (L2VPN)
                           11    0

-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

A full mesh of RSVP-TE LSPs is configured between the PE routers. On PE-1, the MPLS interface and LSP configuration are as follows:

```
# on PE-1:
configure {
    router "Base" {
        mpls {
            admin-state enable
            interface "int-PE-1-P-4" {
            }
            interface "int-PE-1-PE-2" {
            }
            path "loose" {
                admin-state enable
            }
            lsp "LSP-PE-1-PE-2" {
                admin-state enable
                type p2p-rsvp
                to 192.0.2.2
                primary "loose" {
                }
            }
            lsp "LSP-PE-1-PE-3" {
                admin-state enable
                type p2p-rsvp
                to 192.0.2.3
                primary "loose" {
                }
            }
```

The MPLS and LSP configuration for PE-2 and PE-3 are similar to that of PE-1 with the appropriate interfaces and LSP names configured.

BGP VPLS PE configuration

Pseudowire templates

Pseudowire templates are used by BGP to dynamically instantiate SDP bindings for a service to signal the egress service de-multiplexer labels used by remote PEs to reach the local PE.

The template determines the signaling parameters of the pseudowire, control word presence, MAC-pinning, filters, and so on, plus other usage characteristics such as split horizon groups (SHGs).

The MPLS transport tunnel between PEs can be signaled using LDP or RSVP-TE.

LDP based pseudowires can be automatically instantiated. RSVP-TE based SDPs have to be pre-provisioned.

Pseudowire templates for auto-SDP creation using LDP

In order to use an LDP transport tunnel for data flow between PEs, link layer LDP must be configured between all PEs/Ps, so that a transport label for each PE’s system interface is available.

```
# on PE-1:
configure (
    router "Base" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        ldp {
            interface-parameters {
                interface "int-PE-1-P-4" {
                    ipv4 {
                    }
                }
                interface "int-PE-1-PE-2" {
                    ipv4 {
                    }
                }
            }
        }
```

Using this mechanism, SDPs can be auto-instantiated with SDP-IDs starting at the higher end of the SDP numbering range, such as 32767. Any subsequent SDPs created use SDP-IDs decrementing from this value.

A pseudowire template is required containing an SHG. Each SDP created with this template is contained within an SHG so that traffic cannot be forwarded between them.

```
# on PE-1:
configure {
    service {
        pw-template "PW1" {
            pw-template-id 1
            split-horizon-group {
                name "VPLS-SHG"
            }
        }
```

The pseudowire template also has the following options available when used for BGP-VPLS:

```
[ex:configure service]
A:admin@PE-1# pw-template "PW1" ?

 ---snip---

 control-word          - Enable/Disable the use of ControlWord
 ---snip---

 force-vc-forwarding   - VC forwarding action
 ---snip---

 vc-type               - Type of virtual circuit associated with the SDP bind.
 ---snip---
```

- The control word will determine whether the C flag is set in the Layer 2 extended community and, therefore, if a control word is used in the pseudowire.
- The **force-vlan-vc-forwarding** command will add a tag (equivalent to **vc-type vlan**) and will allow for customer QoS transparency (dot1p + Drop Eligibility (DE) bits).

```
[ex:configure service pw-template "PW1"]
A:admin@PE-1# force-vc-forwarding ?

 force-vc-forwarding <keyword>
 <keyword>  - (vlan|qinq-c-tag-c-tag|qinq-s-tag-c-tag)
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- The encap type in the Layer 2 extended community is always 19 (VPLS encap), therefore, the **vc-type** will always be **ether** regardless of the configured value on the vc-type.

```
[ex:configure service pw-template "PW1"]
A:admin@PE-1# vc-type ?

 vc-type <keyword>
 <keyword>  - (ether|vlan)
 Default    - ether

    Type of virtual circuit associated with the SDP bind.
```

Pseudowire templates for provisioned SDPs using RSVP-TE

To use an RSVP-TE tunnel as transport between PEs, it is necessary to bind the RSVP-TE LSP between PEs to an SDP.

The following SDP is created from PE-1 to PE-2:

```
# on PE-1:
configure {
    service {
        sdp 12 {
            admin-state enable
            description "SDP-PE-1-PE-2_RSVP_BGP"
            delivery-type mpls
            signaling bgp
            far-end {
                ip-address 192.0.2.2
            }
            lsp "LSP-PE-1-PE-2" { }
        }
```

The **signaling bgp** parameter is required for BGP-VPLS to be able to use this SDP. Conversely, SDPs that are bound to RSVP-based LSPs with signaling set to the default value of **tldp** will not be used as SDPs within BGP-VPLS.

BGP VPLS using auto-provisioned SDPs

Figure 57: BGP VPLS using auto-provisioned SDPs shows a VPLS instance where SDPs are auto-provisioned. In this case, the transport tunnels are LDP-signaled.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 57: BGP VPLS using auto-provisioned SDPs

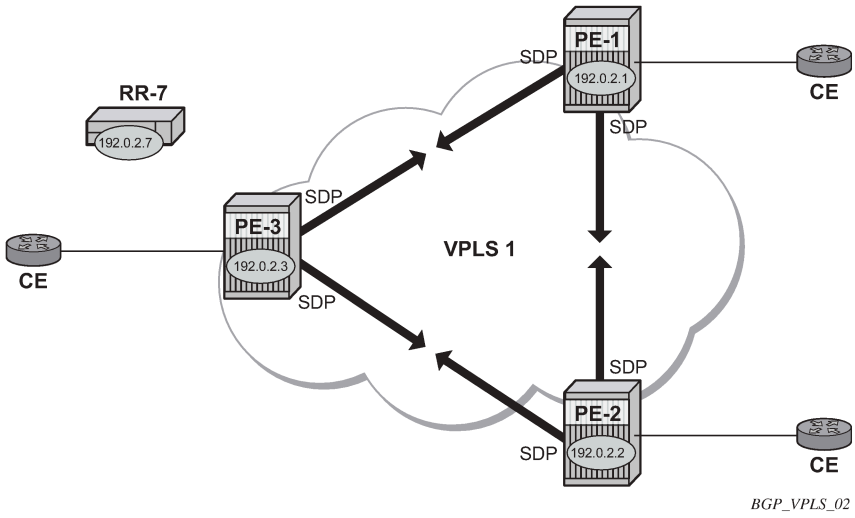

The following shows the configuration required on PE-1 for a BGP-VPLS service using a pseudowire template configured for auto-provisioning of SDPs.

```
# on PE-1:
configure {
    service {
        pw-template "PW1" {
            pw-template-id 1
            split-horizon-group {
                name "VPLS-SHG"
            }
        }
        vpls "VPLS1_PE-1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
                route-distinguisher "65536:1"
                route-target {
                    export "target:65536:1"
                    import "target:65536:1"
                }
                pw-template-binding "PW1" {
                }
            }
            bgp-vpls {
                admin-state enable
                maximum-ve-id 10
                ve {
                    name "PE-1"
                    id 1
                }
            }
            sap 1/1/4:1.0 {
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The **bgp** context specifies parameters which are valid for all of the VPLS BGP applications, such as BGP multi-homing (BGP-MH), BGP auto-discovery (BGP-AD), and BGP-VPLS.

Within the **bgp** context, parameters are configured that are used by neighboring PEs to determine membership of a VPLS instance, such as the auto-discovery of PEs containing the same VPLS instance; the route distinguisher (RD) is configured, along with the route target (RT) extended communities.

RT communities are used to determine membership of a VPLS instance. The import RT at the BGP level is mandatory. The pseudowire template bind is then applied by the service manager on the received routes matching the RT value.

Within the **bgp-vpls** context, the signaling parameters are configured. These determine the service labels required for the data plane of the VPLS instance.

The VPLS edge ID (VE-ID) is a numerical value assigned to each PE within a VPLS instance. This value should be unique for a VPLS instance; no two PEs within the same instance should have the same VE-ID values.

A more specific RT can be applied to a pseudowire template in order to define a specific pseudowire topology, rather than only a full mesh, using the command within the **bgp** context:

```
[ex:configure service vpls "VPLS1_PE-1" bgp 1 pw-template-binding "PW1"]
A:admin@PE-1# import-rt ?

 import-rt <value>
 import-rt [<value>...] - 0..5 system-ordered values separated by spaces enclosed by
                          brackets

 <value> - <string>
 <string> - <10..28 characters>

    Import route-target communities
```

Changes to the import policies are not taken once the pseudowire has been set up (changes on RT are refreshed though). Pseudowire templates can be re-evaluated with the command **tools perform service eval-pw-template**. The **eval-pw-template** command checks whether all the bindings using this pseudowire template policy are still meant to use this policy.

If the policy has changed and **allow-service-impact** is true, then the old binding is removed and it is re-added with the new template.

VE-ID and BGP label allocations

The choice of VE-ID is crucial in ensuring efficient allocation of de-multiplexer labels. The most efficient choice is for VE-IDs to be allocated starting at 1 and incrementing for each PE as the following section explains.

The **maximum-ve-id** *value* determines the range of the VE-ID value that can be configured. If a PE receives a BGP-VPLS update containing a VE-ID with a greater value than the configured **maximum-ve-id**, then the update is dropped and no service labels are installed for this VE-ID.

The **maximum-ve-id** command also checks the locally-configured VE-ID, and prevents a higher value from being used.

Each PE allocates blocks of labels per VPLS instance to remote PEs, in increments of eight labels. It achieves this by advertising three parameters in a BGP update message,

- A label base (LB) which is the lowest label in the block

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- A VE Block Size (VBS) which is always eight labels, and cannot be changed
- A VE Base Offset (VBO).

This defines a block of labels in the range (LB, LB+1, ..., LB+VBS-1).

As an example, if the label base (LB) = 524272, then the range for the block is 524272 to 524279, which is exactly eight labels, as per the block size. (The last label in the block is calculated as 524272+8-1 = 524279)

The label allocated by the PE to each remote PE within the VPLS is chosen from this block and is determined by its VE-ID. In this way, each remote PE has a unique de-multiplexer label for that VPLS.

To reduce label wastage, contiguous VE-IDs in the range (N..N+7) per VPLS should be chosen, where N>0.

Assuming a collection of PEs with contiguous VE-IDs, the following labels will be chosen by PEs from the label block allocated by PE-1 which has a VE-ID =1.

Table 3: VE-IDs and Labels

VE-ID	Label
2	524273
3	524274
4	524275
5	524276
6	524277
7	524278
8	524279

This shows that the label allocated to a PE is (LB+VEID-1). The "1" is the VE block offset (VBO).

This means that the label allocated to a PE router within the VPLS can now be written as (LB + VEID - VBO), which means that (VEID - VBO) calculation must always be at least zero and be less than the block size, which is always 8.

For VE-ID < 8, a label will be allocated from this block.

For the next block of 8 VE-IDs (VE-ID 9 to VE-ID 16) a new block of 8 labels must be allocated, so a new BGP update is sent, with a new label base, and a block offset of 9.

Table 4: VE-IDs and Number of Labels shows how the choice of VE-IDs can affect the number of label blocks allocated, and therefore the number of labels:

Table 4: VE-IDs and Number of Labels

VE-ID	Block Offset	Labels Allocated
1-8	1	8
9-16	9	8

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

VE-ID	Block Offset	Labels Allocated
17-24	17	8
25-32	25	8
33-40	33	8
41-48	41	8
49-56	49	8

This shows that the most efficient use of labels occurs when the VE-IDs for a set of PEs are chosen from the same block offset.

If VE-IDs are chosen that map to different block offsets, then each PE will have to send multiple BGP updates to signal service labels. Each PE sends label blocks in BGP updates to each of its BGP neighbors for all label blocks in which at least one VE-ID has been seen by this PE (it does not advertise label blocks which do not contain an active VE-ID, where active VE-ID means the VE-ID of this PE or any other PE in this VPLS).

The **maximum-ve-id** must be configured first, and determines the maximum value of the VE-ID that can be configured within the PE. The VE-ID value cannot be higher than this within the PE configuration, VE-ID <= maximum VE-ID. Similarly, if the VE-ID within a received NLRI is higher than the maximum VE-ID value, it will not be accepted as valid consequently the maximum VE-ID configured on all PEs must be greater than or equal to any VE-ID used in the VPLS.

Only one VE-ID value can be configured. If the VE-ID value is changed, BGP withdraws the NLRI and sends a route-refresh.

If the same VE-ID is used in different PEs for the same VPLS, a Designated Forwarder (DF) election takes place.

Executing the **admin-state disable** command triggers an MP-UNREACH-NLRI from the PE to all BGP peers.

The **admin-state enable** command triggers an MP-REACH-NLRI to the same peers.

PE-2 service creation

On PE-2, a VPLS service using pseudowire template 1 is created. In order to make the label allocation more efficient, PE-2 has been allocated a VE-ID value of 2. For completeness, the pseudowire template is also shown.

```
# on PE-2:
configure {
    service {
        pw-template "PW1" {
            pw-template-id 1
            split-horizon-group {
                name "VPLS-SHG"
            }
        }
         vpls "VPLS1_PE-2" {
            admin-state enable
            service-id 1
            customer "1"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            bgp 1 {
                route-distinguisher "65536:1"
                route-target {
                    export "target:65536:1"
                    import "target:65536:1"
                }
                pw-template-binding "PW1" {
                }
            }
            bgp-vpls {
                admin-state enable
                maximum-ve-id 10
                ve {
                    name "PE-2"
                    id 2
                }
            }
            sap 1/1/4:1.0 {
            }
        }
```

The **maximum-ve-id** value is set to 10 to allow an increase in the number of PEs that could be a part of this VPLS instance.

PE-3 service creation

The following configuration creates a VPLS instance on PE-3, using a VE-ID value of 3.

```
# on PE-3:
configure {
    service {
        pw-template "PW1" {
            pw-template-id 1
            split-horizon-group {
                name "VPLS-SHG"
            }
        }
        vpls "VPLS1_PE-3" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
                route-distinguisher "65536:1"
                route-target {
                    export "target:65536:1"
                    import "target:65536:1"
                }
                pw-template-binding "PW1" {
                }
            }
            bgp-vpls {
                admin-state enable
                maximum-ve-id 10
                ve {
                    name "PE-3"
                    id 3
                }
            }
            sap 1/1/4:1.0 {
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
```

PE-1 service operation verification

The following command shows that the BGP-VPLS site is enabled on PE-1.

```
[]
A:admin@PE-1# show service id 1 bgp-vpls

===============================================================================
BGP VPLS Information
===============================================================================
Max Ve Id            : 10                  Admin State      : Enabled
VE Name              : PE-1                VE Id            : 1
PW Tmpl used         : 1
===============================================================================
```

The following command shows that the service is operationally up on PE-1:

```
[]
A:admin@PE-1# show service id 1 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 1                   Vpn Id            : 0
Service Type      : VPLS
MACSec enabled    : no
Name              : VPLS1_PE-1
---snip---

Admin State       : Up                  Oper State        : Up
MTU               : 1514
SAP Count         : 1                   SDP Bind Count    : 2
---snip---

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                             Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/4:1.0                          qinq         1522    1522    Up   Up
sdp:32766:4294967294 SB(192.0.2.2)     BgpVpls      0       1556    Up   Up
sdp:32767:4294967295 SB(192.0.2.3)     BgpVpls      0       1556    Up   Up
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

The SAP and SDPs are all operationally up. The **SB** flags for the SDPs signify Spoke-SDP and BGP.

The ingress labels for PE-2 and PE-3—the labels allocated by PE-1—can be seen as follows:

```
[]
A:admin@PE-1# show service id 1 sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
32766:4294967294 BgpVpls  192.0.2.2       Up      Up        524273    524270
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
32767:4294967295 BgpVpls  192.0.2.3       Up      Up        524274    524272
-------------------------------------------------------------------------------
Number of SDPs : 2
-------------------------------------------------------------------------------
===============================================================================
```

As can be seen from the following output, a BGP-VPLS NLRI update is sent to the route reflector (192.0.2.7) and is received by each PE.

PE-1 has sent the following BGP NLRI update for VPLS 1 to RR-7.

```
1 2021/01/26 10:54:39.689 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.7
"Peer 1: 192.0.2.7: UPDATE
Peer 1: 192.0.2.7 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 72
    Flag: 0x90 Type: 14 Len: 28 Multiprotocol Reachable NLRI:
        Address Family L2VPN
        NextHop len 4 NextHop 192.0.2.1
        [VPLS/VPWS] preflen 17, veid: 1, vbo: 1, vbs: 8, label-base: 524272, RD 65536:1
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x80 Type: 4 Len: 4 MED: 0
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:65536:1
        l2-vpn/vrf-imp:Encap=19: Flags=none: MTU=1514: PREF=0
"
```

The control flags within the extended community indicate the status of the VPLS instance.

The control flag D indicates that all attachment circuits are Down, or the VPLS is disabled. The flags are used in BGP-MH when determining which PEs are DF, see chapter BGP Multi-Homing for VPLS Networks.

When flags=none, then all attachment circuits are up. In the preceding example, no flags are present, but should all SAPs become operationally down, then the control flag D would be seen in the debug message. To simulate this, the SAP 1/1/4:1 is disabled on PE-1:

```
# on PE-1:
configure {
    service {
        vpls "VPLS1_PE-1" {
            sap 1/1/4:1.0 {
                admin-state disable
            }
```

All SAPs in VPLS 1 on PE-1 are operationally down, so PE-1 sends a BGP update message with control flag D set, as follows:

```
5 2021/01/26 11:09:10.688 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.7
"Peer 1: 192.0.2.7: UPDATE
Peer 1: 192.0.2.7 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 72
    Flag: 0x90 Type: 14 Len: 28 Multiprotocol Reachable NLRI:
        Address Family L2VPN
        NextHop len 4 NextHop 192.0.2.1
        [VPLS/VPWS] preflen 17, veid: 1, vbo: 1, vbs: 8, label-base: 524272, RD 65536:1
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x80 Type: 4 Len: 4 MED: 0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:65536:1
        l2-vpn/vrf-imp:Encap=19: Flags=D: MTU=1514: PREF=0
"
```

The SAP is re-enabled with the following command on PE-1:

```
# on PE-1:
configure {
    service {
        vpls "VPLS1_PE-1" {
            sap 1/1/4:1.0 {
                admin-state enable
            }
```

The BGP VPLS signaling parameters are also present in the BGP update message, namely the VE-ID of the PE within the VPLS instance, the VBO and VBS, and the label base. The target indicates the VPLS instance, which must be matched against the import RTs of the receiving PEs.

The signaling parameters can be seen within the BGP update with following command:

```
[]
A:admin@PE-1# show router bgp routes l2-vpn rd 65536:1 hunt
===============================================================================
 BGP Router ID:192.0.2.1        AS:65536       Local AS:65536
===============================================================================
---snip---

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
Route Type     : VPLS
Route Dist.    : 65536:1
VeId           : 1                      Block Size     : 8
Base Offset    : 1                      Label Base     : 524272
Nexthop        : 192.0.2.1
To             : 192.0.2.7
Res. Nexthop   : n/a
Local Pref.    : 100                    Interface Name : NotAvailable
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : 0
AIGP Metric    : None                   IGP Cost       : n/a
Connector      : None
Community      : target:65536:1
                 l2-vpn/vrf-imp:Encap=19: Flags=none: MTU=1514: PREF=0
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.7
Origin         : IGP
AS-Path        : No As-Path
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0

-------------------------------------------------------------------------------
Routes : 4
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

In this configuration example, PE-1 (192.0.2.1) with VE-ID =1 has sent an update with base offset (VBO) =1, block size (VBS) = 8, and label base 524272. This means that labels 524272 (LB) to 524279 (LB +VBS-1) are available as de-multiplexer labels, egress labels to be used to reach PE-1 for VPLS 1.

PE-2 receives this update from PE-1. This is seen as a valid VPLS BGP route from PE-1 through the route reflector with next-hop 192.0.2.1.

```
[]
A:admin@PE-2# show router bgp routes l2-vpn rd 65536:1
===============================================================================
 BGP Router ID:192.0.2.2        AS:65536       Local AS:65536
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP L2VPN Routes
===============================================================================
Flag  RouteType                    Prefix                              MED
      RD                           SiteId                              Label
      Nexthop                      VeId                    BlockSize   LocalPref
      As-Path                      BaseOffset              vplsLabelBa
                                                           se
-------------------------------------------------------------------------------
u*>i  VPLS                         -                       -           0
      65536:1                      -                                   -
      192.0.2.1                    1                       8           100
      No As-Path                   1                       524272
i     VPLS                         -                       -           0
      65536:1                      -                                   -
      192.0.2.2                    2                       8           100
      No As-Path                   1                       524270
u*>i  VPLS                         -                       -           0
      65536:1                      -                                   -
      192.0.2.3                    3                       8           100
      No As-Path                   1                       524272
-------------------------------------------------------------------------------
Routes : 3
===============================================================================
```

PE-2 uses this information in conjunction with its own VE-ID to calculate the egress label toward PE-1, using the condition VBO < VE-ID < (VBO+VBS).

The VE-ID of PE-2 is in the Label Block covered by VBO =1, thus,

Label calculation = label base + local VE-ID - Base offset

= 524272 + 2 - 1

Egress label used = 524273

This is verified using the following command on PE-2 where the egress label toward PE-1 (192.0.2.1) is 524273.

```
[]
A:admin@PE-2# show service id 1 sdp

===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
32766:4294967294 BgpVpls  192.0.2.3       Up      Up        524272    524273
32767:4294967295 BgpVpls  192.0.2.1       Up      Up        524270    524273
-------------------------------------------------------------------------------
Number of SDPs : 2
-------------------------------------------------------------------------------
===============================================================================
```

PE-3 also receives this update from PE-1 by the RR. This is seen as a valid VPLS BGP route from PE-1 with next-hop 192.0.2.1.

```
[]A:admin@PE-3# show router bgp routes l2-vpn rd 65536:1
===============================================================================
 BGP Router ID:192.0.2.3        AS:65536       Local AS:65536
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP L2VPN Routes
===============================================================================
Flag  RouteType                  Prefix                           MED
      RD                         SiteId                           Label
      Nexthop                    VeId                  BlockSize  LocalPref
      As-Path                    BaseOffset            vplsLabelBa
                                                       se
-------------------------------------------------------------------------------
u*>i  VPLS                       -                     -          0
      65536:1                    -                                -
      192.0.2.1                  1                     8          100
      No As-Path                 1                     524272
u*>i  VPLS                       -                     -          0
      65536:1                    -                                -
      192.0.2.2                  2                     8          100
      No As-Path                 1                     524270
i     VPLS                       -                     -          0
      65536:1                    -                                -
      192.0.2.3                  3                     8          100
      No As-Path                 1                     524272
-------------------------------------------------------------------------------
Routes : 3
===============================================================================
```

The VE-ID of PE-3 is also in the label block covered by block offset VBO =1.

Label calculation = label base + local VE-ID - VBO

= 524272 + 3 - 1

Egress label used = 524274

This is verified using the following command on PE-3 where egress label toward 192.0.2.1 is 524274.

```
[]
A:admin@PE-3# show service id 1 sdp
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
32766:4294967294 BgpVpls  192.0.2.2       Up      Up        524273    524272
32767:4294967295 BgpVpls  192.0.2.1       Up      Up        524272    524274
-------------------------------------------------------------------------------
Number of SDPs : 2
-------------------------------------------------------------------------------
===============================================================================
```

PE-2 service operation verification

The service is operationally up on PE-2, as follows.

```
[]
A:admin@PE-2# show service id 1 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 1                   Vpn Id            : 0
Service Type      : VPLS
MACSec enabled    : no
Name              : VPLS1_PE-2
---snip---

Admin State       : Up                  Oper State        : Up
MTU               : 1514
SAP Count         : 1                   SDP Bind Count    : 2
---snip---

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                              Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/4:1.0                           qinq         1522    1522    Up   Up
sdp:32766:4294967294 SB(192.0.2.3)      BgpVpls      0       1556    Up   Up
sdp:32767:4294967295 SB(192.0.2.1)      BgpVpls      0       1556    Up   Up
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

PE-2 de-multiplexer label calculation

In the same way that PE-1 allocates a label base (LB), block size (VBS), and base offset (VBO), PE-2 also allocates the same parameters for PE-1 and PE-3 to calculate the egress service label required to reach PE-2.

```
[]
A:admin@PE-2# show router bgp routes l2-vpn rd 65536:1 hunt
===============================================================================
 BGP Router ID:192.0.2.2        AS:65536       Local AS:65536
===============================================================================
---snip---
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
Route Type     : VPLS
Route Dist.    : 65536:1
VeId           : 2                      Block Size     : 8
Base Offset    : 1                      Label Base     : 524270
Nexthop        : 192.0.2.2
To             : 192.0.2.7
Res. Nexthop   : n/a
Local Pref.    : 100                    Interface Name : NotAvailable
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : 0
AIGP Metric    : None                   IGP Cost       : n/a
Connector      : None
Community      : target:65536:1
                 l2-vpn/vrf-imp:Encap=19: Flags=none: MTU=1514: PREF=0
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.7
Origin         : IGP
AS-Path        : No As-Path
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0

-------------------------------------------------------------------------------
Routes : 4
===============================================================================
```

This is verified using the following command on PE-1 to show the egress label toward PE-2 (192.0.2.2) where the egress label toward PE-2 = 524270 + 1 – 1 = 524270.

```
[]
A:admin@PE-1# show service id 1 sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
32766:4294967294 BgpVpls  192.0.2.2       Up      Up        524273    524270
32767:4294967295 BgpVpls  192.0.2.3       Up      Up        524274    524272
-------------------------------------------------------------------------------
Number of SDPs : 2
-------------------------------------------------------------------------------
===============================================================================
```

This is also verified using the following command on PE-3 to show the egress label toward PE-2 (192.0.2.2) where the egress label toward PE-2 = 524270 + 3 – 1 = 524272.

```
[]
A:admin@PE-3# show service id 1 sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
32766:4294967294 BgpVpls  192.0.2.2       Up      Up        524273    524272
32767:4294967295 BgpVpls  192.0.2.1       Up      Up        524272    524274
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Number of SDPs : 2
------------------------------------------------------------------------------
==============================================================================
```

PE-3 service operation verification

The following command shows that the service is operationally up on PE-3:

```
[]
A:admin@PE-3# show service id 1 base

==============================================================================
Service Basic Information
==============================================================================
Service Id        : 1                   Vpn Id            : 0
Service Type      : VPLS
MACSec enabled    : no
Name              : VPLS1_PE-3
---snip---

Admin State       : Up                  Oper State        : Up
MTU               : 1514
SAP Count         : 1                   SDP Bind Count    : 2
---snip---

------------------------------------------------------------------------------
Service Access & Destination Points
------------------------------------------------------------------------------
Identifier                             Type         AdmMTU  OprMTU  Adm  Opr
------------------------------------------------------------------------------
sap:1/1/4:1.0                          qinq         1522    1522    Up   Up
sdp:32766:4294967294 SB(192.0.2.2)     BgpVpls      0       1556    Up   Up
sdp:32767:4294967295 SB(192.0.2.1)     BgpVpls      0       1556    Up   Up
==============================================================================
* indicates that the corresponding row element may have been truncated.
```

```
[]
A:admin@PE-3# show service id 1 sdp

==============================================================================
Services: Service Destination Points
==============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
------------------------------------------------------------------------------
32766:4294967294 BgpVpls  192.0.2.2       Up      Up        524273    524272
32767:4294967295 BgpVpls  192.0.2.1       Up      Up        524272    524274
------------------------------------------------------------------------------
Number of SDPs : 2
------------------------------------------------------------------------------
==============================================================================
```

PE-3 de-multiplexer label verification

PE-3 also allocates the required parameters for PE-1 and PE-2 to calculate the egress service label required to reach PE-3.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

This is verified using the following command on PE-1 to show the egress label toward PE-3 (192.0.2.3) (524272) where egress label toward PE-2 = 524270. The Label Base equals 524272 on PE-3 and 524270 on PE-2.

```
[]
A:admin@PE-1# show service id 1 sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
32766:4294967294 BgpVpls  192.0.2.2       Up      Up        524273    524270
32767:4294967295 BgpVpls  192.0.2.3       Up      Up        524274    524272
-------------------------------------------------------------------------------
Number of SDPs : 2
-------------------------------------------------------------------------------
===============================================================================
```

This is also verified using the following command on PE-2 to show the egress label toward PE-3 (192.0.2.3) which is using auto-provisioned SDP 32766.

```
[]
A:admin@PE-2# show service id 1 sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
32766:4294967294 BgpVpls  192.0.2.3       Up      Up        524272    524273
32767:4294967295 BgpVpls  192.0.2.1       Up      Up        524270    524273
-------------------------------------------------------------------------------
Number of SDPs : 2
-------------------------------------------------------------------------------
===============================================================================
```

This example has shown that for VPLS instance with 3 PEs, not all labels allocated by a PE will be used by remote PEs as de-multiplexer service labels. There will be some wastage of label space, so there is a necessity to choose VE-IDs that keep this waste to a minimum.

The next example will show an even more wasteful use of labels by using a random choice of VE-IDs.

BGP VPLS using pre-provisioned SDP

It is possible to configure BGP-VPLS instances that use RSVP-TE transport tunnels. In this case, the SDP must be created with the MPLS LSPs mapped and with signaling set to BGP, as the service labels are signaled using BGP. The pseudowire template configured within the BGP-VPLS instance must be configured with **provisioned-sdp use**. This example also examines the effect of using VE-IDs that are not all within the same contiguous block.

Figure 58: BGP VPLS using pre-provisioned SDP shows an example of a VPLS instance where SDPs are pre-provisioned with RSVP-TE signaled transport tunnels.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 58: BGP VPLS using pre-provisioned SDP

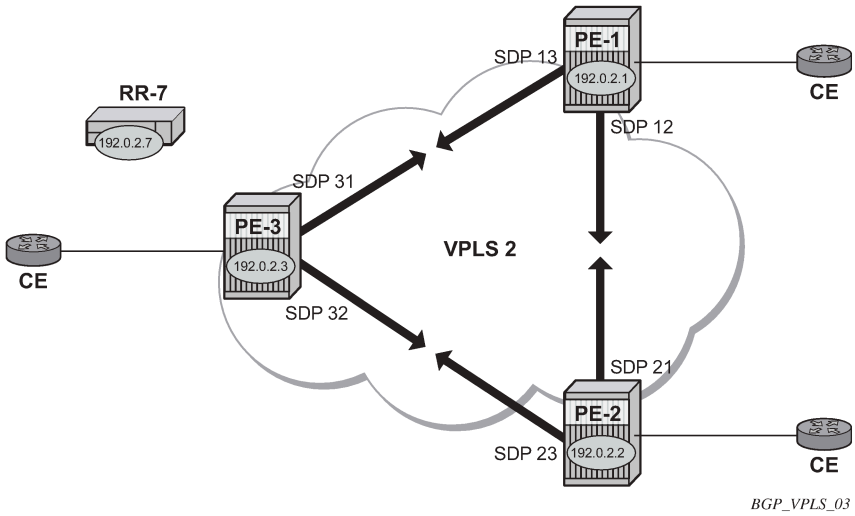

On the PEs, the following SDPs are configured with RSVP transport tunnels.

```
# on PE-1:
configure {
    service {
        sdp 12 {
            admin-state enable
            description "SDP-PE-1-PE-2_RSVP_BGP"
            delivery-type mpls
            signaling bgp
            far-end {
                ip-address 192.0.2.2
            }
            lsp "LSP-PE-1-PE-2" { }
        }
        sdp 13 {
            admin-state enable
            description "SDP-PE-1-PE-3_RSVP_BGP"
            delivery-type mpls
            signaling bgp
            far-end {
                ip-address 192.0.2.3
            }
            lsp "LSP-PE-1-PE-3" { }
        }
```

```
# on PE-2:
configure {
    service {
        sdp 21 {
            admin-state enable
            description "SDP-PE-2-PE-1_RSVP_BGP"
            delivery-type mpls
            signaling bgp
            far-end {
                ip-address 192.0.2.1
            }
            lsp "LSP-PE-2-PE-1" { }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
        sdp 23 {
            admin-state enable
            description "SDP-PE-2-PE-3_RSVP_BGP"
            delivery-type mpls
            signaling bgp
            far-end {
                ip-address 192.0.2.3
            }
            lsp "LSP-PE-2-PE-3" { }
        }
```

```
# on PE-3:
configure {
    service {
        sdp 31 {
            admin-state enable
            description "SDP-PE-3-PE-1_RSVP_BGP"
            delivery-type mpls
            signaling bgp
            far-end {
                ip-address 192.0.2.1
            }
            lsp "LSP-PE-3-PE-1" { }
        }
        sdp 32 {
            admin-state enable
            description "SDP-PE-3-PE-2_RSVP_BGP"
            delivery-type mpls
            signaling bgp
            far-end {
                ip-address 192.0.2.2
            }
            lsp "LSP-PE-3-PE-2" { }
        }
```

Pre-provisioned BGP-SDPs can also be used with BGP-VPLS. For reference, they are configured as follows:

```
# on PE-3:
configure {
    service {
        sdp 332 {
            admin-state enable
            delivery-type mpls
            signaling bgp
            far-end {
                ip-address 192.0.2.2
            }
        }
```

To create an SDP within a service that uses the RSVP transport tunnel, a pseudowire template is required that has the **provisioned-sdp use** parameter set. It is also possible to configure the **provisioned-sdp prefer** parameter, see chapter LDP VPLS Using BGP Auto-Discovery — Prefer Provisioned SDP.

Once again, an SHG is included to prevent forwarding between pseudowires.

The following pseudowire template is provisioned on all PEs:

```
# on PE-1, PE-2, and PE-3:
configure {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    service {
        pw-template "PW2" {
            pw-template-id 2
            provisioned-sdp use
            split-horizon-group {
                name "VPLS-SHG"
            }
        }
```

The following output shows the configuration required for a BGP-VPLS service using a pseudowire template configured for using pre-provisioned RSVP-TE SDPs.

```
# on PE-1:
configure {
    service {
        vpls "VPLS2_PE-1" {
            admin-state enable
            service-id 2
            customer "1"
            bgp 1 {
                route-distinguisher "65536:2"
                route-target {
                    export "target:65536:2"
                    import "target:65536:2"
                }
                pw-template-binding "PW2" {
                }
            }
            bgp-vpls {
                admin-state enable
                maximum-ve-id 100
                ve {
                    name "PE-1"
                    id 1
                }
            }
            sap 1/1/4:2.0 {
            }
        }
```

The RD and RT extended community values for VPLS 2 are different from the ones in VPLS 1. The VE-ID value for PE-1 can be the same as the one in VPLS 1, but these must be different within the same VPLS instance on the other PEs — PE-2 should not have VE-ID = 1.

On PE-2, the configuration is as follows with the VE-ID value equal to 20, which will result in a label from a different block:

```
# on PE-2:
configure {
    service {
        vpls "VPLS2_PE-2" {
            admin-state enable
            service-id 2
            customer "1"
            bgp 1 {
                route-distinguisher "65536:2"
                route-target {
                    export "target:65536:2"
                    import "target:65536:2"
                }
                pw-template-binding "PW2" {
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
            bgp-vpls {
                admin-state enable
                maximum-ve-id 100
                ve {
                    name "PE-2"
                    id 20
                }
            }
            sap 1/1/4:2.0 {
            }
        }
```

On PE-3, the configuration is as follows with the VE-ID value equal to 3:

```
# on PE-3:
configure {
    service {
        vpls "VPLS2_PE-3" {
            admin-state enable
            service-id 2
            customer "1"
            bgp 1 {
                route-distinguisher "65536:2"
                route-target {
                    export "target:65536:2"
                    import "target:65536:2"
                }
                pw-template-binding "PW2" {
                }
            }
            bgp-vpls {
                admin-state enable
                maximum-ve-id 100
                ve {
                    name "PE-3"
                    id 3
                }
            }
            sap 1/1/4:2.0 {
            }
        }
```

The service is operationally up on PE-1, as follows:

```
[]
A:admin@PE-1# show service id 2 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 2                   Vpn Id            : 0
Service Type      : VPLS
MACSec enabled    : no
Name              : VPLS2_PE-1
---snip---

Admin State       : Up                  Oper State        : Up
MTU               : 1514
SAP Count         : 1                   SDP Bind Count    : 2
---snip---

-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Service Access & Destination Points
------------------------------------------------------------------------------
Identifier                               Type         AdmMTU  OprMTU  Adm  Opr
------------------------------------------------------------------------------
sap:1/1/4:2.0                            qinq         1522    1522    Up   Up
sdp:12:4294967292 S(192.0.2.2)           BgpVpls      0       1556    Up   Up
sdp:13:4294967293 S(192.0.2.3)           BgpVpls      0       1556    Up   Up
==============================================================================
* indicates that the corresponding row element may have been truncated.
```

The SDPs 12 and 13 are the pre-provisioned SDPs on PE-1.

The service is operationally up on PE-2, as follows:

```
[]
A:admin@PE-2# show service id 2 base

==============================================================================
Service Basic Information
==============================================================================
Service Id        : 2                    Vpn Id            : 0
Service Type      : VPLS
MACSec enabled    : no
Name              : VPLS2_PE-2
---snip---

Admin State       : Up                   Oper State        : Up
MTU               : 1514
SAP Count         : 1                    SDP Bind Count    : 2
---snip---

------------------------------------------------------------------------------
Service Access & Destination Points
------------------------------------------------------------------------------
Identifier                               Type         AdmMTU  OprMTU  Adm  Opr
------------------------------------------------------------------------------
sap:1/1/4:2.0                            qinq         1522    1522    Up   Up
sdp:21:4294967292 S(192.0.2.1)           BgpVpls      0       1556    Up   Up
sdp:23:4294967293 S(192.0.2.3)           BgpVpls      0       1556    Up   Up
==============================================================================
* indicates that the corresponding row element may have been truncated.
```

The service is operationally up on PE-3, as follows:

```
[]
A:admin@PE-3# show service id 2 base

==============================================================================
Service Basic Information
==============================================================================
Service Id        : 2                    Vpn Id            : 0
Service Type      : VPLS
MACSec enabled    : no
Name              : VPLS2_PE-3
---snip---

Admin State       : Up                   Oper State        : Up
MTU               : 1514
SAP Count         : 1                    SDP Bind Count    : 2
---snip---

------------------------------------------------------------------------------
Service Access & Destination Points
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
Identifier                              Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/4:2.0                           qinq         1522    1522    Up   Up
sdp:31:4294967293 S(192.0.2.1)          BgpVpls      0       1556    Up   Up
sdp:32:4294967292 S(192.0.2.2)          BgpVpls      0       1556    Up   Up
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

PE-1 de-multiplexer label calculation

In the case of VPLS 1, all VE-IDs are in the range of a single label block. In the case of VPLS 2, the VE-IDs are in different blocks, for example, the VE-ID 20 is in a different block to VE-IDs 1 and 3.

As the label allocation is block-dependent, multiple label blocks must be advertised by each PE to encompass this.

Consider PE-1's BGP update NLRIs.

```
[]
A:admin@PE-1# show router bgp routes l2-vpn rd 65536:2 hunt
===============================================================================
 BGP Router ID:192.0.2.1        AS:65536       Local AS:65536
===============================================================================
---snip---

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
Route Type     : VPLS
Route Dist.    : 65536:2
VeId           : 1                     Block Size     : 8
Base Offset    : 1                     Label Base     : 524264
Nexthop        : 192.0.2.1
To             : 192.0.2.7
Res. Nexthop   : n/a
Local Pref.    : 100                   Interface Name : NotAvailable
Aggregator AS  : None                  Aggregator     : None
Atomic Aggr.   : Not Atomic            MED            : 0
AIGP Metric    : None                  IGP Cost       : n/a
Connector      : None
Community      : target:65536:2
                 l2-vpn/vrf-imp:Encap=19: Flags=none: MTU=1514: PREF=0
Cluster        : No Cluster Members
Originator Id  : None                  Peer Router Id : 192.0.2.7
Origin         : IGP
AS-Path        : No As-Path
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                     Dest Class     : 0

Route Type     : VPLS
Route Dist.    : 65536:2
VeId           : 1                     Block Size     : 8
Base Offset    : 17                    Label Base     : 524256
Nexthop        : 192.0.2.1
To             : 192.0.2.7
Res. Nexthop   : n/a
Local Pref.    : 100                   Interface Name : NotAvailable
Aggregator AS  : None                  Aggregator     : None
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Atomic Aggr.   : Not Atomic             MED            : 0
AIGP Metric    : None                   IGP Cost       : n/a
Connector      : None
Community      : target:65536:2
                 l2-vpn/vrf-imp:Encap=19: Flags=none: MTU=1514: PREF=0
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.7
Origin         : IGP
AS-Path        : No As-Path
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0

-------------------------------------------------------------------------------
Routes : 8
===============================================================================
```

Two NLRIs updates are sent to the route reflector, with the following label parameters:

1. LB = 524264, VBS = 8, VBO = 1

2. LB = 524256, VBS = 8, VBO = 17

PE-2 has a VE-ID of 20. Applying the condition VBO < VE-ID < (VBO+VBS)

- Update 1: LB = 524264, VBS = 8, VBO = 1
- VBO < VE-ID for VE-ID = 20 is true
- VE-ID < (VBO+VBS) for VE-ID = 20 is false.
- PE-2 cannot choose a label from this block.
- Update 2: LB = 524256, VBS = 8, VBO = 17
- VBO < VE-ID for VE-ID = 20 is true
- VE-ID < (VBO+VBS) for VE-ID = 20 is true.
- PE-2 chooses label 524256 + 20 - 17 = 524259 (LB + VEID - VBO)

The egress label chosen is verified by examining the egress label toward PE-1 (192.0.2.1) on PE-2.

```
[]
A:admin@PE-2# show service id 2 sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
21:4294967292    BgpVpls  192.0.2.1       Up      Up        524254    524259
23:4294967293    BgpVpls  192.0.2.3       Up      Up        524256    524259
-------------------------------------------------------------------------------
Number of SDPs : 2
-------------------------------------------------------------------------------
===============================================================================
```

PE-3 has a VE-ID of 3. Applying the condition VBO < VE-ID < (VBO+VBS)

- Update 1: LB = 524264, VBS = 8, VBO = 1
- VBO < VE-ID for VE-ID = 3 is true
- VE-ID < (VBO+VBS) for VE-ID = 3 is true.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- PE-3 chooses label 524264 + 3 - 1 = 524266 (LB + VEID - VBO)
- Update 2: LB = 524256, VBS = 8, VBO = 17
- VBO < VE-ID for VE-ID = 3 is false
- VE-ID < (VBO+VBS) for VE-ID = 3 is true.
- PE-3 cannot choose a label from this block.

The egress label chosen is verified by examining the egress label toward PE-1 (192.0.2.1) on PE-3.

```
[]
A:admin@PE-3# show service id 2 sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
31:4294967293    BgpVpls  192.0.2.1       Up      Up        524264    524266
32:4294967292    BgpVpls  192.0.2.2       Up      Up        524259    524256
-------------------------------------------------------------------------------
Number of SDPs : 2
-------------------------------------------------------------------------------
===============================================================================
```

To illustrate the allocation of label blocks by a PE, against the actual use of the same labels, consider the following. When BGP updates from each PE signal the multiplexer labels in blocks of eight, the allocated label values are added to the in-use pool. First check what label range can be allocated dynamically.

```
[]
A:admin@PE-1# show router mpls-labels label-range

===============================================================================
Label Ranges
===============================================================================
Label Type      Start Label End Label   Aging       Available   Total
-------------------------------------------------------------------------------
Static          32          18431       -           18400       18400
Dynamic         18432       524287      0           505824      505856
    Seg-Route   0           0           -           0           0
===============================================================================
```

Verify which labels in the dynamic range are in use. The label pool of PE-1 can be verified as per the following output which shows labels used along with the associated protocol:

```
[]
A:admin@PE-1# show router mpls-labels label 18432 524287 in-use

===============================================================
MPLS Labels from 18432 to 524287 (In-use)
===============================================================
Label               Label Type          Label Owner
---------------------------------------------------------------
524256              dynamic             BGP
524257              dynamic             BGP
524258              dynamic             BGP
524259              dynamic             BGP
524260              dynamic             BGP
524261              dynamic             BGP
524262              dynamic             BGP
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
524263              dynamic             BGP
524264              dynamic             BGP
524265              dynamic             BGP
524266              dynamic             BGP
524267              dynamic             BGP
524268              dynamic             BGP
524269              dynamic             BGP
524270              dynamic             BGP
524271              dynamic             BGP
524272              dynamic             BGP
524273              dynamic             BGP
524274              dynamic             BGP
524275              dynamic             BGP
524276              dynamic             BGP
524277              dynamic             BGP
524278              dynamic             BGP
524279              dynamic             BGP
524280              dynamic             RSVP
524281              dynamic             RSVP
524282              dynamic             ILDP
524283              dynamic             ILDP
524284              dynamic             ILDP
524285              dynamic             ILDP
524286              dynamic             ILDP
524287              dynamic             ILDP
-------------------------------------------------------------------
In-use labels (Owner: All) in specified range   : 32
In-use labels in entire range                   : 32
===================================================================
```

This shows that 24 labels have been allocated for use by BGP. Of this number, 16 labels have been allocated for use by PEs within VPLS 2 to communicate with PE-1, the blocks with label base 524256 and with label base 524264.

There are only two neighboring PEs within this VPLS instance, so only two labels will ever be used in the data plane for traffic destined for PE-1. These are 524259 and 524266. The remaining labels have no PE with the associated VE-ID that can use them.

Once again, this case emphasizes that to reduce label wastage, contiguous VE-IDs in the range (N..N+7) per VPLS should be chosen, where N>0.

Conclusion

BGP-VPLS allows the delivery of Layer 2 VPN services to customers where BGP is commonly used. The examples presented in this chapter show the configuration of BGP-VPLS together with the associated show outputs which can be used for verification and troubleshooting.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Black-hole MAC for EVPN Loop Protection

This chapter provides information about Black-hole MAC for EVPN Loop Protection.

Topics in this chapter include:

Applicability

This chapter was initially written based on SR OS Release 15.0.R4, but the MD-CLI in the current edition corresponds to SR OS Release 21.2.R2. Black-hole MAC for EVPN loop protection is supported in SR OS Release 15.0.R1, and later.

Chapters Auto-Learn MAC Protect in EVPN and Conditional Static Black-Hole MAC in EVPN are prerequisite reading.

Overview

Service providers are migrating VPLS networks to EVPN and require the same or better loop protection mechanisms, such as **mac-move** or **auto-learn-mac-protect** (ALMP). Chapter Auto-Learn MAC Protect in EVPN describes how traffic is protected in "static" networks, where the CEs do not move to a different port or PE, and MAC addresses are always learned first on the correct SAP/SDP-bindings. However, ALMP does not provide a loop protection solution in EVPN networks that require mobility and ALMP has issues with all-active multi-homing. Since mobility and all-active multi-homing are two of the key advantages of EVPN compared to VPLS, an alternate loop protection mechanism is required. This chapter describes an example for the black-hole based loop protection solution, based on *draft-snr-bess-evpn-loop-protect*.

Figure 59: Black-hole MAC for EVPN loop protection shows a topology using black-hole MAC for EVPN loop protection.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 59: Black-hole MAC for EVPN loop protection

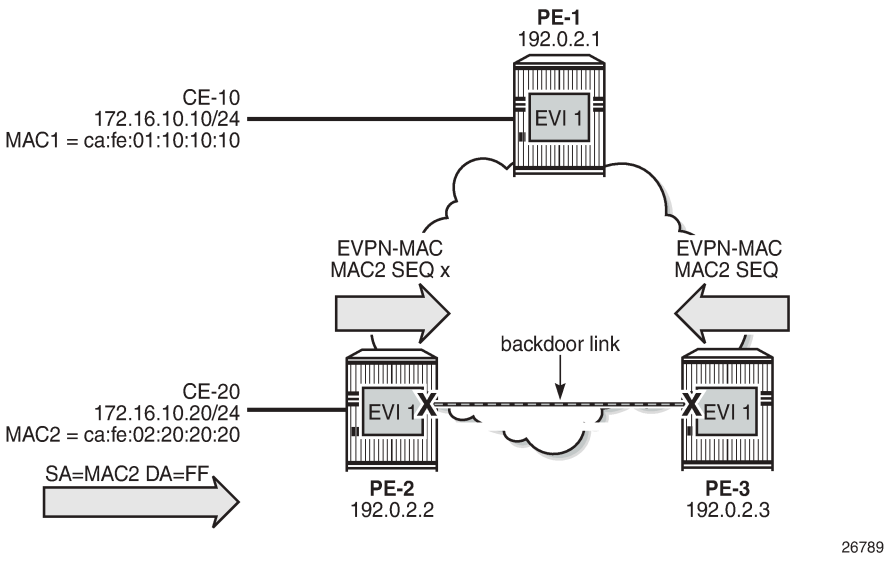

VPLS 1 with EVI 1 is configured on all PEs. A backdoor link exists between PE-2 and PE-3 (in this case, caused by misconfiguration: additional SAPs are configured in VPLS 1). When CE-20 sends Broadcast, Unknown unicast, or Multicast (BUM) traffic, its source address MAC2 is learned by PE-2, which sends an EVPN-MAC route for MAC2 to its BGP peers. PE-2 floods the frame to its EVPN-MPLS destinations (PE-1 and PE-3) as well as its local SAPs (including the backdoor link to PE-3).

PE-3 receives the EVPN-MAC route from PE-2, but due to the backdoor link, it also learns MAC2 on its local SAP. Following the MAC mobility procedures, PE-3 advertises MAC2 with a higher sequence number to its BGP peers. PE-3 floods the frame to its EVPN-MPLS destinations and to its local SAPs.

Note:
The preceding simplified description assumes that PE-3 receives the EVPN-MAC route prior to learning MAC2 from the backdoor link, which may or may not be the case. Regardless of how MAC2 is learned, the MAC duplication procedures are invoked.

PE-2 and PE-3 keep learning and advertising MAC2 until the configured number of MAC moves (**num-moves**) has been reached. Then, MAC2 is detected as duplicate and will not be advertised again until the **retry** interval has expired.

If the **mac-duplication blackhole** option is enabled, MAC2 will be added to the FDB as black-hole MAC, so traffic with MAC DA = MAC2 will be discarded. Also, MAC addresses assigned to a black-hole destination are considered as protected, so traffic with MAC SA = MAC2 will not be forwarded due to one of the following reasons:

- When the SAPs/SDP-bindings or BGP-EVPN MPLS/VXLAN destinations are configured with **fdb>protected-src-mac-violation-action discard**, the frames are discarded before any MAC SA is learned or the MAC DA is looked up.
- When the SAP is configured with fdb>protected-src-mac-violation-action sap-oper-down, an incoming frame with MAC SA = black-hole MAC causes the system to bring down the corresponding SAP.

Assuming PE-3 detects MAC2 as duplicate and installs it as black-hole MAC, PE-3 will discard the broadcast frames with MAC SA = MAC2, so the loop is broken, whereas the legitimate traffic between CE-10 and CE-20 is allowed (assuming PE-2 does not black-hole MAC2).

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Black-hole MAC duplication is enabled with **blackhole true** in the **mac-duplication** context, as follows:

```
[ex:/configure service vpls "VPLS 1" bgp-evpn]
A:admin@PE-3# mac-duplication ?

 mac-duplication

 blackhole             - Enable black hole dup MAC configuration
 detect                + Enter the detect context
 retry                 - BGP EVPN MAC duplication retry
```

```
# on PE-3:
configure {
    service {
        vpls "VPLS 1" {
            bgp-evpn {
                mac-duplication {
                    blackhole true
```

When enabled, the operation is as follows:

- Each node that learns a MAC address that has been advertised by a BGP peer will send an EVPN-MAC route for that MAC address with a higher sequence number. When the number of MAC moves exceeds the configured threshold (by default, five MAC moves in three minutes), the MAC address is detected as duplicate and no EVPN-MAC routes will be sent for that MAC address until the retry interval (default nine minutes) has elapsed.
- When MAC2 is detected as duplicate, the system will:
  - Add MAC2 to the duplicate MAC list
  - Add MAC2 in the FDB as protected MAC associated with a black-hole endpoint (type **EvpnD:P** and source identifier **black-hole**)
    - Incoming frames with MAC DA = MAC2 will be discarded based on a MAC lookup in the FDB.
    - MAC addresses assigned to a black-hole destination are protected and incoming frames with MAC SA = MAC2 will be discarded or the system will bring down the SAP/SDP-binding, depending on the **protected-src-mac-violation-action** on the SAP/SDP/EVPN endpoint.

The following output shows the FDB with black-hole MAC address ca:fe:02:20:20:20 (type EvpnD:P):

```
[/]
A:admin@PE-3# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
1          ca:fe:01:10:10:10 mpls:                   Evpn     04/28/21 09:59:12
                             192.0.2.1:524284
           ldp:65537
1          ca:fe:02:20:20:20 black-hole              EvpnD:P  04/28/21 09:59:12
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The duplicate MAC address will be removed from the FDB and the process will be restarted in the following cases:

- Retry interval events:
  - When the retry interval expires.
  - When the user configures **retry never** on the service that detected the duplicate MAC address.
- MAC relearning events:
  - When the remote PE withdraws the MAC address (due to aging or **clear service fdb**). Local attempts to clear a black-hole MAC (via **clear service fdb**) will fail because the type of the MAC entry is not "learned", but "EvpnD:P".
  - When configuring a local conditional static MAC address (CStatic:P) prevents the EvpnD:P entry for the same MAC address from being installed in the FDB as black-hole, if the SAP/SDP-binding where the MAC is configured is operationally up.
- CPM switchover event

Configuration

Figure 60: Example topology shows the example topology with three PEs and two CEs. A loop will occur when CE-20 sends Broadcast, Unknown unicast, or Multicast (BUM) traffic. Traffic between PE-2 and PE-3 will be sent over the regular router interfaces between the PEs, but also over the backdoor link (SAP 1/1/2:1 in VPLS 1 on PE-2 and SAP 1/1/1:1 in VPLS 1 on PE-3).

Figure 60: Example topology

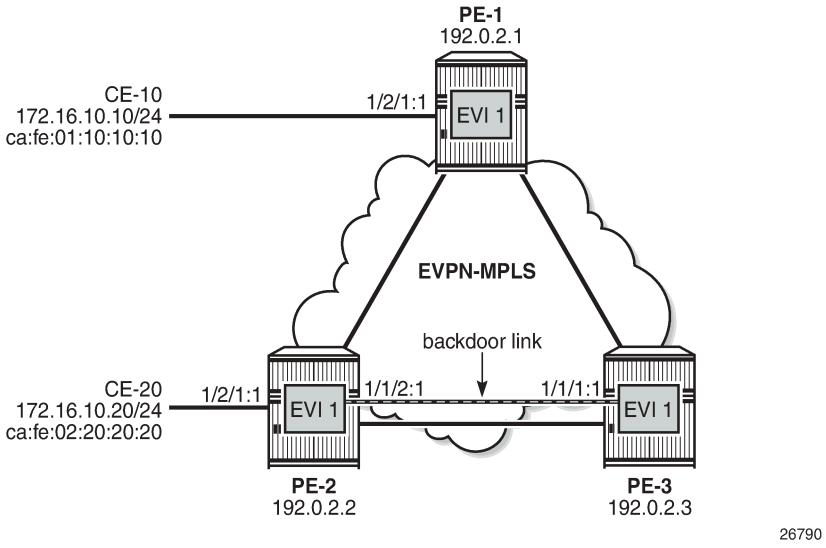

The initial configuration includes:

- Cards, MDAs, ports
- Router interfaces
- IS-IS on all router interfaces (alternatively, OSPF can be used)
- LDP on all router interfaces

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Enable black-hole MAC duplication detection in EVPN

BGP is configured for address family EVPN on all PEs with PE-3 as route reflector. The following is the BGP configuration on PE-3:

```
# on PE-3:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            rapid-withdrawal true
            split-horizon true
            rapid-update {
                evpn true
            }
            group "internal" {
                peer-as 64500
                family {
                    evpn true
                }
                cluster {
                    cluster-id 192.0.2.3
                }
            }
            neighbor "192.0.2.1" {
                group "internal"
            }
            neighbor "192.0.2.2" {
                group "internal"
            }
        }
```

VPLS 1 is configured on all PEs with BGP-EVPN and MAC duplication enabled; on PE-2, as follows:

```
# on PE-2:
configure {
    service {
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
            }
            bgp-evpn {
                evi 1
                mac-duplication {
                    retry 2     # Duplicate MACs are released after retry interval
                    blackhole true
                    detect {
                        num-moves 3     # speed up MAC-duplication detection
                        window 1        # speed up MAC-duplication detection
                    }
                }
                mpls 1 {
                    admin-state enable
                    auto-bind-tunnel {
                        resolution any
                    }
                    fdb {
                        protected-src-mac-violation-action discard
                    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
            }
            sap 1/1/2:1 {          # backdoor link to PE-3
            }
            sap 1/2/1:1 {          # to CE-20
            }
        }
```

To speed up MAC duplication detection, MAC duplication is detected after three MAC moves (default: five MAC moves). To shorten the retry interval, the time window is reduced to one minute (default: three minutes). When a MAC address has been detected as duplicated, the system removes the duplicate MAC entry after a retry interval of two minutes (default: nine minutes). The retry interval must be at least twice the time window for MAC duplication detection.

On the EVPN-MPLS endpoints, **protected-src-mac-violation-action discard** must be configured. When MAC address ca:fe:02:20:20:20 is detected on PE-3 as a duplicate MAC address that is black-holed, the EVPN-MPLS endpoints on PE-3 should discard all frames with MAC SA ca:fe:02:20:20:20.

The configuration on the other PEs is similar; only the SAPs are different. VPLS 1 on PE-1 has SAP 1/2/1:1 to CE-10, but no SAP to a backdoor link; VPLS 1 on PE-3 has SAP 1/1/1:1 to the backdoor link to PE-2, but no SAP to a CE.

When CE-20 sends BUM traffic, its MAC SA ca:fe:02:20:20:20 is learned by PE-2 and advertised in EVPN-MAC routes. Because of the backdoor link to PE-3, PE-3 also learns MAC SA ca:fe:02:20:20:20 and advertises it to its BGP peers. The MAC-mobility sequence number is increased until the threshold of three MAC moves is reached. The following BGP EVPN-MAC route with sequence number 2 is sent by PE-2 to PE-3:

```
# on PE-2:
17 2021/04/28 09:59:11.599 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 89
    Flag: 0x90 Type: 14 Len: 44 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
        Type: EVPN-MAC Len: 33 RD: 192.0.2.2:1 ESI: ESI-0, tag: 0, mac len: 48
                       mac: ca:fe:02:20:20:20, IP len: 0, IP: NULL, label1: 8388544
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:1
        bgp-tunnel-encap:MPLS
        mac-mobility:Seq:2
"
```

The FDB on PE-2 shows that MAC ca:fe:02:20:20:20 has been learned on the SAP toward CE-20 (but it could also have been learned on the backdoor SAP or even be black-holed), as follows:

```
[/]
A:admin@PE-2# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
1          ca:fe:01:10:10:10 mpls:                   Evpn     04/28/21 09:59:12
                             192.0.2.1:524284
           ldp:65537
1          ca:fe:02:20:20:20 sap:1/2/1:1             L/0      04/28/21 09:59:12
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

The following FDB on PE-3 shows that MAC ca:fe:02:20:20:20 has been detected as a duplicate and protected MAC (type EvpnD:P) associated with a black-hole endpoint:

```
[/]
A:admin@PE-3# show service id 1 fdb mac ca:fe:02:20:20:20

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
1          ca:fe:02:20:20:20 black-hole              EvpnD:P  04/28/21 09:59:12
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

The following BGP-EVPN information for VPLS 1 on PE-3 shows the settings for MAC duplication detection, and the number of and list of detected duplicate MAC addresses:

```
[/]
A:admin@PE-3# show service id 1 bgp-evpn

===============================================================================
BGP EVPN Table
===============================================================================
MAC Advertisement  : Enabled            Unknown MAC Route  : Disabled
CFM MAC Advertise  : Disabled
Creation Origin    : manual
MAC Dup Detn Moves : 3                  MAC Dup Detn Window: 1
MAC Dup Detn Retry : 2                  Number of Dup MACs : 1
MAC Dup Detn BH    : Enabled
IP Route Advert    : Disabled
Sel Mcast Advert   : Disabled

EVI                : 1
Ing Rep Inc McastAd: Enabled
Accept IVPLS Flush : Disabled

-------------------------------------------------------------------------------
Detected Duplicate MAC Addresses            Time Detected
-------------------------------------------------------------------------------
ca:fe:02:20:20:20                           04/28/2021 09:59:12
-------------------------------------------------------------------------------
===============================================================================
---snip---
```

The following message is logged in log "99" on PE-3 when VPLS 1 has detected duplicate MACs:

```
# on PE-3:
69 2021/04/28 10:04:40.266 UTC MINOR: SVCMGR #2331 Base
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
"VPLS Service 1 has MAC(s) detected as duplicates by EVPN mac-duplication detection."
```

MAC address ca:fe:02:20:20:20 remains in the FDB as duplicate and black-holed until the retry interval expires, as follows:

```
[ex:/configure service vpls "VPLS 1" bgp-evpn mac-duplication]
A:admin@PE-3# retry ?

 retry (<number> | <keyword>)
 <number>   - <2..60>  - minutes
 <keyword>  - never    - minutes
 Default    - 9

    BGP EVPN MAC duplication retry
```

By default, the retry interval is nine minutes, but in this example, it is set to two minutes, which is the minimum value. The retry interval must be at least twice the time window for MAC duplication detection, which is by default three minutes, but reduced to one minute in this example. The following error is raised when attempting to configure a retry interval of two minutes for a detection time window of three minutes:

```
*[ex:/configure service vpls "VPLS 1" bgp-evpn mac-duplication]
A:admin@PE-3# commit
MINOR: MGMT_CORE #3001: configure service vpls "VPLS 1" bgp-evpn mac-duplication retry - mac-
duplication detection window should be less than or equal to half of retry time
```

After the retry interval expires, the MAC duplication is released.

Log "99" shows the following message when VPLS 1 no longer has duplicate MAC addresses:

```
# on PE-3:
70 2021/04/28 10:06:43.398 UTC MINOR: SVCMGR #2332 Base
"VPLS Service 1 no longer has MAC(s) detected as duplicates by EVPN mac-duplication detection."
```

MAC address ca:fe:02:20:20:20 remains in the FDB with type Evpn instead of EvpnD:P. BGP routes only disappear after a withdraw message has been received, whereas locally learned MAC addresses are flushed.

```
[/]
A:admin@PE-3# show service id 1 fdb mac ca:fe:02:20:20:20

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
1          ca:fe:02:20:20:20 mpls:                   Evpn     04/28/21 10:06:43
                             192.0.2.2:524284
           ldp:65538
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Clear commands

The following FDB entry on PE-3 of type EvpnD:P cannot be cleared with a normal FDB **clear** command:

```
[/]
A:admin@PE-3# show service id 1 fdb mac ca:fe:02:20:20:20

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
1          ca:fe:02:20:20:20 black-hole              EvpnD:P  04/28/21 10:07:52
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

The following error is raised when attempting to clear this FDB entry:

```
[/]
A:admin@PE-3# clear service id 1 fdb mac ca:fe:02:20:20:20
MAJOR: LOG #1202: Cannot perform clear operation - Entry is not of learned type
```

Log "99" shows the following message:

```
72 2021/04/28 10:08:17.960 UTC INDETERMINATE: LOGGER #2010 Base Clear SVCMGR
"Clear function clearSvcIdFdbMac has been run with parameters: svc-id="1" mac=
"ca:fe:02:20:20:20".  The completion result is: failure.  Additional error text, if any, is:
 Entry is not of learned type"
```

The following **clear** command releases the MAC duplication from the entry in the FDB, but it does not remove the entry from the FDB if it was learned from EVPN. The type is changed from EvpnD:P to Evpn.

```
[/]
A:admin@PE-3# clear service id 1 evpn mac-dup-detect ieee-address ca:fe:02:20:20:20
```

```
[/]
A:admin@PE-3# show service id 1 fdb mac ca:fe:02:20:20:20

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
1          ca:fe:02:20:20:20 mpls:                   Evpn     04/28/21 10:09:50
                             192.0.2.2:524284
           ldp:65538
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

Instead of clearing the MAC duplication state for one specific MAC address, all duplicate MAC addresses can be cleared by the following command:

```
[/]
A:admin@PE-3# clear service id 1 evpn mac-dup-detect all
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

When the MAC duplication is released, VPLS 1 no longer has duplicate MAC addresses detected, as follows:

```
[/]
A:admin@PE-3# show service id 1 bgp-evpn | match "Detected" pre-lines 1 post-lines 4
-------------------------------------------------------------------------------
Detected Duplicate MAC Addresses             Time Detected
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
===============================================================================
===============================================================================
```

Log "99" shows the following messages related to the **clear** commands:

```
76 2021/04/28 10:10:13.078 UTC INDETERMINATE: LOGGER #2010 Base Clear SVCMGR
"Clear function cliClearSvcIdEvpnDupDetMacAll has been run with parameters: svc-id="1".  The
 completion result is: success.  Additional error text, if any, is: "

75 2021/04/28 10:09:49.947 UTC INDETERMINATE: LOGGER #2010 Base Clear SVCMGR
"Clear function cliClearSvcIdEvpnDupDetMac has been run with parameters: svc-id="1"mac=
"ca:fe:02:20:20:20".  The completion result is: success.  Additional error text, if any, is: "
```

Restrict Protected Source option

By default, the frames with MAC SA or DA equal to the duplicate MAC address are discarded, but the SAP/SDP-binding where the frame enters the VPLS remains operationally up. With the **protected-src-mac-violation-action sap-oper-down**, the system will bring the SAP down where the frame with duplicate source MAC enters. The configuration on PE-2 and PE-3 is modified with **protected-src-mac-violation-action sap-oper-down** on the SAP to the backdoor link, as follows:

```
# on PE-2:
configure {
    service {
        vpls "VPLS 1" {
            sap 1/1/2:1 {
                fdb {
                    protected-src-mac-violation-action sap-oper-down

# on PE-3:
configure {
    service {
        vpls "VPLS 1" {
            sap 1/1/1:1 {
                fdb {
                    protected-src-mac-violation-action sap-oper-down
```

When CE-20 sends BUM traffic, PE-3 detects MAC ca:fe:02:20:20:20 as duplicate. Log "99" shows that a duplicate MAC address has been detected, that protected MAC address ca:fe:02:20:20:20 has been received on SAP 1/1/1:1 in VPLS 1, and that the status of SAP 1/1/1:1 in VPLS 1 is changed to operationally down, with flag **RxProtSrcMac** indicating that a protected source MAC has been received.

```
80 2021/04/28 10:11:40.885 UTC MINOR: SVCMGR #2203 Base
"Status of SAP 1/1/1:1 in service 1 (customer 1) changed to admin=up oper=down flags=RxProtSrc
Mac "

79 2021/04/28 10:11:40.885 UTC MINOR: SVCMGR #2208 Base
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
"Protected MAC ca:fe:02:20:20:20 received on SAP 1/1/1:1 in service 1. The SAP will be
 disabled."

78 2021/04/28 10:11:39.886 UTC MINOR: SVCMGR #2331 Base
"VPLS Service 1 has MAC(s) detected as duplicates by EVPN mac-duplication detection."
```

The following shows that SAP 1/1/1:1 in VPLS 1 on PE-3 is operationally down with flag RxProtSrcMac:

```
[/]
A:admin@PE-3# show service id 1 sap 1/1/1:1

===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 1
SAP                : 1/1/1:1                 Encap             : q-tag
Description        : (Not Specified)
Admin State        : Up                      Oper State        : Down
Flags              : RxProtSrcMac
Multi Svc Site     : None
Last Status Change : 04/28/2021 10:11:41
Last Mgmt Change   : 04/28/2021 10:11:19
===============================================================================
```

The only way to re-enable the SAP is to disable and enable the SAP, as follows:

```
# on PE-3:
configure exclusive
    service {
        vpls "VPLS 1" {
            sap 1/1/1:1
                admin-state disable
                commit
                admin-state enable
                commit
```

```
[/]
A:admin@PE-3# show service id 1 sap

===============================================================================
SAP(Summary), Service 1
===============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.  Egr.   Adm  Opr
                                          QoS   Fltr    QoS   Fltr
-------------------------------------------------------------------------------
1/1/1:1                         1          1     none    1     none   Up   Up
-------------------------------------------------------------------------------
Number of SAPs : 1
-------------------------------------------------------------------------------
===============================================================================
```

Black-hole MAC duplication in all-active multi-homing

Figure 61: Example topology with all-active multi-homing shows the example topology with all-active multi-homing.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 61: Example topology with all-active multi-homing

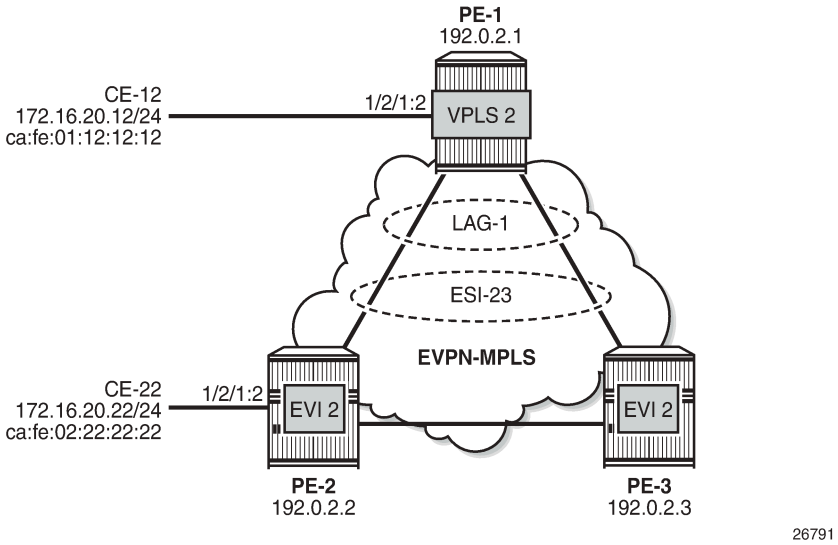

In this topology, the backdoor link is removed. On PE-1, VPLS 2 is configured without EVPN; on PE-2 and PE-3, VPLS 2 is configured with EVPN-MPLS. LAG 1 is configured on the PEs and Ethernet Segment (ES) ESI-23 is created on PE-2 and PE-3, as follows:

```
# on PE-2, PE-3:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "ESI-23" {
                        admin-state enable
                        esi 01:00:00:00:00:23:00:00:00:01
                        multi-homing-mode all-active
                        df-election {
                            es-activation-timer 3
                        }
                        association {
                            lag "lag-1" {
                            }
                        }
                    }
```

The reason why black-hole MAC duplication should be configured instead of ALMP is the following. When ALMP is configured on SAP lag-1:2 on PE-2 and PE-3, MAC address ca:fe:01:12:12:12 of CE-12 is learned and protected on the SAP on both PEs. Traffic sent from CE-12 to CE-22 that is hashed over the direct link between PE-1 and PE-2 will reach its destination. Traffic that is hashed over the link between PE-1 and PE-3 will be forwarded by PE-3 to PE-2, but PE-2 will drop the traffic because it contains a MAC SA that is protected locally, as shown in Figure 62: Traffic dropped when ALMP is configured in all-active multi-homing.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 62: Traffic dropped when ALMP is configured in all-active multi-homing

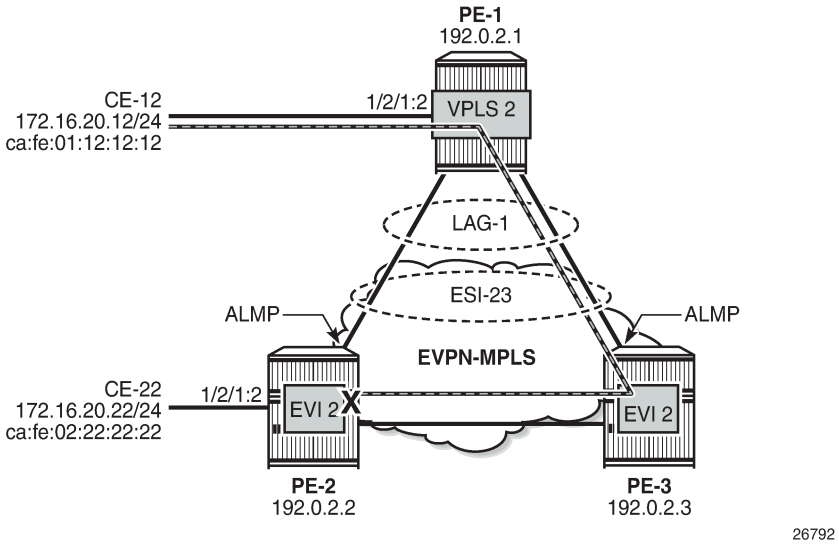

When black-hole MAC duplication is configured instead of ALMP, traffic hashed on the link to PE-3 is forwarded to PE-2 and to CE-22. This is because MAC duplication is ES-aware and the same MAC seen on the same ES in two different PEs will never be detected as duplicate.

The configuration of VPLS 2 in PE-2 is as follows:

```
# on PE-2:
configure {
    service {
        vpls "VPLS 2" {
            admin-state enable
            service-id 2
            customer "1"
            bgp 1 {
            }
            bgp-evpn {
                evi 2
                mac-duplication {
                    blackhole true
                }
                mpls 1 {
                    admin-state enable
                    auto-bind-tunnel {
                        resolution any
                    }
                    fdb {
                        protected-src-mac-violation-action discard
                    }
                }
            }
            sap 1/2/1:2 {
            }
            sap lag-1:2 {
            }
        }
```

The configuration of VPLS 2 on PE-3 is similar.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Conclusion

Black-hole MAC for EVPN MAC duplication protects EVPN services against customer-created backdoors or loops, while supporting MAC mobility and all-active multi-homing.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Conditional Static Black-Hole MAC in EVPN

This chapter provides information about Conditional Static Black-Hole MAC in EVPN.

Topics in this chapter include:

Applicability

This chapter was initially written for SR OS Release 14.0.R6, but the MD-CLI in the current edition is based on SR OS Release 21.2.R1. Conditional static black-hole MAC is supported on EVPN services only, including EVPN-VXLAN and EVPN-MPLS, in SR OS Release 14.0.R1, and later.

Overview

A static black-hole MAC address is a local FDB record associated with a black-hole instead of a SAP or SDP-binding. Black-hole MAC addresses offer a scalable way to filter frames in the data plane based on MAC DA or SA, regardless of how the frame is arriving in the system. Black-hole MAC addresses can be configured in EVPN in the following ways:

- Static configured black-hole MAC address
- Anti-spoof MAC address in proxy Address Resolution Protocol/Neighbor Discovery (proxy-ARP/ND)
- MAC-duplication black-hole (supported in SR OS Release 15.0.R1, and later), see chapter Black-hole MAC for EVPN Loop Protection

When a specific MAC address is configured as a static black-hole MAC address, all frames with MAC DA equal to this black-hole MAC address will be dropped. Also, black-hole MAC addresses are treated as protected MAC addresses, which allows filtering on MAC SA; see chapter Auto-Learn MAC Protect in EVPN.

The default behavior on the SAP/SDP-bindings is Restricted Protected Source Discard Frame (RPS-DF). Therefore, all frames with MAC SA equal to the black-hole MAC address will, by default, be dropped on the SAP/SDP-binding where the frames enter the service. Instead of dropping the frames, the entire SAP/SDP-binding can be brought operationally down, if the SAP/SDP-binding is explicitly configured with Restricted Protected Source (RPS) with **sap-oper-down/sdp-bind-oper-down**. The SAP/SDP-binding can only be brought up manually by disabling and re-enabling the SAP/SDP-binding. On the EVPN endpoints between PEs, it is possible to configure RPS-DF, not RPS. When configured, the EVPN endpoint will drop frames with MAC SA equal to the black-hole MAC address.

Black-hole MAC addresses can be used as an alternative to MAC filters, which simplifies the deployment of proxy-ARP/ND with anti-spoof MAC addresses. ARP/ND spoofing is a technique whereby an attacker sends fake ARP/ND messages to a broadcast domain. Generally, the aim is to get the routers in the broadcast domain to associate the attacker's MAC address with the IP address of another host, causing any traffic destined to that IP address to be sent to the attacker instead. To prevent this from happening, a proxy-ARP/ND with duplicate IP detection monitors the number of times the MAC changes for an offending

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

IP address. When a certain number of MAC moves are detected in a defined period, the system flags the proxy-ARP entry as duplicate for a defined hold time and an alarm is sent to log 99.

Chapter EVPN for MPLS Tunnels describes the proxy-ARP/ND configuration with the option to define an anti-spoof MAC (AS-MAC) address for EVPN-MPLS networks using MAC filters, including some recommended settings. The AS-MAC address will be advertised with the duplicate IP address in gratuitous ARP (GARP) and ARP replies to all CEs in the EVPN (in the case of proxy-ND, unsolicited Neighbor Advertisement messages are sent instead of GARP messages).

ARP/ND broadcast traffic is a security issue for Internet eXchange Providers (IXPs) and service providers with large Layer 2 domains. In such networks, administrators try to avoid ARP/ND flooding. Figure 63: Proxy-ARP/ND and ARP spoofing shows the proxy-ARP/ND feature where local ARP/ND requests are responded by the system on behalf of the IP interface owners.

Figure 63: Proxy-ARP/ND and ARP spoofing

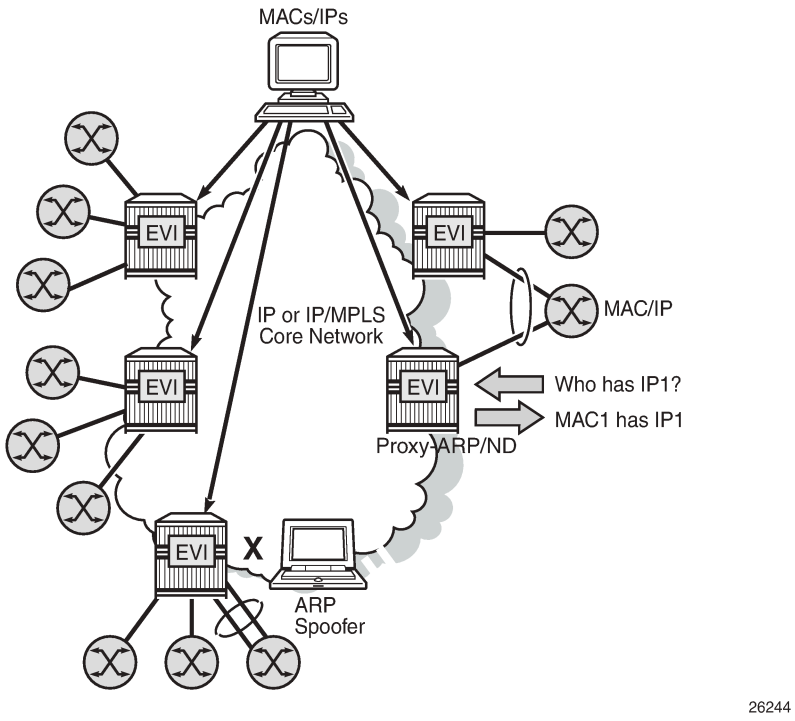

EVPN can suppress ARP/ND flooding within an EVPN service if all the attached hosts advertise their presence. Therefore, EVPN is preferred in IXPs to mitigate and even eliminate the ARP/ND flooding issue. The proxy-ARP/ND agent responds to local ARP/ND requests using a proxy-ARP/ND table per service. This table is populated by EVPN entries (MAC-IP pairs), static entries configured in the service, and dynamic entries snooped from ARP/GARP/ND messages sent by the ISP routers. The static entries and snooped dynamic entries are also advertised in EVPN-MAC routes.

As well as the proxy-ARP/ND, SR OS supports an anti-spoofing mechanism that can detect and block an ARP spoofing attack or a misconfigured duplicated IP address. When using MAC filters, the same anti-spoof-mac option must be configured in all the PEs and this filter may be configured on all the PE SAPs/SDP-bindings to discard all the frames with MAC DA equal to the anti-spoof MAC address. This requires a lot of configuration and is prone to configuration errors.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Conditional static black-hole MAC addresses can be configured for the anti-spoof MAC address so that frames with MAC DA equal to the anti-spoof MAC address can be discarded based on a MAC address lookup in the FDB, as opposed to a MAC filter entry. Less configuration is required and this simplifies the deployment of proxy-ARP/ND with AS-MAC. The configuration example in this chapter includes proxy-ARP, but the behavior is similar for proxy-ND.

Configuration

Figure 64: Example topology shows the example topology. Traffic will be sent between the CEs and may be dropped in the PEs if the MAC DA or MAC SA matches a black-hole MAC address. IP address 172.16.0.10/24 is duplicate (CE-10 and CE-11).

Figure 64: Example topology

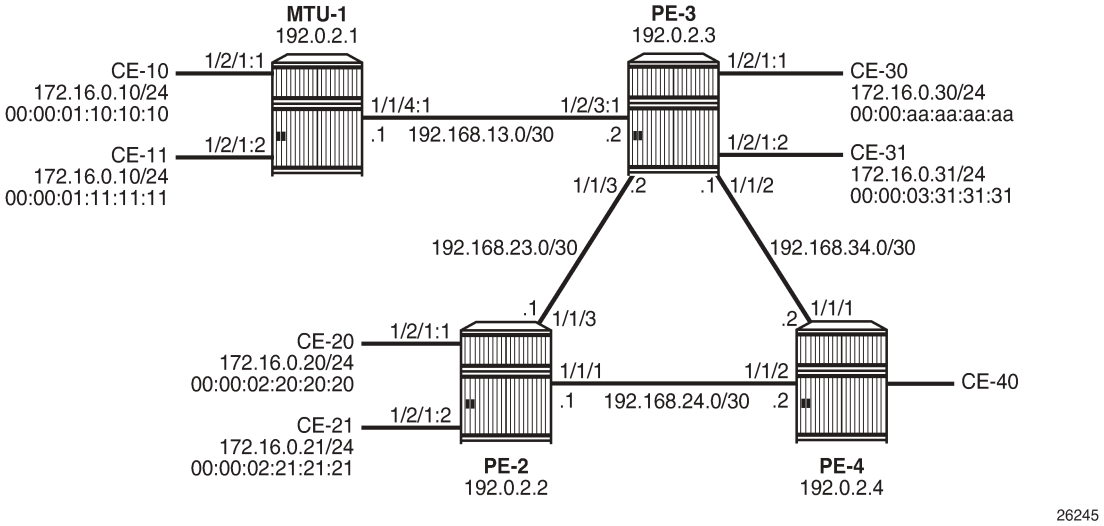

The initial configuration on the nodes includes:

- Cards, MDAs, ports
- Router interfaces
- IS-IS between PEs
- LDP between PEs

BGP is configured between the PEs for address family EVPN with PE-2 as route reflector (RR). Instead of an RR, a full mesh can also be configured between the PEs. The BGP configuration on PE-2 is as follows:

```
# on RR PE-2:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            rapid-withdrawal true
            split-horizon true
            rapid-update {
                evpn true
            }
            group "internal" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                peer-as 64500
                family {
                    evpn true
                }
                cluster {
                    cluster-id 1.1.1.1
                }
            }
            neighbor "192.0.2.3" {
                group "internal"
            }
            neighbor "192.0.2.4" {
                group "internal"
            }
        }
```

VPLS 1 is configured on all PEs and on MTU-1 (MTU-1's VPLS 1 is connected to PE-3 by a SAP). The VPLS configuration on the PEs includes EVPN-MPLS, as follows:

```
# on PE-3:
configure {
    service {
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
            }
            bgp-evpn {
                evi 1
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            sap 1/2/1:1 {
            }
            sap 1/2/3:1 {
            }
        }
```

Conditional static black-hole MAC

Conditional static black-hole MAC address is an extension to the conditional static MAC address, but with the blackhole keyword. It is a scalable way to filter MAC DA or SA in the data plane, regardless of how the frame is arriving at the system (SAP/SDP-bindings or EVPN termination endpoints).

When the static black-hole MAC is added to the FDB, all Ethernet frames with MAC DA equal to the black-hole MAC are dropped. Filtering based on the MAC SA is explained in the next section: Conditional static black-hole MAC in combination with restrict protected source.

Figure 65: Conditional static black-hole MAC shows the example setup with conditional static black-hole MAC 00:00:aa:aa:aa:aa.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 65: Conditional static black-hole MAC

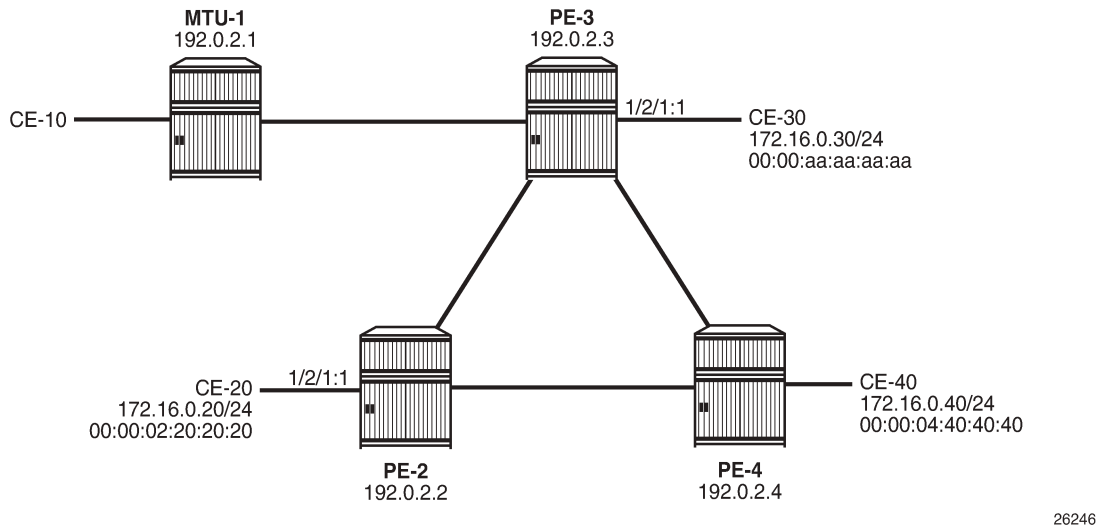

When no conditional static black-hole MAC is configured, CE-30 can receive and send traffic from and to the other CEs; for instance, from and toward CE-20, as follows:

```
[/]
A:admin@PE-2# ping 172.16.0.30 router-instance "VPRN 10"
PING 172.16.0.30 56 data bytes
64 bytes from 172.16.0.30: icmp_seq=1 ttl=64 time=7.31ms.
64 bytes from 172.16.0.30: icmp_seq=2 ttl=64 time=3.33ms.
---snip---
```

```
[/]
A:admin@PE-3# ping 172.16.0.20 router-instance "VPRN 10"
PING 172.16.0.20 56 data bytes
64 bytes from 172.16.0.20: icmp_seq=1 ttl=64 time=3.45ms.
64 bytes from 172.16.0.20: icmp_seq=2 ttl=64 time=3.13ms.
---snip---
```

In this example, CE-20 and CE-30 correspond to VPRN 10 configured on PE-2 and PE-3 (using a hairpin to loop the traffic back to the PE).

Conditional static black-hole MAC 00:00:aa:aa:aa:aa (which corresponds to the MAC address of CE-30) is configured in VPLS 1 on PE-3 as follows:

```
# on PE-3:
configure {
    service {
        vpls "VPLS 1" {
            fdb {
                static-mac {
                    mac 00:00:aa:aa:aa:aa {
                        blackhole
                    }
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The black-hole MAC is added as a conditional static (CStatic) MAC that is protected (P), as follows:

```
[/]
A:admin@PE-3# show service id 1 fdb mac 00:00:aa:aa:aa:aa

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
1          00:00:aa:aa:aa:aa black-hole              CStatic: 03/26/21 11:49:43
                                                     P
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

The source identifier is black-hole and it is applicable to frames that enter the VPLS on this node, regardless of how they enter the VPLS (SAP, SDP-binding, or EVPN endpoint).

The conditional static black-hole MAC is advertised to the BGP peers in a BGP-EVPN MAC route with the sticky/static bit set, as follows:

```
# on PE-3:
9 2021/03/26 11:49:42.675 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 89
    Flag: 0x90 Type: 14 Len: 44 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.3
        Type: EVPN-MAC Len: 33 RD: 192.0.2.3:1 ESI: ESI-0, tag: 0, mac len: 48
                       mac: 00:00:aa:aa:aa:aa, IP len: 0, IP: NULL, label1: 8388544
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:1
        bgp-tunnel-encap:MPLS
        mac-mobility:Seq:0/Static
"
```

The MAC route is added to the FDB on the other PEs as a static (S) and protected (P) MAC; for example, on PE-2, as follows:

```
[/]
A:admin@PE-2# show service id 1 fdb mac 00:00:aa:aa:aa:aa

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
1          00:00:aa:aa:aa:aa mpls:                   EvpnS:P  03/26/21 11:49:43
                             192.0.2.3:524284
           ldp:65537
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

When CE-20 sends an ICMP request to CE-30, the MAC DA 00:00:aa:aa:aa:aa is black-holed on PE-3, and no ICMP request succeeds, as follows:

```
[/]
A:admin@PE-2# ping 172.16.0.30 router-instance "VPRN 10"
PING 172.16.0.30 56 data bytes
Request timed out. icmp_seq=1.
Request timed out. icmp_seq=2.
Request timed out. icmp_seq=3.
Request timed out. icmp_seq=4.
Request timed out. icmp_seq=5.

---- 172.16.0.30 PING Statistics ----
5 packets transmitted, 0 packets received, 100% packet loss
```

The port statistics show that the traffic was sent from PE-2 to PE-3, where it entered on port 1/1/3, then got discarded. To verify this, the port statistics are cleared on PE-2 and PE-3, then 1000 ICMP packets are sent from CE-20, as follows:

```
[/]
A:admin@PE-2# ping 172.16.0.30 router-instance "VPRN 10" count 1000 interval 0.1 output-format
 summary
---snip---
1000 packets transmitted, 0 packets received, 100% packet loss
```

The 1000 packets are received at SAP 1/2/1:1 on PE-2, as follows:

```
[/]
A:admin@PE-2# show port 1/2/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/2/1                                  1000                          106000
                                          0                               0
===============================================================================
```

These packets are forwarded to port 1/1/3 toward PE-3, as follows:

```
[/]
A:admin@PE-2# show port 1/1/3 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/3                                    13                            1306
                                       1013                          125254
===============================================================================
```

On the interfaces between the PEs, other packets are sent besides the ICMP requests, such as IS-IS messages; therefore, the number of packets is slightly greater than 1000.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

On PE-3, these packets are received on port 1/1/3, as follows:

```
[/]
A:admin@PE-3# show port 1/1/3 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/3                                 1024                          126444
                                        24                            2351
===============================================================================
```

The FDB entry for this MAC DA is black-holed and no traffic is received on SAP 1/2/1:1 toward CE-30; therefore, the statistics for port 1/2/1 are empty and nothing is displayed, as follows:

```
[/]
A:admin@PE-3# show port 1/2/1 statistics
```

It is possible to configure the black-hole MAC address on a different PE; for example, on PE-4 instead of PE-3. The conditional static black-hole MAC address configuration in VPLS 1 on PE-3 is removed, as follows:

```
# on PE-3:
configure {
    service {
        vpls "VPLS 1" {
            fdb {
                static-mac {
                    delete mac 00:00:aa:aa:aa:aa {
                    }
                }
            }
```

The conditional static black-hole MAC is configured on PE-4 instead, as follows:

```
# on PE-4:
configure {
    service {
        vpls "VPLS 1" {
            fdb {
                static-mac {
                    mac 00:00:aa:aa:aa:aa {
                        blackhole
                    }
                }
            }
```

PE-4 sends EVPN-MAC updates to its peers. PE-2 learns that all traffic with MAC DA 00:00:aa:aa:aa:aa should be redirected to PE-4, as shown in the FDB on PE-2:

```
[/]
A:admin@PE-2# show service id 1 fdb mac 00:00:aa:aa:aa:aa

===============================================================================
Forwarding Database, Service 1
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
1          00:00:aa:aa:aa:aa mpls:                   EvpnS:P  03/26/21 11:51:59
                             192.0.2.4:524284
           ldp:65538
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

The port statistics are cleared on all PEs and 1000 ICMP packets are sent from CE-20 to CE-30, as follows:

```
[/]
A:admin@PE-2# ping 172.16.0.30 router-instance "VPRN 10" count 1000 interval 0.1 output-format
 summary
---snip---
1000 packets transmitted, 0 packets received, 100% packet loss
```

On PE-2, traffic is not forwarded on the direct link (port 1/1/3) toward PE-3, but redirected to PE-4 (port 1/1/1) instead, as follows:

```
[/]
A:admin@PE-2# show port 1/1/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/1                                    15                           1464
                                       1014                         125360
===============================================================================

[/]
A:admin@PE-2# show port 1/1/3 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/3                                    15                           1433
                                         16                           1589
===============================================================================
```

On PE-4, traffic is received from PE-2 on port 1/1/2, then discarded because the MAC DA equals the static black-hole MAC in the FDB, as follows. No traffic is forwarded to port 1/1/1 toward PE-3, where CE-30 is attached.

```
[/]
A:admin@PE-4# show port 1/1/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
1/1/1                                   22                            2192
                                        22                            2192
===============================================================================

[/]
A:admin@PE-4# show port 1/1/2 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                         Ingress Packets                 Ingress Octets
Id                            Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/2                                 1025                          126476
                                        24                            2351
===============================================================================
```

The configuration is restored with conditional static black-hole MAC in VPLS 1 on PE-3, not on PE-4, as follows:

```
# on PE-3:
configure {
    service {
        vpls "VPLS 1" {
            fdb {
                static-mac {
                    mac 00:00:aa:aa:aa:aa {
                        blackhole
                    }
                }
            }
```

```
# on PE-4:
configure {
    service {
        vpls "VPLS 1" {
            fdb {
                static-mac {
                    delete mac 00:00:aa:aa:aa:aa {
                    }
                }
            }
```

Conditional static black-hole MAC in combination with restrict protected source

For Ethernet frames with MAC SA equal to the static black-hole MAC, the treatment is the same as for protected MACs (see chapter Auto-Learn MAC Protect in EVPN), but for conditional static black-hole MACs, ALMP need not be enabled on the SAP or SDP-binding:

- When a frame is received with MAC SA equal to the black-hole MAC, it is dropped, because RPS-DF is enabled on the SAP or SDP-binding, by default. RPS-DF need not be enabled explicitly. An error message is raised when the following command is entered:

```
[ex:/configure service vpls "VPLS 1" sap 1/2/1:1 fdb]
A:admin@PE-3# protected-src-mac-violation-action discard

*[ex:/configure service vpls "VPLS 1" sap 1/2/1:1 fdb]
A:admin@PE-3# commit
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
MINOR: SVCMGR #12: configure service vpls "VPLS 1" sap 1/2/1:1 fdb protected-src-mac-
violation-action - Inconsistent Value error - not supported on bgp-evpn services - configure
 service vpls "VPLS 1" bgp-evpn
```

- When RPS is enabled instead of RPS-DF, the SAP or SDP-binding where the frame was received, with MAC SA equal to the black-hole MAC, is brought operationally down. The SAP or SDP-binding can be brought up manually by disabling and re-enabling the SAP or SDP-binding. RPS is enabled on SAP 1/2/1:1 as follows:

```
# on PE-3:
configure {
    service {
        vpls "VPLS 1" {
            sap 1/2/1:1 {
                fdb {
                    protected-src-mac-violation-action sap-oper-down
                }
            }
```

- Optionally, RPS-DF can be enabled on the EVPN-MPLS endpoint or EVPN-VXLAN endpoint. When enabled, the EVPN endpoint will discard frames with MAC SA equal to the black-hole MAC. RPS cannot be configured instead of RPS-DF on EVPN endpoints. It is not an option to bring the EVPN endpoint down when a frame is received with MAC SA equal to the static black-hole MAC. The commands to enable RPS-DF on the EVPN-MPLS endpoints and EVPN-VXLAN endpoints are as follows:

```
[ex:configure service vpls "VPLS 1" bgp-evpn mpls 1 fdb]
A:admin@PE-3# protected-src-mac-violation-action ?

 protected-src-mac-violation-action <keyword>
 <keyword> - discard

    Action when a relearn request for a protected MAC is received
```

```
*[ex:configure service vpls "VPLS 1" vxlan instance 1 fdb]
A:admin@PE-3# protected-src-mac-violation-action ?

 protected-src-mac-violation-action <keyword>
 <keyword> - discard

    Action when a relearn request for a protected MAC is received
```

With the default configuration (RPS-DF on SAP/SDP-bindings), the behavior is as follows for conditional static black-hole MAC 00:00:aa:aa:aa:aa configured in VPLS 1 on PE-3. All traffic from CE-30 with MAC SA 00:00:aa:aa:aa:aa is black-holed on SAP 1/2/1:1 on PE-3, because the default behavior on SAP 1/2/1:1 is RPS-DF, and the frame is discarded. The packets are received on port 1/2/1 (SAP 1/2/1:1) and dropped. No packets are forwarded to port 1/1/3 toward PE-2 or any other port.

```
[/]
A:admin@PE-3# ping 172.16.0.20 router-instance "VPRN 10" count 1000 interval 0.1 output-format
 summary
---snip---
1000 packets transmitted, 0 packets received, 100% packet loss
```

```
[/]
A:admin@PE-3# show port 1/1/2 statistics      # port toward PE-4
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/2                                    17                           1732
                                         17                           1732
===============================================================================
```

```
[/]
A:admin@PE-3# show port 1/1/3 statistics      # port toward PE-2

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/3                                    20                           1995
                                         18                           1787
===============================================================================
```

```
[/]
A:admin@PE-3# show port 1/2/1 statistics      # SAP 1/2/1:1 toward CE-30

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/2/1                                  1000                         106000
                                          0                              0
===============================================================================
```

If the static MAC is configured in VPLS 1 on PE-4 and not on PE-3, PE-3 will still discard the packets with MAC SA 00:00:aa:aa:aa:aa arriving on SAP 1/2/1:1, because it learned from the EVPN-MAC updates that MAC 00:00:aa:aa:aa:aa is a protected MAC on PE-4. Therefore, traffic with this MAC SA is not expected and not allowed on PE-3, as follows:

```
# on PE-4:
configure {
    service {
        vpls "VPLS 1" {
            fdb {
                static-mac {
                    mac 00:00:aa:aa:aa:aa {
                        blackhole
                    }
                }
            }
```

```
# on PE-3:
configure {
    service {
        vpls "VPLS 1" {
            fdb {
                static-mac {
                    delete mac 00:00:aa:aa:aa:aa {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    }
                }
            }
```

```
[/]
A:admin@PE-3# show service id 1 fdb mac 00:00:aa:aa:aa:aa

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
1          00:00:aa:aa:aa:aa mpls:                   EvpnS:P  03/26/21 11:55:22
                             192.0.2.4:524284
           ldp:65538
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

```
[/]
A:admin@PE-3# ping 172.16.0.20 router-instance "VPRN 10" count 1000 interval 0.1 output-format
 summary
---snip---
1000 packets transmitted, 0 packets received, 100% packet loss
```

```
[/]
A:admin@PE-3# show port 1/1/2 statistics      # port toward PE-4

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/2                                    22                           2192
                                         22                           2192
===============================================================================
```

```
[/]
A:admin@PE-3# show port 1/1/3 statistics      # port toward PE-2

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/3                                    25                           2476
                                         24                           2351
===============================================================================
```

```
[/]
A:admin@PE-3# show port 1/2/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
1/2/1                                1000                          106000
                                        0                               0
===============================================================================
```

The configuration is restored as follows:

```
# on PE-3:
configure {
    service {
        vpls "VPLS 1" {
            fdb {
                static-mac {
                    mac 00:00:aa:aa:aa:aa {
                        blackhole
                    }
                }
            }
```

```
# on PE-4:
configure {
    service {
        vpls "VPLS 1" {
            fdb {
                static-mac {
                    delete mac 00:00:aa:aa:aa:aa {
                    }
                }
            }
```

Optionally, RPS-DF can be configured on the EVPN-MPLS endpoints on the PEs, as follows:

```
configure {
    service {
        vpls "VPLS 1" {
            bgp-evpn {
                mpls 1 {
                    fdb {
                        protected-src-mac-violation-action discard
                    }
                }
            }
```

When RPS-DF is configured on the EVPN-MPLS endpoints, frames with MAC SA 00:00:aa:aa:aa:aa can be discarded by the EVPN endpoints between the PEs. However, in this example this is not required, because any frame with MAC SA 00:00:aa:aa:aa:aa will be dropped by the local SAP before it can be forwarded to an EVPN endpoint.

It is possible to configure RPS with **sap-oper-down** on SAP 1/2/1:1 on PE-3, as follows:

```
# on PE-3:
configure {
    service {
        vpls "VPLS 1" {
            sap 1/2/1:1 {
                fdb {
                    protected-src-mac-violation-action sap-oper-down
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

When CE-30 sends traffic with MAC SA equal to a protected MAC address (black-hole or not), the entire SAP 1/2/1:1 will be brought operationally down, as follows:

```
[/]
A:admin@PE-3# ping 172.16.0.20 router-instance "VPRN 10"
PING 172.16.0.20 56 data bytes
Request timed out. icmp_seq=1.
Request timed out. icmp_seq=2.
---snip---
---- 172.16.0.20 PING Statistics ----
5 packets transmitted, 0 packets received, 100% packet loss
```

```
[/]
A:admin@PE-3# show service id 1 sap

===============================================================================
SAP(Summary), Service 1
===============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.  Egr.   Adm  Opr
                                           QoS   Fltr    QoS   Fltr
-------------------------------------------------------------------------------
1/2/1:1                         1          1     none    1     none   Up   Down
1/2/3:1                         1          1     none    1     none   Up   Up
-------------------------------------------------------------------------------
Number of SAPs : 2
-------------------------------------------------------------------------------
===============================================================================
```

The following information for SAP 1/2/1:1 in VPLS 1 shows that this SAP is operationally down because a protected source MAC address was received on this SAP (Flags: RxProtSrcMac), as follows:

```
[/]
A:admin@PE-3# show service id 1 sap 1/2/1:1

===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 1
SAP                : 1/2/1:1                 Encap             : q-tag
Description        : (Not Specified)
Admin State        : Up                      Oper State        : Down
Flags              : RxProtSrcMac
Multi Svc Site     : None
Last Status Change : 03/26/2021 11:54:35
Last Mgmt Change   : 03/26/2021 11:53:50
===============================================================================
```

Log 99 shows that a protected MAC was received on SAP 1/2/1:1 and the SAP went operationally down with flag RxProtSrcMac, as follows:

```
90 2021/03/26 11:54:35.164 CET MINOR: SVCMGR #2208 Base
"Protected MAC 00:00:aa:aa:aa:aa received on SAP 1/2/1:1 in service 1. The SAP will be
 disabled."

91 2021/03/26 11:54:35.164 CET MINOR: SVCMGR #2203 Base
"Status of SAP 1/2/1:1 in service 1 (customer 1) changed to admin=up oper=down flags=RxProtSrc
Mac "
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The SAP can only be brought up manually by disabling and re-enabling the SAP, as follows:

```
# on PE-3:
configure {
    service {
        vpls "VPLS 1" {
            sap 1/2/1:1 {
                admin-state disable
                commit
                admin-state enable
                commit
```

```
[/]
A:admin@PE-3# show service id 1 sap

===============================================================================
SAP(Summary), Service 1
===============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.  Egr.   Adm  Opr
                                           QoS   Fltr    QoS   Fltr
-------------------------------------------------------------------------------
1/2/1:1                         1          1     none    1     none   Up   Up
1/2/3:1                         1          1     none    1     none   Up   Up
-------------------------------------------------------------------------------
Number of SAPs : 2
-------------------------------------------------------------------------------
===============================================================================
```

The default behavior of SAP 1/2/1:1 is RPS-DF, which is configured by removing the RPS configuration, as follows:

```
# on PE-3:
configure {
    service {
        vpls "VPLS 1" {
            sap 1/2/1:1 {
                fdb {
                    delete protected-src-mac-violation-action
                }
            }
```

The conditional static black-hole MAC configuration is removed as follows:

```
# on PE-3:
configure {
    service {
        vpls "VPLS 1" {
            fdb {
                static-mac {
                    delete mac 00:00:aa:aa:aa:aa {
                    }
                }
            }
```

Black-hole MAC in services with proxy-ARP/ND

In this example, only proxy-ARP is shown, not proxy-ND. However, the configuration and procedures for proxy-ND would be equivalent.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

First, the implementation of proxy-ARP and AS-MAC is described without static black-hole MAC addresses. MAC filters will be required to drop or redirect traffic, but these are not shown in the example. Configuring MAC filters and applying them on SAP/SDP-bindings is labor-intensive and can be error-prone. Afterward, the implementation with AS-MAC as static black-hole is described.

Services with proxy-ARP and AS-MAC - no static black-hole MAC

IP duplication works when the IP address moves between:

- Dynamic (learned on SAP) and EVPN
- EVPN and dynamic
- Dynamic and dynamic

The following example shows IP address moves from dynamic to dynamic between SAP 1/2/1:1 (to CE-10) and SAP 1/2/1:2 (to CE-11) in VPLS 1 on MTU-1. However, the duplicate IP address could have been in PE-3 and MTU-1 instead (EVPN or dynamic) and still the IP address would have been detected as duplicate.

Figure 66: VPLS 1 with proxy-ARP and AS-MAC shows the example setup with duplicate IP address 172.16.0.10/24 for CE-10 and CE-11. VPLS 1 is configured with proxy-ARP with duplicate IP detection in PE-2 and PE-3 (and possibly also in other PEs). MAC address 00:00:bb:bb:bb:bb is configured as AS-MAC, which will be used when a duplicate IP address has been detected.

Figure 66: VPLS 1 with proxy-ARP and AS-MAC

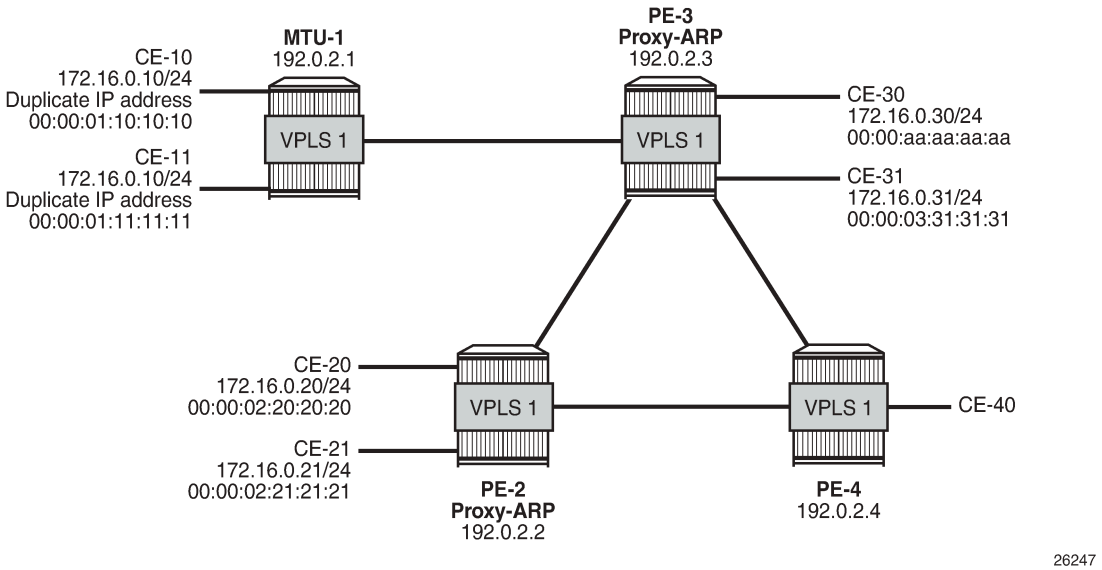

For IP duplication detection, the following parameters can be customized so that the system can react to particular conditions in the network. The syntax is as follows:

```
*[ex:/configure service vpls "VPLS 1" proxy-arp]
A:admin@PE-3# duplicate-detect ?

 duplicate-detect

 anti-spoof-mac      - MAC address to replace the proxy-ARP/ND offending entry's MAC
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
 hold-down-time      - Hold down time for a duplicate entry
 num-moves           - Number of moves required to declare a duplicate entry
 static-blackhole    - Consider anti-spoof MAC as black-hole static MAC in FDB
 window              - Time to monitor the MAC address in the anti-spoofing mechanism
```

In VPLS 1 on PE-2 and PE-3, a proxy-ARP with duplicate IP detection is configured, including an optional anti-spoof MAC (AS-MAC) 00:00:bb:bb:bb:bb for offending IP addresses, as follows:

```
# on PE-2, PE-3:
configure {
    service {
        vpls "VPLS 1" {
            proxy-arp {
                admin-state enable
                dynamic-populate true
                duplicate-detect {
                    anti-spoof-mac 00:00:bb:bb:bb:bb
                    window 3
                    num-moves 3
                    hold-down-time max
                }
                static-arp {
                    ip-address 172.16.0.20 {
                        mac 00:00:02:20:20:20
                    }
                }
            }
```

The proxy-ARP table contains one static entry (for IP 172.16.0.20). In this case, dynamic ARP populate is enabled. Therefore, the proxy-ARP table will be updated with ARP entries for IP 172.16.0.10 and MAC 00:00:01:10:10:10 or MAC 00:00:01:11:11:11 for frames originating from CE-10 or CE-11.

When a duplicate IP is detected for IP 172.16.0.10 (after three changes of MAC for IP 172.16.0.10 in a period of three minutes), the corresponding ARP entry contains the duplicate IP address 172.16.0.10 and the AS-MAC 00:00:bb:bb:bb:bb and its type is duplicate (dup). Therefore, this ARP entry is always active until it is removed. Until now, this configuration does not include a static black-hole MAC, and this option is by default disabled. This configuration for duplicate IP detection can be used in combination with MAC filters. The configuration with static black-hole MAC is shown in the section Services with proxy-ARP and AS-MAC configured as static black-hole MAC.

The configured AS-MAC will be advertised in an EVPN-MAC route with the sticky/static bit set and without any IP address (because there is no IP duplication detected yet), as follows:

```
# on PE-3:
25 2021/03/26 11:59:16.417 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 124
    Flag: 0x90 Type: 14 Len: 79 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.3
        Type: EVPN-MAC Len: 33 RD: 192.0.2.3:1 ESI: ESI-0, tag: 0, mac len: 48
                       mac: 02:17:ff:00:03:3a, IP len: 0, IP: NULL, label1: 8388544
        Type: EVPN-MAC Len: 33 RD: 192.0.2.3:1 ESI: ESI-0, tag: 0, mac len: 48
                       mac: 00:00:bb:bb:bb:bb, IP len: 0, IP: NULL, label1: 8388544
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        bgp-tunnel-encap:MPLS
        mac-mobility:Seq:0/Static
"
```

Without the option static black-hole, the configured AS-MAC is not added to the local FDB, but this MAC address is treated as a local MAC. The FDB on PE-3 does not contain AS-MAC 00:00:bb:bb:bb:bb, as follows:

```
[/]
A:admin@PE-3# show service id 1 fdb mac 00:00:bb:bb:bb:bb

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

Debugging is enabled for proxy-ARP for IP address 172.16.0.10 in VPLS 1 on PE-3, as follows (in classic CLI):

```
A:PE-3# debug service id 1 proxy-arp ip 172.16.0.10
```

When traffic is sent from CE-11 to CE-21, a dynamic ARP entry for IP address 172.16.0.10 and MAC 00:00:01:11:11:11 is added to the proxy-ARP table for VPLS 1 in PE-3, and an EVPN-MAC update is sent to the peer PEs, as follows:

```
# on PE-3:
36 2021/03/26 12:06:35.179 CET MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.10 mac: 00:00:01:11:11:11 evpn advertise"

37 2021/03/26 12:06:35.179 CET MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.10 type: Dyn mac: 00:00:01:11:11:11 Added"

38 2021/03/26 12:06:35.179 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 85
    Flag: 0x90 Type: 14 Len: 48 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.3
        Type: EVPN-MAC Len: 37 RD: 192.0.2.3:1 ESI: ESI-0, tag: 0, mac len: 48
                mac: 00:00:01:11:11:11, IP len: 4, IP: 172.16.0.10, label1: 8388544
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:1
        bgp-tunnel-encap:MPLS
"
```

There is no duplicate IP detected yet.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

CE-10 and CE-11 have the same IP address for different MAC addresses. When CE-10 sends traffic to CE-20, the ARP entry for IP 172.16.0.10 changes MAC from 00:00:01:11:11:11 to 00:00:01:10:10:10, and an EVPN-MAC withdraw message is sent, as follows:

```
# on PE-3:
39 2021/03/26 12:06:35.489 CET MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.10 mac: 00:00:01:11:11:11 evpn withdraw"

40 2021/03/26 12:06:35.489 CET MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.10 Mac Change: 00:00:01:11:11:11->00:00:01:10:10:10 "

41 2021/03/26 12:06:35.489 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 46
    Flag: 0x90 Type: 15 Len: 42 Multiprotocol Unreachable NLRI:
        Address Family EVPN
        Type: EVPN-MAC Len: 37 RD: 192.0.2.3:1 ESI: ESI-0, tag: 0, mac len: 48
                       mac: 00:00:01:11:11:11, IP len: 4, IP: 172.16.0.10, label1: 0
"
```

When the MAC changes, the system sends an ARP request for confirmation of the old MAC 00:00:01:11:11:11 for IP 172.16.0.10, as follows:

```
# on PE-3:
42 2021/03/26 12:06:35.542 CET MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.10 mac: 00:00:01:11:11:11 confirm"
```

When MAC 00:00:01:11:11:11 is confirmed, the MAC in the ARP entry is changed once again to 00:00:01:11:11:11 and another ARP request is sent asking to confirm MAC 00:00:01:10:10:10 for IP 172.16.0.10, as follows:

```
# on PE-3:
43 2021/03/26 12:06:35.798 CET MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.10 Mac Change: 00:00:01:10:10:10->00:00:01:11:11:11 "

44 2021/03/26 12:06:35.842 CET MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.10 mac: 00:00:01:10:10:10 confirm"
```

When CE-10 confirms MAC 00:00:01:10:10:10 for IP 172.16.0.10, IP duplication is detected for IP address 172.16.0.10 (after three MAC moves in a detection period of three minutes), and the following message is raised in log 99 after a duplicate proxy-ARP entry was detected for IP 172.16.0.10:

```
# log "99" on PE-3:
107 2021/03/26 12:06:36.108 CET MINOR: SVCMGR #2346 Base
"A duplicate proxy ARP entry was detected with new MAC 00:00:01:10:10:10 for entry IP
 172.16.0.10 MAC 00:00:01:11:11:11 in service 1"
```

The following proxy-ARP debug messages show that the ARP entry for IP 172.16.0.10 in the proxy-ARP table changed MAC to the AS-MAC 00:00:bb:bb:bb:bb, and the type from dynamic to duplicate:

```
# on PE-3:
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
45 2021/03/26 12:06:36.108 CET MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.10 mac: 00:00:bb:bb:bb:bb evpn advertise"

46 2021/03/26 12:06:36.108 CET MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.10 Mac Change: 00:00:01:11:11:11->00:00:bb:bb:bb:bb Type Change:
 Dyn->Dup "

47 2021/03/26 12:06:36.108 CET MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.10 type: Dup    Dup Detected"
```

If a duplicate IP is detected, AS-MAC 00:00:bb:bb:bb:bb is advertised with duplicate IP address 172.16.0.10 in an EVPN-MAC update to the BGP peers with the sticky/static bit set, as follows:

```
48 2021/03/26 12:06:36.108 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 93
    Flag: 0x90 Type: 14 Len: 48 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.3
        Type: EVPN-MAC Len: 37 RD: 192.0.2.3:1 ESI: ESI-0, tag: 0, mac len: 48
                   mac: 00:00:bb:bb:bb:bb, IP len: 4, IP: 172.16.0.10, label1: 8388544
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:1
        bgp-tunnel-encap:MPLS
        mac-mobility:Seq:0/Static
"
```

The difference with the first EVPN-MAC update for AS-MAC is the IP address. Immediately after the AS-MAC was configured, it was also advertised to the BGP-EVPN peers, but without any IP address.

The proxy-ARP entry is shown with type duplicate (dup) and active status in the proxy-ARP table for VPLS 1 on PE-3, as follows:

```
[/]
A:admin@PE-3# show service id 1 proxy-arp detail
-------------------------------------------------------------------------------
Proxy Arp
-------------------------------------------------------------------------------
Admin State       : enabled
Dyn Populate      : enabled
Age Time          : disabled            Send Refresh      : disabled
Table Size        : 250                 Total             : 2
Static Count      : 1                   EVPN Count        : 0
Dynamic Count     : 0                   Duplicate Count   : 1

Dup Detect
-------------------------------------------------------------------------------
Detect Window     : 3 mins              Num Moves         : 3
Hold down         : max
Anti Spoof MAC    : 00:00:bb:bb:bb:bb

EVPN
-------------------------------------------------------------------------------
Garp Flood        : enabled             Req Flood         : enabled
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Static Black Hole : disabled
EVPN Route Tag    : 0
-------------------------------------------------------------------------------

===============================================================================
VPLS Proxy Arp Entries
===============================================================================
IP Address          Mac Address         Type      Status    Last Update
-------------------------------------------------------------------------------
172.16.0.10         00:00:bb:bb:bb:bb   dup       active    03/26/2021 12:06:36
172.16.0.20         00:00:02:20:20:20   stat      inActv    03/26/2021 11:59:16
-------------------------------------------------------------------------------
Number of entries : 2
===============================================================================
```

A duplicate entry is always active, regardless of the AS-MAC. When the entry with the duplicate IP address and the AS-MAC address are installed in the proxy-ARP table as active, every ARP request for the duplicate IP address will be replied by the system. The entry in the proxy-ARP table is treated as active, even if the AS-MAC address is not in the FDB (AS-MAC addresses do not consume FDB space). The AS-MAC address, along with the duplicate IP address, is advertised in EVPN with the sticky/static bit set, as shown earlier. GARP messages with AS-MAC/IP information are flooded locally to make the CEs update their ARP caches to use the AS-MAC address for traffic to the duplicate IP 172.16.0.10, as follows.

```
# on PE-3:
49 2021/03/26 12:06:36.142 CET MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.10 type: Dup mac: 00:00:bb:bb:bb:bb Gratuitous Update"
```

Note:
The AS-MAC address will always be "unique" in the system. When the AS-MAC is configured, the system will flush any entry with the same MAC address learned through EVPN or dynamic sources. Conditional static MAC addresses or OAM MAC addresses with the same value as the AS-MAC address are only allowed when they are configured as black-hole, which is not the case yet.

When the duplicate proxy-ARP entry is cleared from the list (hold-down timer expires, or clear command, or replacement of the duplicate entry for a static entry), an ARP request asking who has IP 172.16.0.10 is flooded by the proxy-ARP agent. This ARP refresh triggers an ARP reply from the IP owner, which will be learned in the proxy-ARP table and advertised in EVPN. The system will also send a GARP to local SAP/SDP-bindings. This will correct all host ARP caches in the network. In this example, the duplicate proxy-ARP entry is manually cleared, as follows:

```
[/]
A:admin@PE-3# clear service id 1 proxy-arp duplicate
```

Log 99 shows that the clear function has been run and the duplicate proxy-ARP entry 172.16.0.10 is cleared. The system forces a refresh and, if the condition with the duplicate IP address remains, this is detected almost immediately and a message is logged that a duplicate proxy-ARP entry was detected, as follows:

```
# on PE-3:
108 2021/03/26 12:07:54.958 CET INDETERMINATE: LOGGER #2010 Base Clear SVCMGR
"Clear function clearSvcIdProxyArpDups has been run with parameters: svc-id="1" ip-address="".
 The completion result is: success.  Additional error text, if any, is: "

109 2021/03/26 12:07:54.958 CET MINOR: SVCMGR #2347 Base
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
"A duplicate proxy ARP entry 172.16.0.10 is cleared in service 1"

110 2021/03/26 12:07:55.146 CET MINOR: SVCMGR #2346 Base
"A duplicate proxy ARP entry was detected with new MAC 00:00:01:11:11:11 for entry IP
 172.16.0.10 MAC 00:00:01:10:10:10 in service 1"
```

The following debug messages for proxy-ARP on PE-3 show the process in more detail. Initially, an EVPN-MAC route withdraw message is sent and the proxy-ARP entry is deleted.

```
# on PE-3:
50 2021/03/26 12:07:54.958 CET MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.10 mac: 00:00:bb:bb:bb:bb evpn withdraw"

51 2021/03/26 12:07:54.958 CET MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.10 type: Dup mac: 00:00:bb:bb:bb:bb Deleted"
```

The following BGP-EVPN MAC update is sent by PE-3 to indicate that the AS-MAC is withdrawn for IP 172.16.0.10 (multiprotocol unreachable NLRI):

```
# on PE-3:
53 2021/03/26 12:07:54.958 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 46
    Flag: 0x90 Type: 15 Len: 42 Multiprotocol Unreachable NLRI:
        Address Family EVPN
        Type: EVPN-MAC Len: 37 RD: 192.0.2.3:1 ESI: ESI-0, tag: 0, mac len: 48
                       mac: 00:00:bb:bb:bb:bb, IP len: 4, IP: 172.16.0.10, label1: 0
"
```

Removing the active duplicate entry from the proxy-ARP table triggers an ARP flooding request asking who has IP 172.16.0.10 in VPLS 1, as follows:

```
# on PE-3:
52 2021/03/26 12:07:54.958 CET MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.10 flood request"
```

The result of the ARP flooding request is that the IP owners reply with their MAC, at the local or a remote PE. In this case, the reply from CE-10 is received first (IP 172.16.0.10 - MAC 00:00:01:10:10:10), a dynamic proxy-ARP entry is added, and the MAC/IP route is advertised, as follows:

```
# on PE-3:
54 2021/03/26 12:07:54.961 CET MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.10 mac: 00:00:01:10:10:10 evpn advertise"

55 2021/03/26 12:07:54.961 CET MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.10 type: Dyn mac: 00:00:01:10:10:10 Added"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

When CE-11 answers with its MAC 00:00:01:11:11:11, the MAC/IP route is withdrawn for IP 172.16.0.10, and the MAC address in the proxy-ARP entry for IP 172.16.0.10 is changed from MAC 00:00:01:10:10:10 to MAC 00:00:01:11:11:11, as follows:

```
# on PE-3:
56 2021/03/26 12:07:54.961 CET MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.10 mac: 00:00:01:10:10:10 evpn withdraw"

57 2021/03/26 12:07:54.961 CET MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.10 Mac Change: 00:00:01:10:10:10->00:00:01:11:11:11 "
```

Any change of MAC address in a proxy-ARP entry triggers an ARP request asking for confirmation of the old MAC address for IP 172.16.0.10, in this case for MAC 00:00:01:10:10:10, as follows:

```
# on PE-3:
58 2021/03/26 12:07:55.042 CET MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.10 mac: 00:00:01:10:10:10 confirm"
```

MAC address 00:00:01:10:10:10 is confirmed for IP address 172.16.0.10; therefore, the MAC address is changed in the proxy-ARP entry from 00:00:01:11:11:11 to 00:00:01:10:10:10, and an ARP confirmation is asked for the old MAC address 00:00:01:11:11:11, as follows:

```
# on PE-3:
59 2021/03/26 12:07:55.045 CET MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.10 Mac Change: 00:00:01:11:11:11->00:00:01:10:10:10 "

60 2021/03/26 12:07:55.142 CET MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.10 mac: 00:00:01:11:11:11 confirm"
```

MAC address 00:00:01:11:11:11 is confirmed and, therefore, three MAC moves occurred within three minutes. Duplicate IP 172.16.0.10 is detected and the proxy-ARP entry has the AS-MAC 00:00:bb:bb:bb:bb and type duplicate (Dup), as follows:

```
# on PE-3:
61 2021/03/26 12:07:55.146 CET MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.10 mac: 00:00:bb:bb:bb:bb evpn advertise"

62 2021/03/26 12:07:55.146 CET MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.10 Mac Change: 00:00:01:10:10:10->00:00:bb:bb:bb:bb Type Change:
 Dyn->Dup "

63 2021/03/26 12:07:55.146 CET MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.10 type: Dup    Dup Detected"

64 2021/03/26 12:07:55.146 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 93
    Flag: 0x90 Type: 14 Len: 48 Multiprotocol Reachable NLRI:
        Address Family EVPN
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        NextHop len 4 NextHop 192.0.2.3
        Type: EVPN-MAC Len: 37 RD: 192.0.2.3:1 ESI: ESI-0, tag: 0, mac len: 48
                  mac: 00:00:bb:bb:bb:bb, IP len: 4, IP: 172.16.0.10, label1: 8388544
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:1
        bgp-tunnel-encap:MPLS
        mac-mobility:Seq:0/Static
"
```

A GARP update is sent for IP 172.16.0.10 and AS-MAC 00:00:bb:bb:bb:bb, as follows:

```
# on PE-3:
65 2021/03/26 12:07:55.242 CET MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.10 type: Dup mac: 00:00:bb:bb:bb:bb Gratuitous Update"
```

The AS-MAC address is optionally configured and populates all the host ARP caches when a duplicate IP address is detected. All traffic destined to the suspicious IP address 172.16.0.10 will have the AS-MAC address 00:00:bb:bb:bb:bb as MAC DA. The user can configure MAC filters on all SAP/SDP-bindings where the CEs are connected to drop, log, or redirect traffic destined to the AS-MAC. This will block any interception or man-in-the-middle attack (due to ARP spoofing) in the network.

The AS-MAC address is independently configured on each PE for the same service. When a different AS-MAC address is configured per PE for the same service, the user will need to filter all the AS-MAC addresses in the service at each PE, which increases the complexity of the filters. Nokia recommends using the same AS-MAC address for the same service in all the PES where duplicate detect is active and MAC filters need to be configured. However, this recommendation is suspended when the AS-MAC address is configured as static black-hole MAC address, as described in the following section.

Services with proxy-ARP and AS-MAC configured as static black-hole MAC

With the AS-MAC address configured as static black-hole MAC address, MAC-filters do not need to be configured to discard frames with MAC DA equal to the AS-MAC address. Instead, the user can decide whether to use the same AS-MAC address on all the PEs. This scalability is not limited by the number of filters, but by the number of FDB entries.

The **static-blackhole** parameter is optional and disabled by default. In the example, the static-black-hole option is not configured yet for the AS-MAC address and the behavior is as follows:

- The AS-MAC address is added to the MAC DB as local, but not programmed in the FDB.
- The AS-MAC address is advertised in EVPN (initially without an IP address, and with an IP address as soon as the IP is detected as duplicate).
- The AS-MAC address cannot be overridden by any other MAC address.
- The AS-MAC address value cannot be configured on a static MAC address, because that MAC address is reserved for the proxy-ARP, as follows:

```
# on PE-3:
*[ex:/configure service vpls "VPLS 1" fdb static-mac mac 00:00:bb:bb:bb:bb]
A:admin@PE-3# sap 1/2/3:1 monitor forward-status

*[ex:/configure service vpls "VPLS 1" fdb static-mac mac 00:00:bb:bb:bb:bb]
A:admin@PE-3# commit
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
MINOR: MGMT_CORE #4001: configure service vpls "VPLS 1" proxy-arp duplicate-detect anti-spoof-
mac - antispoof-mac conflicts with static-mac - configure service vpls "VPLS 1" fdb static-mac
 mac 00:00:bb:bb:bb:bb

# on PE-3:
*[ex:/configure service vpls "VPLS 1" fdb static-mac mac 00:00:bb:bb:bb:bb]
A:admin@PE-3# blackhole

*[ex:/configure service vpls "VPLS 1" fdb static-mac mac 00:00:bb:bb:bb:bb]
A:admin@PE-3# commit
MINOR: MGMT_CORE #4001: configure service vpls "VPLS 1" proxy-arp duplicate-detect anti-spoof-
mac - antispoof-mac conflicts with static-mac - configure service vpls "VPLS 1" fdb static-mac
 mac 00:00:bb:bb:bb:bb
```

When the **static-blackhole** option is not configured, the AS-MAC address is considered as a local MAC address and cannot be overridden. The MAC address priority is as follows:

1. Local MAC address (including AS-MAC addresses without static-black-hole, es-bmacs, src-bmacs, OAM, and so on)
2. Conditional static MAC addresses (including AS-MAC addresses with static-black-hole)
3. Auto-Learn Protected MAC addresses
4. EVPN-MAC addresses with sticky/static bit set
5. Data plane learned MAC addresses (regular learning on SAP/SDP-binding)
6. EVPN-MAC addresses without sticky/static bit set

To configure an AS-MAC address with static-black-hole option, a static black-hole MAC address needs to be configured. The following error is raised when no static black-hole MAC has been configured for AS-MAC 00:00:bb:bb:bb:bb:

```
# on PE-3:
*[ex:/configure service vpls "VPLS 1" proxy-arp duplicate-detect]
A:admin@PE-3# static-blackhole true

*[ex:/configure service vpls "VPLS 1" proxy-arp duplicate-detect]
A:admin@PE-3# commit
MINOR: MGMT_CORE #4001: configure service vpls "VPLS 1" proxy-arp duplicate-detect static-
blackhole - blackhole conditional static mac needs to be configured - configure service vpls
 "VPLS 1"
```

The VPLS service is configured with proxy-ARP and AS-MAC as static black-hole on PE-2 and PE-3, as follows:

```
# on PE-2, PE-3:
configure {
    service {
        vpls "VPLS 1" {
            fdb {
                static-mac {
                    mac 00:00:bb:bb:bb:bb {
                        blackhole
                    }
                }
            }
            proxy-arp {
                admin-state enable
                dynamic-populate true
                duplicate-detect {
                    anti-spoof-mac 00:00:bb:bb:bb:bb
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    window 3
                    num-moves 5
                    hold-down-time max
                    static-blackhole true
                }
                static-arp {
                    ip-address 172.16.0.20 {
                        mac 00:00:02:20:20:20
                    }
                }
```

When the AS-MAC address is configured with the static black-hole option, the AS-MAC will be added not only to the MAC DB, but also to the FDB as CStatic, and associated with a black-hole endpoint, as follows:

```
[/]
A:admin@PE-3# show service id 1 fdb mac 00:00:bb:bb:bb:bb

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
1          00:00:bb:bb:bb:bb black-hole              CStatic: 03/26/21 12:09:35
                                                     P
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

Any frame with MAC DA equal to the AS-MAC with static black-hole will be dropped, regardless of the ingress endpoint and without any need for a filter. This mechanism is the only way to filter MAC DAs on EVPN endpoints, because MAC filters cannot be configured on EVPN endpoints.

The AS-MAC with static black-hole will be advertised in EVPN with the sticky/static bit set, as follows:

```
# on PE-3:
87 2021/03/26 12:09:34.970 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 89
    Flag: 0x90 Type: 14 Len: 44 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.3
        Type: EVPN-MAC Len: 33 RD: 192.0.2.3:1 ESI: ESI-0, tag: 0, mac len: 48
                       mac: 00:00:bb:bb:bb:bb, IP len: 0, IP: NULL, label1: 8388544
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:1
        bgp-tunnel-encap:MPLS
        mac-mobility:Seq:0/Static
"
```

When a duplicate IP address is detected, the EVPN-MAC update contains the IP address 172.16.0.10, as follows:

```
91 2021/03/26 12:10:10.674 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Send BGP UPDATE:
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    Withdrawn Length = 0
    Total Path Attr Length = 93
    Flag: 0x90 Type: 14 Len: 48 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.3
        Type: EVPN-MAC Len: 37 RD: 192.0.2.3:1 ESI: ESI-0, tag: 0, mac len: 48
                       mac: 00:00:bb:bb:bb:bb, IP len: 4, IP: 172.16.0.10, label1: 8388544
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:1
        bgp-tunnel-encap:MPLS
        mac-mobility:Seq:0/Static
"
```

The local CEs receive a GARP update with the AS-MAC address. The ARP table of CE-30 and CE-31 have an entry for the duplicate IP address 172.16.0.10 with the AS-MAC address 00:00:bb:bb:bb:bb, as follows:

```
[/]
A:admin@PE-3# show router 10 arp

===============================================================================
ARP Table (Service: 10)
===============================================================================
IP Address      MAC Address       Expiry    Type   Interface
-------------------------------------------------------------------------------
172.16.0.10     00:00:bb:bb:bb:bb 03h43m02s Dyn[I] int-CE-30-PE-3
172.16.0.30     00:00:aa:aa:aa:aa 00h00m00s Oth[I] int-CE-30-PE-3
-------------------------------------------------------------------------------
No. of ARP Entries: 2
===============================================================================
```

```
[/]
A:admin@PE-3# show router 11 arp

===============================================================================
ARP Table (Service: 11)
===============================================================================
IP Address      MAC Address       Expiry    Type   Interface
-------------------------------------------------------------------------------
172.16.0.10     00:00:bb:bb:bb:bb 03h47m43s Dyn[I] int-CE-31-PE-3
172.16.0.31     00:00:03:31:31:31 00h00m00s Oth[I] int-CE-31-PE-3
-------------------------------------------------------------------------------
No. of ARP Entries: 2
===============================================================================
```

CE-30 and CE-31 cannot reach CE-10 or CE-11, because the MAC DA will be the AS-MAC address and all traffic to this MAC DA is black-holed instead of forwarded to SAP 1/2/3:1 toward CE-10 or CE-11. When 1000 ICMP packets are sent by CE-30, they arrive in SAP 1/2/1:1 on PE-3 and are then discarded, as follows:

```
[/]
A:admin@PE-3# ping 172.16.0.10 router-instance "VPRN 10" count 1000 interval 0.1 output-format
 summary
PING 172.16.0.10 56 data bytes
---snip---
---- 172.16.0.10 PING Statistics ----
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
1000 packets transmitted, 0 packets received, 100% packet loss
```

```
[/]
A:admin@PE-3# show port 1/1/2 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/2                                    13                           1274
                                         13                           1274
===============================================================================
```

```
[/]
A:admin@PE-3# show port 1/1/3 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/3                                    16                           1537
                                         16                           1537
===============================================================================
```

```
[/]
A:admin@PE-3# show port 1/2/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/2/1                                  1000                         106000
                                          0                              0
===============================================================================
```

No packets were forwarded to SAP 1/2/3:1 toward MTU-1; therefore, there are no statistics for port 1/2/3.

Conclusion

Static black-hole MAC addresses can be applied in EVPN for security as a scalable alternative to MAC filters. Static black-hole MAC addresses are programmed in the FDB and all frames with MAC DA equal to the static black-hole MAC address are dropped, regardless of how the frame arrived at the system (SAP/SDP-binding or EVPN endpoint). Also, static black-hole MAC addresses are treated like protected MAC addresses and, in combination with RPS(-DF), filtering on MAC SA is performed in the data plane. Black-hole MAC addresses can be an option for an AS-MAC address in services with proxy-ARP/ND enabled, which simplifies the configuration because MAC filters are not required.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Data Center Interconnect Using Dual EVPN-VXLAN Instance VPLS

This chapter provides information about Data Center Interconnect using dual EVPN-VXLAN instance VPLS.

Topics in this chapter include:

Applicability

This chapter was initially written for SR OS Release 16.0.R7, but the MD-CLI in the current edition is based on SR OS Release 21.7.R1. Dual EVPN-VXLAN instances are supported in SR OS Release 16.0.R2, or later.

This chapter describes the redundancy based on an Anycast solution, as supported in SR OS Release 16.0, and later. For I-ES based redundancy scenarios as supported in SR OS Release 19.10, and later, see the EVPN Interconnect Ethernet Segments in Dual EVPN-VXLAN Instance VPLS Services chapter.

Overview

Chapter EVPN-MPLS Interconnect for EVPN-VXLAN VPLS Services describes a Data Center Interconnect (DCI) scenario using VXLAN in the DCs and MPLS in the WAN. This chapter describes a similar scenario, where the core is an IP network that does not use MPLS, and where end-to-end VXLAN is used instead. The DC Gateways (GWs) contain VPLS services with two EVPN-VXLAN instances and two BGP instances: one EVPN-VXLAN instance faces the DC and the other EVPN-VXLAN instance faces the WAN.

Figure 67: Dual EVPN-VXLAN instance VPLS 1 shows the example topology with two DCs. On PE-1 and PE-6, VPLS 1 is configured with one VXLAN instance and one BGP instance. On the DC GWs, VPLS 1 is configured with two VXLAN instances and two BGP instances: one toward the DC and one toward the WAN.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 67: Dual EVPN-VXLAN instance VPLS 1

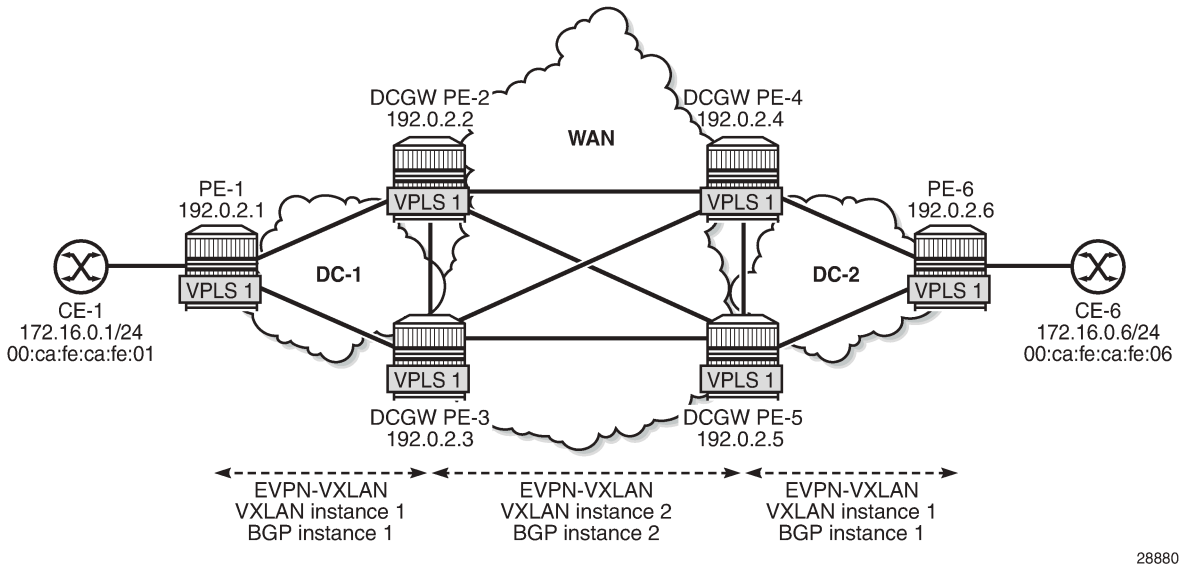

For example, on DC GW PE-2, VPLS 1 is configured with VXLAN instance 1 using BGP instance 1 and VXLAN 2 using BGP instance 2. In this example, the BGP instance ID matches the VXLAN instance ID, but that is not required. Each VXLAN instance has a different VNI and a different BGP instance.

```
# on PE-2:
configure {
    service {
        system {
            bgp-auto-rd-range {
                ip-address 10.0.0.1
                community-value {
                    start 60000
                    end 65000
                }
            }
        }
        vpls "VPLS 1" {
            admin-state enable
            description "dual evpn-vxlan VPLS"
            service-id 1
            customer "1"
            vxlan {
                instance 1 {
                    vni 11
                }
                instance 2 {
                    vni 12
                }
            }
            bgp 1 {
                route-distinguisher auto-rd
                route-target {
                    export "target:64500:11"
                    import "target:64500:11"
                }
            }
            bgp 2 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                route-distinguisher auto-rd
                route-target {
                    export "target:64500:12"
                    import "target:64500:12"
                }
            }
            bgp-evpn {
                evi 1
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
                vxlan 2 {
                    admin-state enable
                    vxlan-instance 2
                }
            }
        }
```

When different BGP instances are configured, the auto-derived route distinguishers (RDs) in BGP instance 1 and BGP instance 2 are different, as follows:

```
[/]
A:admin@PE-2# show service id 1 bgp 1 | match "Route Dist"
Route Dist           : auto-rd
Oper Route Dist      : 10.0.0.1:60000

[/]
A:admin@PE-2# show service id 1 bgp 2 | match "Route Dist"
Route Dist           : auto-rd
Oper Route Dist      : 10.0.0.1:60001
```

Dual EVPN-VXLAN instance VPLSs can contain SAPs in SR OS Release 19.10.R1, and later. However, dual EVPN-VXLAN instance VPLSs cannot contain any SDP bindings in SR OS Release 21.7.R1, as follows:

```
*[ex:/configure service vpls "VPLS 1" spoke-sdp 21:1]
A:admin@PE-2# commit
MINOR: MGMT_CORE #4001: configure service vpls "VPLS 1" - multiple bgp-evpn instances not
 supported with local mesh or spoke sdp
MINOR: MGMT_CORE #4001: configure service vpls "VPLS 1" - multi-instance vxlan not supported
 with sdp-bindings in service
```

Note:
This chapter describes the redundancy based on an Anycast solution, as supported in SR OS Release 16.0, and later. For I-ES based redundancy scenarios as supported in SR OS Release 19.10, and later, see chapter *EVPN Multi-Homing on Dual EVPN-VXLAN BGP Instance VPLS*.

To provide DC GW redundancy, an anycast IP address can be configured for the dual EVPN-VXLAN instance VPLSs on the DC GWs.

EVPN route types 2 and 3 are processed by dual EVPN-VXLAN VPLS services as follows:

- Route type 2: MAC/IP routes
  - MAC/IP routes received in a BGP instance will be imported and — according to the selection rules — installed in the FDB.
  - Active MAC routes are re-advertised in the other BGP instance with new BGP attributes (RD, route target (RT), and so on).

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

    - Only the best EVPN MAC route is redistributed.
    - The MAC/IP information and the sticky bit are propagated. The only exception is the Ethernet Segment Identifier (ESI). A non-zero ESI will be reset unless the auto-disc advertise command is enabled.
    - When an attribute has changed for a redistributed MAC route, the MAC route will be updated if it is still the best route. For example, an update of the sequence number or the sticky bit can trigger a redistribution.
- Route type 3: inclusive multicast routes
  – EVPN inclusive multicast routes are generated independently for each instance with the proper BGP extended communities.
  – Ingress Replication (IR) or Assisted Replication (AR) Inclusive Multicast Ethernet Tag (IMET) routes are supported.
  – The inclusive multicast originating IP can be configured with an anycast address:
    - The configured originating IP address is encoded in the originating IP field of the IMET-IR routes; the originating IP field of the IMET-AR routes is still derived from the assisted replication IP value in the service system settings for VXLAN.
    - If a router receives two IMET routes with the same originating IP address, different RDs, and different next-hops, it sets up two bindings: one to each next-hop.
    - If a router receives two IMET routes with the same originating IP address, the same RD, but different next-hops, it sets up one binding to the next-hop with the lowest IP address.
    - If a router receives two IMET routes with the same originating IP address, different RDs, but the same next-hop, it sets up one binding to the next-hop.
    - A DC GW will not set up a binding to its DC GW peer if the received originating IP equals its own originating IP, regardless of whether the local RD and the remote RD are the same or different.

Configuration

Figure 68: Example topology with VPLS 1 and anycast addresses shows the example topology. Redundancy is based on anycast: on the DC GWs PE-2 and PE-3, anycast address 23.23.23.23 is configured as inclusive multicast originating IP; on PE-4 and PE-5 in DC-2, the anycast address is 45.45.45.45. However, no Ethernet segments are used in this example.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 68: Example topology with VPLS 1 and anycast addresses

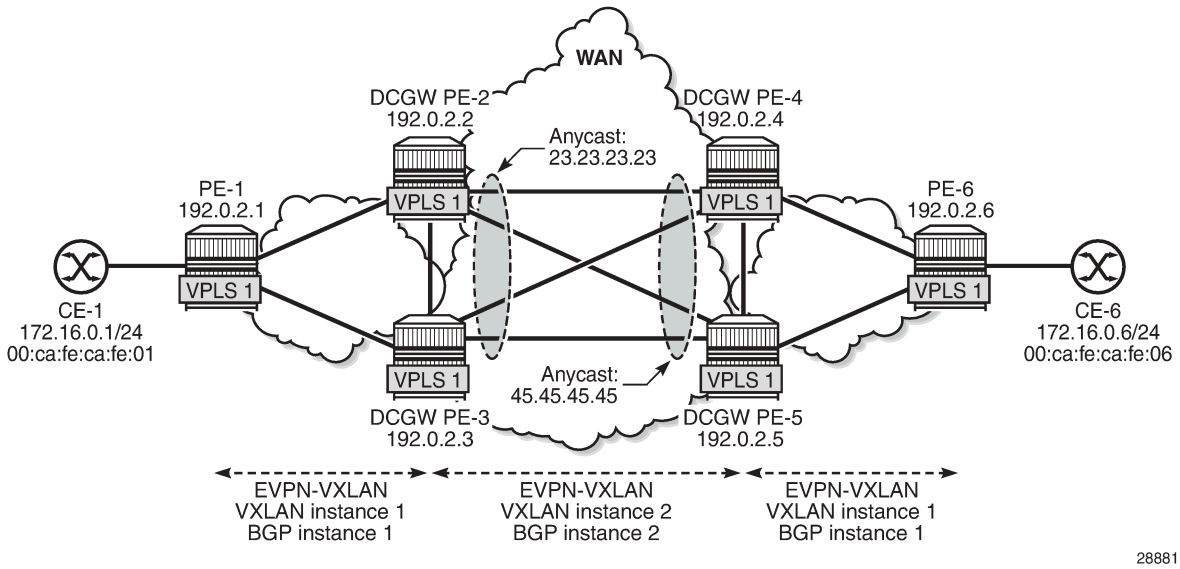

The initial configuration includes:

- Cards, MDAs, ports
- Router interfaces
- IS-IS as IGP (level 1 in the DCs and level 2 in the WAN)

MPLS is not configured in any of these networks.

BGP configuration

BGP is configured for the EVPN address family on all nodes. Figure 69: Example topology with BGP groups shows the BGP groups: on the DC GWs, both BGP group "DC" and "WAN" are defined. Route policies ensure that only DC routes are forwarded to DC neighbors and only WAN routes are forwarded to WAN neighbors. Also, DC GWs need to drop BGP-EVPN routes from the local peer DC GW.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 69: Example topology with BGP groups

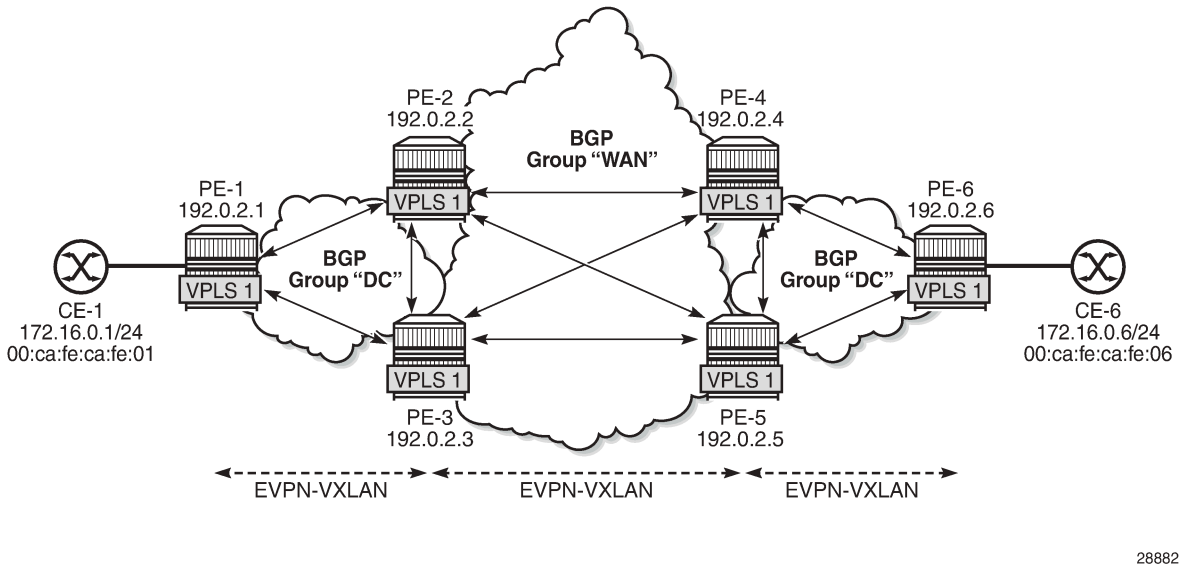

The BGP configuration on PE-1 is as follows. The configuration on PE-6 is similar.

```
# on PE-1:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            family {
                ipv4 false
                evpn true
            }
            rapid-update {
                evpn true
            }
            group "DC" {
                type internal
            }
            neighbor "192.0.2.2" {
                group "DC"
            }
            neighbor "192.0.2.3" {
                group "DC"
            }
```

On the DC GWs, two BGP groups are defined: one for the DC group and one for the WAN group. Export policies ensure that only DC routes are exported to the DC group and only WAN routes are exported to the WAN group. Import policies ensure that routes from the local DC GW are dropped; for example, PE-2 drops routes from PE-3, and vice versa. The policies will be described later. The BGP configuration on PE-2 is as follows. The BGP configuration on the other DC GWs is similar.

```
# on PE-2:
configure {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            family {
                ipv4 false
                evpn true
            }
            rapid-update {
                evpn true
            }
            group "DC" {
                type internal
                import {
                    policy ["drop SOO-DCGW-23"]
                }
                export {
                    policy ["allow only DC and add SOO"]
                }
            }
            group "WAN" {
                type internal
                import {
                    policy ["drop SOO-DCGW-23"]
                }
                export {
                    policy ["allow only WAN and add SOO"]
                }
            }
            neighbor "192.0.2.1" {
                group "DC"
            }
            neighbor "192.0.2.3" {
                group "DC"
            }
            neighbor "192.0.2.4" {
                group "WAN"
            }
            neighbor "192.0.2.5" {
                group "WAN"
            }
```

Route policies

The route policies are equivalent to the policies described in the chapter EVPN-MPLS Interconnect for EVPN-VXLAN VPLS Services. In this example, no filtering can be done based on the encapsulation extended community (VXLAN versus MPLS), because only VXLAN is used in the DCs and the WAN. Therefore, the route tag is used as a criterion instead in the route policies "allow only DC and add SOO" (Site of Origin) and "allow only WAN and add SOO". When two BGP instances for the same encapsulation are configured in a VPLS, different route tags in each BGP instance are required. In this example, route tag 11 is used in BGP instance 1 in the DCs and route tag 12 is used in BGP instance 2 in the WAN.

When redistributing to the other BGP instance, route filtering toward DC or WAN will be based on the route tags. Export policy "allow only DC and add SOO" drops routes with WAN route tag 12. Likewise, export policy "allow only WAN and add SOO" drops routes with DC route tag 11. Filtering matching on route tags on EVPN BGP instances is scalable, because only two route tags are required per PE. Filtering matching

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

on route target (RT) is also possible, but in that case, two RTs per service are required. This does not scale well and is cumbersome.

The export policy "allow only DC and add SOO" ensures that EVPN routes with route tag 12 are dropped and only DC routes are forwarded. This route policy is applied in the BGP group "DC" context. Likewise, the export policy "allow only WAN and add SOO" drops EVPN routes with route tag 11, so that only WAN EVPN routes are forwarded.

Both policies also add a site of origin, such as "SOO-23" for PE-2 and PE-3, and "SOO-45" for PE-4 and PE-5. This SOO is used for filtering in the import policies "drop SOO-DCGW-23" and "drop SOO-DCGW-45" to ensure that, for instance, PE-2 drops routes advertised by the local peer PE-3 with the same SOO-23, and vice versa. Likewise, PE-4 drops routes advertised by its local peer PE-5 with the same SOO-45, and vice versa.

The following policies are configured on DC GWs PE-2 and PE-3:

```
# on PE-2, PE-3:
configure {
    policy-options {
        community "SOO-23" {
            member "origin:64500:23" { }
        }
        policy-statement "allow only DC and add SOO" {
            entry 10 {
                from {
                    family [evpn]
                    tag 12
                }
                action {
                    action-type reject
                }
            }
            entry 20 {
                from {
                    family [evpn]
                }
                action {
                    action-type accept
                    community {
                        add ["SOO-23"]
                    }
                }
            }
        }
        policy-statement "allow only WAN and add SOO" {
            entry 10 {
                from {
                    family [evpn]
                    tag 11
                }
                action {
                    action-type reject
                }
            }
            entry 20 {
                from {
                    family [evpn]
                }
                action {
                    action-type accept
                    community {
                        add ["SOO-23"]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    }
                }
            }
        }
        policy-statement "drop SOO-DCGW-23" {
            entry 10 {
                from {
                    family [evpn]
                    community {
                        name "SOO-23"
                    }
                }
                action {
                    action-type reject
                }
            }
        }
```

The following policies are configured on DC GWs PE-4 and PE-5:

```
# on PE-4, PE-5:
configure {
    policy-options {
        community "SOO-45" {
            member "origin:64500:45" { }
        }
        policy-statement "allow only DC and add SOO" {
            entry 10 {
                from {
                    family [evpn]
                    tag 12
                }
                action {
                    action-type reject
                }
            }
            entry 20 {
                from {
                    family [evpn]
                }
                action {
                    action-type accept
                    community {
                        add ["SOO-45"]
                    }
                }
            }
        }
        policy-statement "allow only WAN and add SOO" {
            entry 10 {
                from {
                    family [evpn]
                    tag 11
                }
                action {
                    action-type reject
                }
            }
            entry 20 {
                from {
                    family [evpn]
                }
                action {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    action-type accept
                    community {
                        add ["SOO-45"]
                    }
                }
            }
        }
        policy-statement "drop SOO-DCGW-45" {
            entry 10 {
                from {
                    family [evpn]
                    community {
                        name "SOO-45"
                    }
                }
                action {
                    action-type reject
                }
            }
        }
```

VPLS configuration

On PE-2 and PE-3, the service configuration is identical and VPLS 1 is configured as follows. For redundancy, the anycast IP address 23.23.23.23 is configured as inclusive multicast originating IP on PE-2 and PE-3. The RT is the same in all nodes: the RT is 64500:11 in BGP instance 1 of VPLS 1; in BGP instance 2, the RT is 64500:12. The RD is 64500:2311 in BGP instance 1 of VPLS 1 and 64500:2312 in BGP instance 2 of VPLS 1 on PE-2 and PE-3. The RD must be the same in PE-2 and PE-3 because they are part of the anycast group, but the RD in PE-1 must be different.

```
# on PE-2, PE-3:
configure {
    service {
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            vxlan {
                instance 1 {
                    vni 11
                }
                instance 2 {
                    vni 12
                }
            }
            bgp 1 {
                route-distinguisher "64500:2311"
                route-target {
                    export "target:64500:11"
                    import "target:64500:11"
                }
            }
            bgp 2 {
                route-distinguisher "64500:2312"
                route-target {
                    export "target:64500:12"
                    import "target:64500:12"
                }
            }
            bgp-evpn {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                evi 1
                incl-mcast-orig-ip 23.23.23.23
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                    default-route-tag 0xb         # default route tag 11
                }
                vxlan 2 {
                    admin-state enable
                    vxlan-instance 2
                    default-route-tag 0xc         # default route tag 12
                }
            }
        }
```

On PE-1, VPLS 1 is configured with VXLAN instance 1 and BGP instance 1, as follows. The RT is 64500:11 in BGP instance 1 of VPLS 1 on PE-1. The RD (64500:111) in PE-1 is different from the RD (64500:2311) in PE-2 and PE-3.

```
# on PE-1:
configure {
    service {
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            vxlan {
                instance 1 {
                    vni 11
                }
            }
            bgp 1 {
                route-distinguisher "64500:111"
                route-target {
                    export "target:64500:11"
                    import "target:64500:11"
                }
            }
            bgp-evpn {
                evi 1
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
            sap 1/2/1:1 {
            }
        }
```

On PE-4 and PE-5, the service configuration is identical and VPLS 1 is configured as follows. For redundancy, the anycast IP address 45.45.45.45 is configured as inclusive multicast originating IP.

```
# on PE-4, PE-5:
configure {
    service {
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            vxlan {
                instance 1 {
                    vni 11
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
                instance 2 {
                    vni 12
                }
            }
            bgp 1 {
                route-distinguisher "64500:4511"
                route-target {
                    export "target:64500:11"
                    import "target:64500:11"
                }
            }
            bgp 2 {
                route-distinguisher "64500:4512"
                route-target {
                    export "target:64500:12"
                    import "target:64500:12"
                }
            }
            bgp-evpn {
                evi 1
                incl-mcast-orig-ip 45.45.45.45
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                    default-route-tag 0xb         # default route tag 11
                }
                vxlan 2 {
                    admin-state enable
                    vxlan-instance 2
                    default-route-tag 0xc         # default route tag 12
                }
            }
        }
```

Verification

On PE-2, the following EVPN inclusive multicast routes are received. The first route has RD 64500:111, so it applies to BGP instance 1; the last two have RD 64500:4512, so they apply to BGP instance 2. Toward anycast address 45.45.45.45, the lowest IP next-hop 192.0.2.4 (PE-4) is preferred over 192.0.2.5 (PE-5).

```
[/]
A:admin@PE-2# show router bgp routes evpn incl-mcast
===============================================================================
 BGP Router ID:192.0.2.2        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Inclusive-Mcast Routes
===============================================================================
Flag  Route Dist.         OrigAddr
      Tag                 NextHop
-------------------------------------------------------------------------------
u*>i  64500:111           192.0.2.1
      0                   192.0.2.1

u*>i  64500:4512          45.45.45.45
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
      0                   192.0.2.4

*>i   64500:4512          45.45.45.45
      0                   192.0.2.5

-------------------------------------------------------------------------------
Routes : 3
===============================================================================
```

The following shows the VXLAN destinations on PE-2: PE-1 (192.0.2.1) is the VXLAN Tunnel Endpoint (VTEP) in VXLAN instance 1; the VTEP for VXLAN instance 2 is PE-4 (192.0.2.4).

```
[/]
A:admin@PE-2# show service id 1 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
Instance    VTEP Address                            Egress VNI  EvpnStatic Num
 Mcast       Oper State                              L2 PBR     SupBcasDom MACs
-------------------------------------------------------------------------------
1           192.0.2.1                               11          evpn       0
 BUM         Up                                      No          No
2           192.0.2.4                               12          evpn       0
 BUM         Up                                      No          No
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 2
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN-VXLAN Ethernet Segment Dest
===============================================================================
Instance  Eth SegId                      Num. Macs     Last Change
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

The following shows the BGP information for VPLS 1 on PE-1. Only BGP instance 1 is configured. The RD is configured with the value 64500:111 and the RT with the value 64500:11, which are also the operational values. No VSI import or VSI export policies are configured on PE-1.

```
[/]
A:admin@PE-1# show service id 1 bgp

===============================================================================
BGP Information
===============================================================================
Bgp Instance         : 1
Vsi-Import           : None
Vsi-Export           : None
Route Dist           : 64500:111
Oper Route Dist      : 64500:111
Oper RD Type         : configured
Rte-Target Import    : 64500:11            Rte-Target Export: 64500:11
Oper RT Imp Origin   : configured          Oper RT Import   : 64500:11
Oper RT Exp Origin   : configured          Oper RT Export   : 64500:11

PW-Template Id       : None
-------------------------------------------------------------------------------
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

On PE-2, the following information for BGP instance 1 includes the configured and operational RD 64500:2311 and RT 64500:11.

```
[/]
A:admin@PE-2# show service id 1 bgp 1

===============================================================================
BGP Information
===============================================================================
Vsi-Import           : None
Vsi-Export           : None
Route Dist           : 64500:2311
Oper Route Dist      : 64500:2311
Oper RD Type         : configured
Rte-Target Import    : 64500:11             Rte-Target Export: 64500:11
Oper RT Imp Origin   : configured           Oper RT Import   : 64500:11
Oper RT Exp Origin   : configured           Oper RT Export   : 64500:11
PW-Template Id       : None
-------------------------------------------------------------------------------
===============================================================================
```

On PE-2, the following information for BGP instance 1 includes the configured and operational RD 64500:2312 and RT 64500:12.

```
[/]
A:admin@PE-2# show service id 1 bgp 2

===============================================================================
BGP Information
===============================================================================
Vsi-Import           : None
Vsi-Export           : None
Route Dist           : 64500:2312
Oper Route Dist      : 64500:2312
Oper RD Type         : configured
Rte-Target Import    : 64500:12             Rte-Target Export: 64500:12
Oper RT Imp Origin   : configured           Oper RT Import   : 64500:12
Oper RT Exp Origin   : configured           Oper RT Export   : 64500:12
-------------------------------------------------------------------------------
===============================================================================
```

The CEs are simulated by VPRN 11 configured on PE-1 and PE-6. Connectivity between CE-1 and CE-6 is verified as follows:

```
[/]
A:admin@PE-1# ping 172.16.0.6 router-instance "VPRN 11" interval 0.1 output-format summary
PING 172.16.0.6 56 data bytes
!!!!!
---- 172.16.0.6 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 4.19ms, avg = 4.56ms, max = 5.30ms, stddev = 0.393ms
```

The following two EVPN MAC routes are accepted on PE-1, which has only BGP instance 1 and VXLAN instance 1 enabled, and the VNI is 11. The used EVPN MAC route for MAC address 00:ca:fe:ca:fe:06 has PE-2 (192.0.2.2) as next-hop. The second route for the same MAC address has PE-3 (192.0.2.3) as next-hop, but it is not preferred, so it is not used.

```
[/]
A:admin@PE-1# show router bgp routes evpn mac
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
 BGP Router ID:192.0.2.1        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN MAC Routes
===============================================================================
Flag  Route Dist.          MacAddr           ESI
      Tag                  Mac Mobility      Label1
                           Ip Address
                           NextHop
-------------------------------------------------------------------------------
u*>i  64500:2311           00:ca:fe:ca:fe:06 ESI-0
      0                    Seq:0             VNI 11
                           n/a
                           192.0.2.2

*>i   64500:2311           00:ca:fe:ca:fe:06 ESI-0
      0                    Seq:0             VNI 11
                           n/a
                           192.0.2.3

-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

On PE-2, the following three EVPN MAC routes are accepted. The first route has PE-1 (192.0.2.1) as next-hop and is received in BGP instance 1, which corresponds to VXLAN 1 and VNI 11. The latter two routes are received in BGP instance 2 for VXLAN instance 2 with VNI 12. These routes both have RD 64500:4512 and the route with the lowest IP next-hop is preferred, so the route to PE-4 (192.0.2.4) is used. The EVPN MAC routes on PE-3 are similar.

```
[/]
A:admin@PE-2# show router bgp routes evpn mac
===============================================================================
 BGP Router ID:192.0.2.2        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN MAC Routes
===============================================================================
Flag  Route Dist.          MacAddr           ESI
      Tag                  Mac Mobility      Label1
                           Ip Address
                           NextHop
-------------------------------------------------------------------------------
u*>i  64500:111            00:ca:fe:ca:fe:01 ESI-0
      0                    Seq:0             VNI 11
                           n/a
                           192.0.2.1

u*>i  64500:4512           00:ca:fe:ca:fe:06 ESI-0
      0                    Seq:0             VNI 12
                           n/a
                           192.0.2.4
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
*>i   64500:4512          00:ca:fe:ca:fe:06 ESI-0
      0                   Seq:0             VNI 12
                          n/a
                          192.0.2.5

-------------------------------------------------------------------------------
Routes : 3
===============================================================================
```

The EVPN MAC routes on the nodes in DC-2 are also similar.

The following FDB for VPLS 1 on PE-2 shows that MAC address 00:ca:fe:ca:fe:01 is learned from an EVPN MAC route in VXLAN 1 from 192.0.2.1 (PE-1); MAC address 00:ca:fe:ca:fe:06 is learned in VXLAN 2 from 192.0.2.4 (PE-4). For routes with the same RD but different next-hops, the router processes only the route with the lowest IP next-hop.

```
[/]
A:admin@PE-2# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
1          00:ca:fe:ca:fe:01 vxlan-1:                Evpn     08/13/21 15:42:24
                             192.0.2.1:11
1          00:ca:fe:ca:fe:06 vxlan-2:                Evpn     08/13/21 15:42:33
                             192.0.2.4:12
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

Conclusion

With dual EVPN-VXLAN instance VPLS services, service providers can deploy DCI scenarios with end-to-end VXLAN.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Domain Path Attribute for VPRN BGP Routes

This chapter provides information about the domain path attribute for VPRN BGP routes.

Topics in this chapter include:

Applicability

The information and configuration in this chapter are based on SR OS Release 22.7.R1. The domain path (D-path) attribute is supported in SR OS Release 21.10.R1 and later.

Overview

The D-path attribute can be used for route traceability, BGP best path selection, and loop prevention in networks that expand multiple IP-VPN and EVPN domains.

The D-path attribute is a sequence of domain segments, where each domain segment is represented by a domain ID in combination with an inter-subnet forwarding (ISF) subaddress family indicator (SAFI). The D-path attribute is added or modified by gateways (GWs) that import BGP-EVPN route type 5 (RT-5) or IP-VPN routes into a VPRN route table and export these prefixes as BGP-EVPN RT-5 or IP-VPN routes to their neighbors. Any PE that imports a prefix route does not install the route in the VPRN route table if the D-path attribute contains a domain segment where the domain ID matches a local domain ID, as shown in the figure Figure 70: Loop prevention in networks with multiple IP-VPN and EVPN domains.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 70: Loop prevention in networks with multiple IP-VPN and EVPN domains

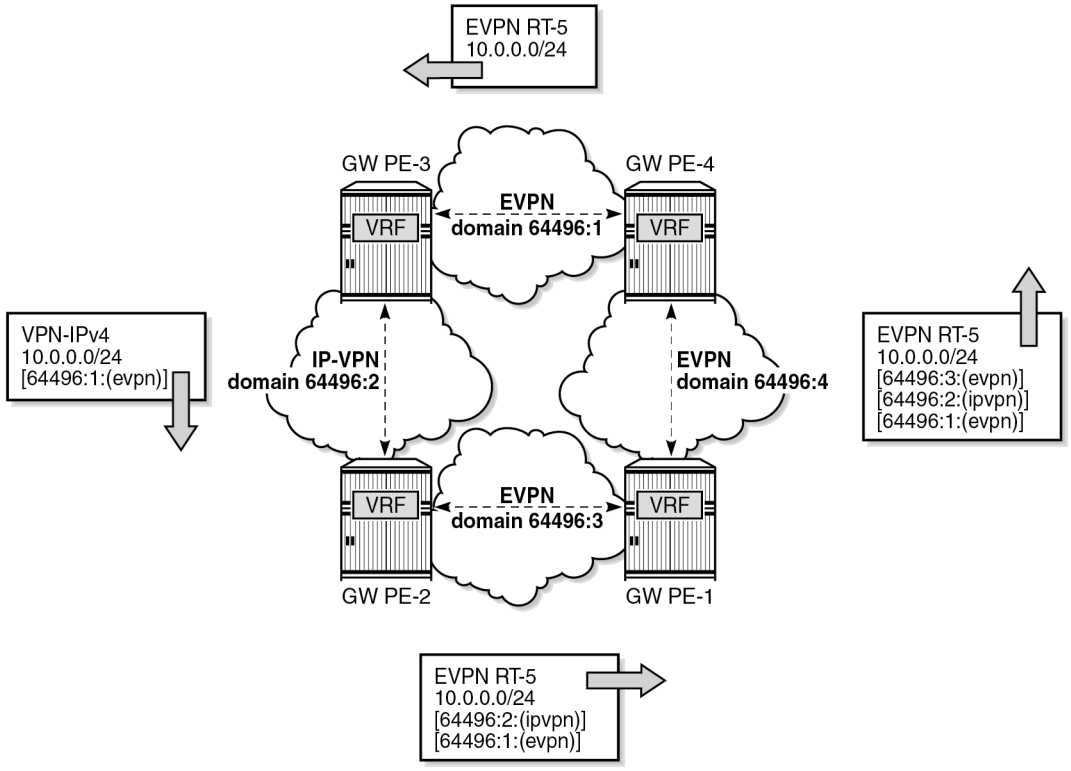

All PEs in the figure Figure 70: Loop prevention in networks with multiple IP-VPN and EVPN domains are GWs. PE-4 exports local prefix 10.0.0.0/24 as an EVPN RT-5 route without the D-path attribute when no domain ID is configured for local routes. PE-3 accepts this route. Domain ID 64496:1 is defined in PE-4 and PE-3, but the domain segment 64496:1:(evpn) is only added by GW PE-3 where the prefix is exported as an IP-VPN route instead of an EVPN RT-5 route. GW PE-2 accepts this route and modifies the D-path attribute by prepending domain segment 64496:2:(ipvpn) when exporting prefix 10.0.0.0/24 as an EVPN RT-5 route. PE-1 accepts this route. When PE-1 exports the prefix as an EVPN RT-5 route to PE-4, it prepends domain segment 64496:3:(evpn) to the D-path attribute. The VRF on PE-4 cannot import this prefix because the D-path attribute contains domain ID 64496:1, which is defined on PE-4.

The figure Figure 71: D-path attribute shows the D-path attribute as defined in *draft-ietf-bess-evpn-ipvpn-interworking*.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 71: D-path attribute

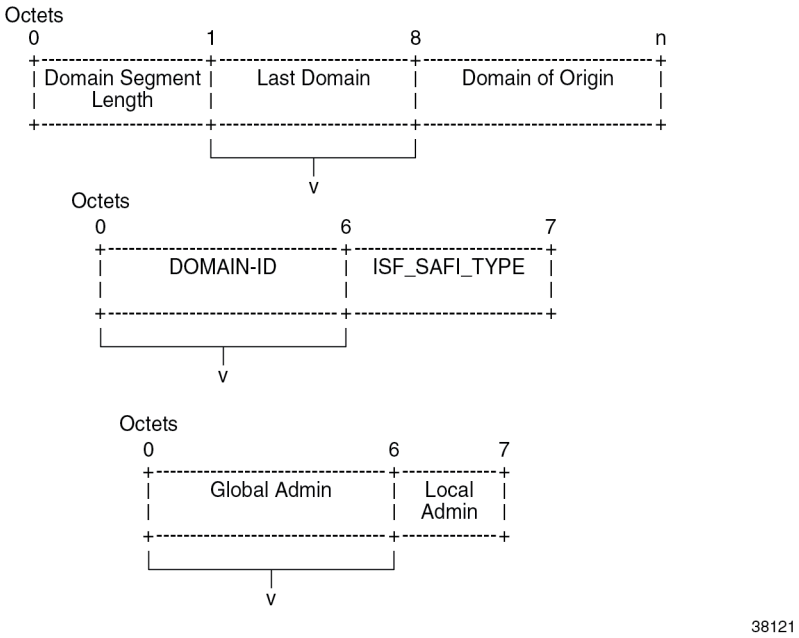

The D-path attribute is composed of a sequence of domain segments. Each domain segment consists of a domain ID and a SAFI type. The domain ID represents the domain and is composed of a 4-octet global administrator subfield and a 2-octet local administrator subfield. The global administrator subfield must have a value that is unique for the domain; for example, an autonomous system number (ASN). The 1-octet SAFI field can have the following values:

- 0 for local ISF routes
- 1 for PE-CE BGP domains
- 70 for EVPN domains
- 128 for IP-VPN domains

The domain ID can be configured on:

- VPRN BGP-EVPN MPLS and BGP-EVPN SRv6 instances (EVPN interface-less (EVPN-IFL))
- VPRN BGP-IPVPN MPLS and BGP-IPVPN SRv6 instances
- R-VPLS BGP-EVPN MPLS and BGP-EVPN VXLAN instances (EVPN interface-ful (EVPN-IFF))
- VPRN BGP neighbors (PE-CE)
- VPRN level (for local routes). When configured on the VPRN level, using the optional **local-routes-domain-id** command, the PE advertises its direct, static, or IGP routes with a D-path attribute.

Domain IDs can be modified while the service is operational. Modifying the domain ID initiates a route refresh for all address families associated with the VPRN.

A PE receiving a prefix route with a D-path attribute containing one of its own domain IDs detects a routing loop and does not install the route in the VPRN route table.

The D-path attribute length can influence the BGP best path selection. In the BGP decision process, the shorter D-path is preferred, unless the **d-path-length-ignore** command is configured.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Configuration

The figure Figure 72: Example topology with VPRN 10 and its domain IDs shows an example topology where PE-6 exports EVPN RT-5 routes 172.31.6.0/24 and 2001:db8::31:6:0/120 to route reflector RR-5, whereas PE-7 exports IP-VPN routes 172.31.7.0/24 and 2001:db8::31:7:0/120 to RR-5. LDP tunnels are used between PE-4, RR-5, PE-6, and PE-7; SRv6 tunnels are used between PE-2, PE-3, and PE-4; SR-OSPF tunnels are used between PE-1, PE-2, and PE-3.

Figure 72: Example topology with VPRN 10 and its domain IDs

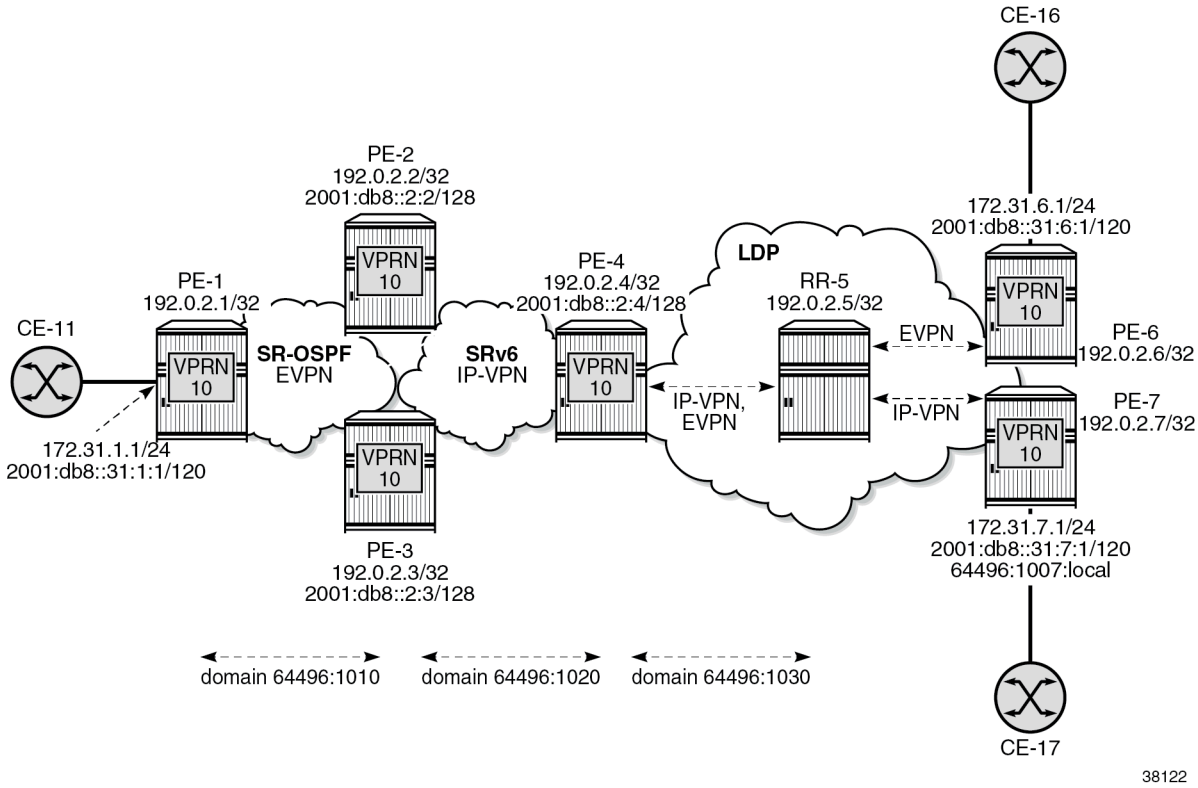

The initial configuration includes:

- cards, MDAs, ports
- router interfaces
- OSPF as IGP on PE-1, PE-2, and PE-3
- IS-IS as IGP on PE-2, PE-3, PE-4, RR-5, PE-6, and PE-7
- SR-OSPF on PE-1, PE-2, and PE-3
- SRv6 on PE-2, PE-3, and PE-4, configured as in the "Segment Routing over IPv6" chapter in the Segment Routing and PCE volume in the *7450 ESS, 7750 SR, and 7950 XRS MD-CLI Advanced Configuration Guide - Part I*.
- LDP on PE-4, RR-5, PE-6, and PE-7

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The BGP configuration on PE-1 is as follows:

```
# on PE-1:
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            split-horizon true
            rapid-update {
                evpn true
            }
            group "internal1" {
                type internal
                family {
                    evpn true
                }
            }
            neighbor "192.0.2.2" {
                group "internal1"
            }
            neighbor "192.0.2.3" {
                group "internal1"
            }
        }
```

```
# on PE-2 (similar configuration on PE-3):
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            router-id 192.0.2.2             # on PE-3: 192.0.2.3
            advertise-inactive true
            rapid-withdrawal true
            peer-ip-tracking true
            split-horizon true
            rapid-update {
                vpn-ipv4 true
                vpn-ipv6 true
                evpn true
            }
            group "internal1" {
                next-hop-self true
                type internal
                local-address 192.0.2.2        # on PE-3: 192.0.2.3
                family {
                    evpn true
                }
            }
            group "internal2" {
                next-hop-self true
                type internal
                local-address 2001:db8::2:2     # on PE-3: 2001:db8::2:3
                family {
                    vpn-ipv4 true
                    vpn-ipv6 true
                }
                extended-nh-encoding {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    vpn-ipv4 true
                    ipv4 true
                }
                advertise-ipv6-next-hops {
                    vpn-ipv6 true
                    vpn-ipv4 true
                }
            }
            neighbor "192.0.2.1" {
                group "internal1"
            }
            neighbor "192.0.2.3" {            # on PE-3: 192.0.2.2
                group "internal1"
            }
            neighbor "2001:db8::2:3" {        # on PE-3: 2001:db8::2:2
                group "internal2"
            }
            neighbor "2001:db8::2:4" {
                group "internal2"
            }
        }
```

```
# on PE-4:
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            router-id 192.0.2.4
            advertise-inactive true
            rapid-withdrawal true
            peer-ip-tracking true
            split-horizon true
            rapid-update {
                vpn-ipv4 true
                vpn-ipv6 true
                evpn true
            }
            group "internal2" {
                next-hop-self true
                type internal
                local-address 2001:db8::2:4
                family {
                    vpn-ipv4 true
                    vpn-ipv6 true
                }
                extended-nh-encoding {
                    vpn-ipv4 true
                    ipv4 true
                }
                advertise-ipv6-next-hops {
                    vpn-ipv6 true
                    vpn-ipv4 true
                }
            }
            group "internal3" {
                next-hop-self true
                type internal
                local-address 192.0.2.4
                family {
                    vpn-ipv4 true
                    vpn-ipv6 true
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    evpn true
                }
            }
            neighbor "192.0.2.5" {
                group "internal3"
            }
            neighbor "2001:db8::2:2" {
                group "internal2"
            }
            neighbor "2001:db8::2:3" {
                group "internal2"
            }
        }
```

```
# on RR-5: only EVPN toward PE-6; only IP-VPN toward PE-7:
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            split-horizon true
            rapid-update {
                vpn-ipv4 true
                vpn-ipv6 true
                evpn true
            }
            group "internal3" {
                type internal
                cluster {
                    cluster-id 192.0.2.5
                }
            }
            neighbor "192.0.2.4" {
                group "internal3"
                family {
                    vpn-ipv4 true
                    vpn-ipv6 true
                    evpn true
                }
            }
            neighbor "192.0.2.6" {
                group "internal3"
                family {
                    evpn true
                }
            }
            neighbor "192.0.2.7" {
                group "internal3"
                family {
                    vpn-ipv4 true
                    vpn-ipv6 true
                }
            }
        }
```

```
# on PE-6:
configure {
    router "Base" {
        autonomous-system 64496
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            split-horizon true
            rapid-update {
                evpn true
            }
            group "internal3" {
                type internal
            }
            neighbor "192.0.2.5" {
                group "internal3"
                family {
                    evpn true
                }
            }
        }
```

```
# on PE-7:
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            split-horizon true
            rapid-update {
                vpn-ipv4 true
                vpn-ipv6 true
            }
            group "internal3" {
                type internal
            }
            neighbor "192.0.2.5" {
                group "internal3"
                family {
                    vpn-ipv4 true
                    vpn-ipv6 true
                }
            }
        }
```

Domain IDs in VPRN BGP-EVPN MPLS and SRv6 instances

On PE-1, VPRN 10 is configured without domain ID in the **bgp-evpn mpls 1** context:

```
# on PE-1:
configure {
    service {
        vprn "VPRN 10" {
            admin-state enable
            service-id 10
            customer "1"
            autonomous-system 64496
            bgp-evpn {
                mpls 1 {
                    admin-state enable
                    route-distinguisher "192.0.2.1:10"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    vrf-target {
                        community "target:64496:10"
                    }
                    auto-bind-tunnel {
                        resolution filter
                        resolution-filter {
                            sr-ospf true
                        }
                    }
                }
            }
            interface "int-PE-1-CE-11" {
                ipv4 {
                    primary {
                        address 172.31.1.1
                        prefix-length 24
                    }
                }
                sap 1/1/c5/1:10 {
                }
                ipv6 {
                    address 2001:db8::31:1:1 {
                        prefix-length 120
                    }
                }
            }
        }
```

Domain ID 64496:1010 is configured in the **bgp-evpn mpls 1** context on GWs PE-2 and PE-3, whereas domain ID 64496:1020 is configured in the **bgp-ipvpn segment-routing-v6** context on PE-2, PE-3, and PE-4. Domain ID 64496:1030 is configured for IP-VPN and for BGP-EVPN on PE-4.

On PE-2, VPRN 10 is configured as follows. The configuration on PE-3 is similar.

```
# on GW PE-2:
configure {
    service {
         vprn "VPRN 10" {
            admin-state enable
            service-id 10
            customer "1"
            autonomous-system 64496
            segment-routing-v6 1 {
                locator "PE-2_loc" {                      # on PE-3:"PE-3_loc"
                    function {
                        end-dt4 {
                        }
                        end-dt6 {
                        }
                    }
                }
            }
            bgp-evpn {
                mpls 1 {
                    admin-state enable
                    route-distinguisher "192.0.2.2:10"   # on PE-3: 192.0.2.3:10
                    domain-id "64496:1010"
                    vrf-target {
                        community "target:64496:10"
                    }
                    auto-bind-tunnel {
                        resolution filter
                        resolution-filter {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        sr-ospf true
                    }
                }
            }
        }
        bgp-ipvpn {
            segment-routing-v6 1 {
                admin-state enable
                route-distinguisher "192.0.2.2:16"   # on PE-3: 192.0.2.3:16
                source-address 2001:db8::2:2         # on PE-3: 2001:db8::2:3
                domain-id "64496:1020"
                vrf-target {
                    community "target:64496:10"
                }
                srv6 {
                    instance 1
                    default-locator "PE-2_loc"       # on PE-3:"PE-3_loc"
                }
            }
        }
    }
```

On GW PE-4, VPRN 10 is configured with two domain IDs: domain ID 1020 for IP-VPN over SRv6 and domain ID 1030 for IP-VPN over MPLS and for EVPN over MPLS.

```
# on GW PE-4:
configure {
    service {
        vprn "VPRN 10" {
            admin-state enable
            service-id 10
            customer "1"
            autonomous-system 64496
            allow-export-bgp-vpn true
            segment-routing-v6 1 {
                locator "PE-4_loc" {
                    function {
                        end-dt4 {
                        }
                        end-dt6 {
                        }
                    }
                }
            }
            bgp-evpn {
                mpls 1 {
                    admin-state enable
                    route-distinguisher "192.0.2.4:10"
                    domain-id "64496:1030"
                    vrf-target {
                        community "target:64496:10"
                    }
                    auto-bind-tunnel {
                        resolution filter
                        resolution-filter {
                            ldp true
                        }
                    }
                }
            }
            bgp-ipvpn {
                mpls {
                    admin-state enable
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    route-distinguisher "192.0.2.4:10"
                    domain-id "64496:1030"
                    vrf-target {
                        community "target:64496:10"
                    }
                    auto-bind-tunnel {
                        resolution filter
                        resolution-filter {
                            ldp true
                        }
                    }
                }
                segment-routing-v6 1 {
                    admin-state enable
                    route-distinguisher "192.0.2.4:16"
                    source-address 2001:db8::2:4
                    domain-id "64496:1020"
                    vrf-target {
                        community "target:64496:10"
                    }
                    srv6 {
                        instance 1
                        default-locator "PE-4_loc"
                    }
                }
            }
```

For completeness, the configuration on VPRN 10 on PE-6 and PE-7 is also shown. PE-6 has no domain ID configured:

```
# on PE-6:
configure {
    service {
        vprn "VPRN 10" {
            admin-state enable
            service-id 10
            customer "1"
            autonomous-system 64496
            bgp-evpn {
                mpls 1 {
                    admin-state enable
                    route-distinguisher "192.0.2.6:10"
                    vrf-target {
                        community "target:64496:10"
                    }
                    auto-bind-tunnel {
                        resolution filter
                        resolution-filter {
                            ldp true
                        }
                    }
                }
            }
            interface "int-PE-6-CE-16" {
                ipv4 {
                    primary {
                        address 172.31.6.1
                        prefix-length 24
                    }
                }
                sap 1/1/c5/1:10 {
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                ipv6 {
                    address 2001:db8::31:6:1 {
                        prefix-length 120
                    }
                }
            }
        }
```

PE-7 does not have a domain ID configured in the **bgp-ipvpn mpls** context, but it has a local domain ID configured: 64496:1007:

```
# on PE-7:
configure {
    service {
        vprn "VPRN 10" {
            admin-state enable
            service-id 10
            customer "1"
            autonomous-system 64496
            local-routes-domain-id "64496:1007"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "192.0.2.7:10"
                    vrf-target {
                        community "target:64496:10"
                    }
                    auto-bind-tunnel {
                        resolution filter
                        resolution-filter {
                            ldp true
                        }
                    }
                }
            }
            interface "int-PE-7-CE-17" {
                ipv4 {
                    primary {
                        address 172.31.7.1
                        prefix-length 24
                    }
                }
                sap 1/1/c5/1:10 {
                }
                ipv6 {
                    address 2001:db8::31:7:1 {
                        prefix-length 120
                    }
                }
            }
        }
```

The following commands on PE-4 display the domain ID for BGP-IPVPN and BGP-EVPN. For BGP-IPVPN, domain ID 64496:1030 is configured in the EVPN-MPLS domain and domain ID 64496:1020 is configured in the SRv6 domain:

```
[/]
A:admin@PE-4# show service id 10 bgp-ipvpn

===============================================================================
Service 10 BGP-IPVPN MPLS Information
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Admin State       : Up
VRF Import        : None
VRF Export        : None
Route Dist.       : None
Oper Route Dist   : 192.0.2.4:10
Oper RD Type      : configured
Route Target      : target:64496:10
Route Target Impor: None
Route Target Expor: None
Domain-Id         : 64496:1030
Dyn Egr Lbl Limit : Disabled

Auto-Bind Tunnel
Resolution        : disabled            Strict Tnl Tag    : False
ECMP              : 0                   Flex Algo FB      : False
Weighted ECMP     : False
BGP Instance      : 1
Filter Tunnel Type: (Not Specified)
===============================================================================

===============================================================================
Service 10 BGP-IPVPN Segment-Routing-V6 Information
===============================================================================

Admin State       : Up
VRF Import        : None
VRF Export        : None
Route Dist.       : 192.0.2.4:16
Oper Route Dist   : 192.0.2.4:16
Oper RD Type      : configured
Route Target      : target:64496:10
Route Target Expor: None
Route Target Impor: None
Def Route Tag     : 0x0
Route Resolution  : route-table

Srv6 Instance     : 1
Default Locator   : PE-4_loc
Source Address    : 2001:db8::2:4
Domain-Id         : 64496:1020

===============================================================================
```

For BGP-EVPN, domain ID 64496:1030 is configured in the EVPN-MPLS domain:

```
[/]
A:admin@PE-4# show service id 10 bgp-evpn

===============================================================================
BGP EVPN MPLS Table
===============================================================================
Admin State        : Up
VRF Import         : None
VRF Export         : None
Route Dist.        : 192.0.2.4:10
Oper Route Dist.   : 192.0.2.4:10
Oper RD Type       : configured
Route Target       : target:64496:10
Route Target Import: None
Route Target Export: None
Default Route Tag  : None
Domain-Id          : 64496:1030
Dyn Egr Lbl Limit  : Disabled
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Advertise          : Disabled
Weighted ECMP      : Disabled

Auto-Bind Tunnel
Resolution         : filter                Strict Tnl Tag : False
ECMP               : 1                     Flex Algo FB   : False
BGP Instance       : 1
Filter Tunnel Types: ldp

Tunnel Encap
MPLS               : True                  MPLSoUDP       : False
===============================================================================
```

VPRN BGP routes for prefix 172.31.6.0/24

PE-6 advertises prefix 172.31.6.0/24 as an EVPN-IFL route without the D-path attribute, as follows:

```
# on PE-6:
2 2022/09/06 10:46:07.053 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.5
"Peer 1: 192.0.2.5: UPDATE
Peer 1: 192.0.2.5 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 82
    Flag: 0x90 Type: 14 Len: 45 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.6
        Type: EVPN-IP-PREFIX Len: 34 RD: 192.0.2.6:10, ESI: ESI-0, tag: 0, ip_prefix:
 172.31.6.0/24 gw_ip 0.0.0.0 Label: 8388528 (Raw Label: 0x7fffb0)
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64496:10
        bgp-tunnel-encap:MPLS
```

RR-5 forwards prefix 172.31.6.0/24 as an EVPN-IFL route without the D-path attribute, as follows:

```
# on RR-5:
12 2022/09/06 10:46:07.053 CEST MINOR: DEBUG #2001 Base Peer 1: 192.0.2.4
"Peer 1: 192.0.2.4: UPDATE
Peer 1: 192.0.2.4 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 96
    Flag: 0x90 Type: 14 Len: 45 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.6
        Type: EVPN-IP-PREFIX Len: 34 RD: 192.0.2.6:10, ESI: ESI-0, tag: 0, ip_prefix:
 172.31.6.0/24 gw_ip 0.0.0.0 Label: 8388528 (Raw Label: 0x7fffb0)
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0x80 Type: 9 Len: 4 Originator ID: 192.0.2.6
    Flag: 0x80 Type: 10 Len: 4 Cluster ID:
        192.0.2.5
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64496:10
        bgp-tunnel-encap:MPLS
"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

PE-4 adds a D-path attribute when advertising prefix 172.31.6.0/24 as a VPN-IPv4 route to PE-2 (or PE-3):

```
29 2022/09/06 10:46:07.055 CEST MINOR: DEBUG #2001 Base Peer 1: 2001:db8::2:2
"Peer 1: 2001:db8::2:2: UPDATE
Peer 1: 2001:db8::2:2 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 98
    Flag: 0x90 Type: 14 Len: 44 Multiprotocol Reachable NLRI:
        Address Family VPN_IPV4
        NextHop len 24 NextHop 2001:db8::2:4
        172.31.6.0/24 RD 192.0.2.4:10 Label 524281 (Raw label 0x7fff91)
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0x80 Type: 9 Len: 4 Originator ID: 192.0.2.6
    Flag: 0x80 Type: 10 Len: 4 Cluster ID:
        192.0.2.5
    Flag: 0xc0 Type: 16 Len: 8 Extended Community:
        target:64496:10
    Flag: 0xc0 Type: 36 Len: 8 D-PATH:[64496:1030:(evpn)]
"
```

PE-2 prepends domain segment 64496:1020:(ipvpn) to the D-path attribute when advertising prefix 172.31.6.0/24 in an EVPN-IFL route to PE-1:

```
# on PE-2:
21 2022/09/06 10:46:07.056 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.1
"Peer 1: 192.0.2.1: UPDATE
Peer 1: 192.0.2.1 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 115
    Flag: 0x90 Type: 14 Len: 45 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
        Type: EVPN-IP-PREFIX Len: 34 RD: 192.0.2.2:10, ESI: ESI-0, tag: 0, ip_prefix:
 172.31.6.0/24 gw_ip 0.0.0.0 Label: 8388528 (Raw Label: 0x7fffb0)
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0x80 Type: 9 Len: 4 Originator ID: 192.0.2.6
    Flag: 0x80 Type: 10 Len: 4 Cluster ID:
        192.0.2.5
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64496:10
        bgp-tunnel-encap:MPLS
    Flag: 0xc0 Type: 36 Len: 16 D-PATH:[64496:1020:(ipvpn)][64496:1030:(evpn)]
"
```

The figure Figure 73: VPRN BGP routes for prefix 172.31.6.0/24 shows the D-path attribute in the BGP routes for prefix 172.31.6.0/24:

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 73: VPRN BGP routes for prefix 172.31.6.0/24

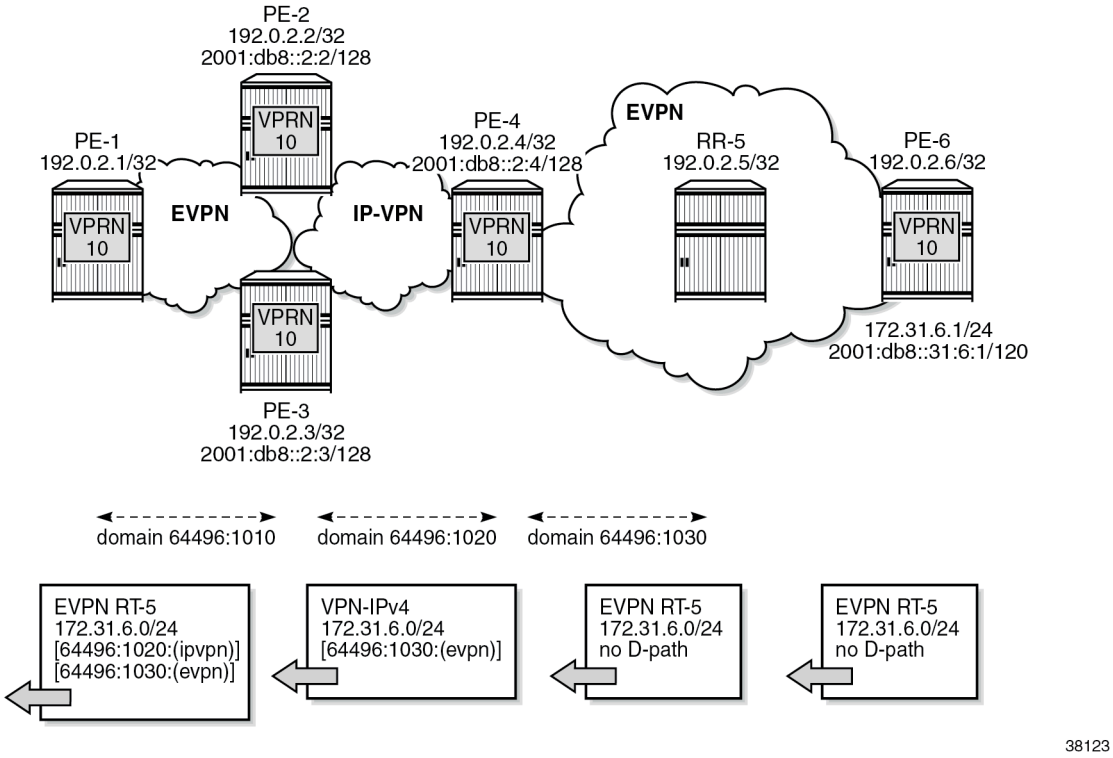

The figure Figure 74: VPRN BGP routes for prefix 172.31.7.0/24 similarly shows the D-path attribute in the BGP routes for prefix 172.31.7.0/24:

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 74: VPRN BGP routes for prefix 172.31.7.0/24

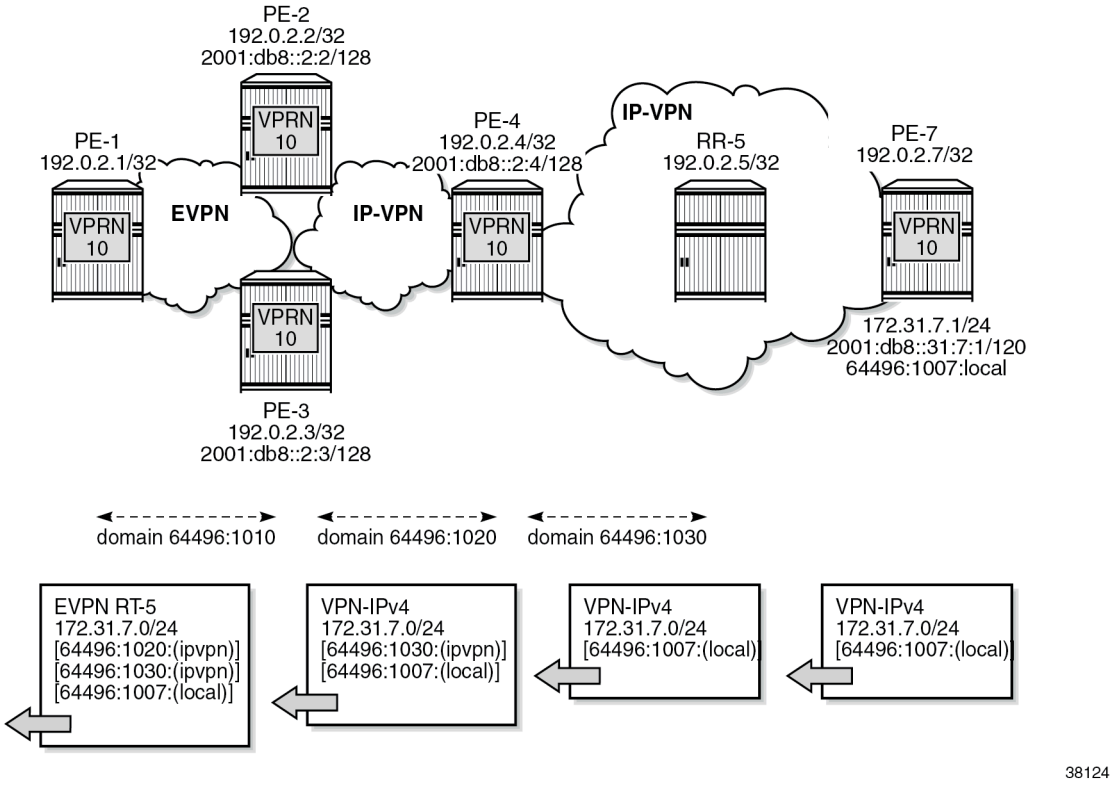

In VPRN 10 on PE-6, no local domain ID is configured, whereas in VPRN 10 on PE-7, the local domain ID 64496:1007 is configured for the routes local to PE-7.

The following BGP update shows that PE-7 advertises prefix 172.31.7.0/24 as a VPN-IPv4 route with a D-path attribute containing the domain segment 64496:1007:(local).

```
# on PE-7:
5 2022/09/06 10:46:12.896 CEST MINOR: DEBUG #2001 Base Peer 1: 192.0.2.5
"Peer 1: 192.0.2.5: UPDATE
Peer 1: 192.0.2.5 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 72
    Flag: 0x90 Type: 14 Len: 32 Multiprotocol Reachable NLRI:
        Address Family VPN_IPV4
        NextHop len 12 NextHop 192.0.2.7
        172.31.7.0/24 RD 192.0.2.7:10 Label 524282 (Raw label 0x7fffa1)
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 8 Extended Community:
        target:64496:10
    Flag: 0xc0 Type: 36 Len: 8 D-PATH:[64496:1007:(local)]
"
```

RR-5 advertises prefix 172.31.7.0/24 as a VPN-IPv4 route with the same D-path attribute. PE-4 prepends the domain segment 64496:1030:(ipvpn) to the D-path attribute of the VPN-IPv4 routes for prefix

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

172.31.7.0/24 to PE-2 (and PE-3). PE-2 advertises prefix 172.31.7.0/24 as an EVPN-IFL route to PE-1 with domain segment 64496:1020:(ipvpn) added to the D-path attribute:

```
# on PE-2:
31 2022/09/06 10:46:12.900 CEST MINOR: DEBUG #2001 Base Peer 1: 192.0.2.1
"Peer 1: 192.0.2.1: UPDATE
Peer 1: 192.0.2.1 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 123
    Flag: 0x90 Type: 14 Len: 45 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
        Type: EVPN-IP-PREFIX Len: 34 RD: 192.0.2.2:10, ESI: ESI-0, tag: 0, ip_prefix:
 172.31.7.0/24 gw_ip 0.0.0.0 Label: 8388528 (Raw Label: 0x7fffb0)
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0x80 Type: 9 Len: 4 Originator ID: 192.0.2.7
    Flag: 0x80 Type: 10 Len: 4 Cluster ID:
        192.0.2.5
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64496:10
        bgp-tunnel-encap:MPLS
    Flag: 0xc0 Type: 36 Len: 24 D-PATH:[64496:1020:(ipvpn)][64496:1030:(ipvpn)][64496:1007:
(local)]
"
```

Loop prevention

Besides traceability, the D-path attribute provides loop prevention in the control plane. Redundant GWs PE-2 and PE-3 cause routing loops and the D-path attribute helps preventing these loops. When PE-2 receives the EVPN-IFL route from PE-3 with a D-path containing domain IDs configured on PE-2, such as 64496:1020, it does not install the route in the VPRN route table, as shown in the figure Figure 75: Loop prevention between PE-2 and PE-3:

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 75: Loop prevention between PE-2 and PE-3

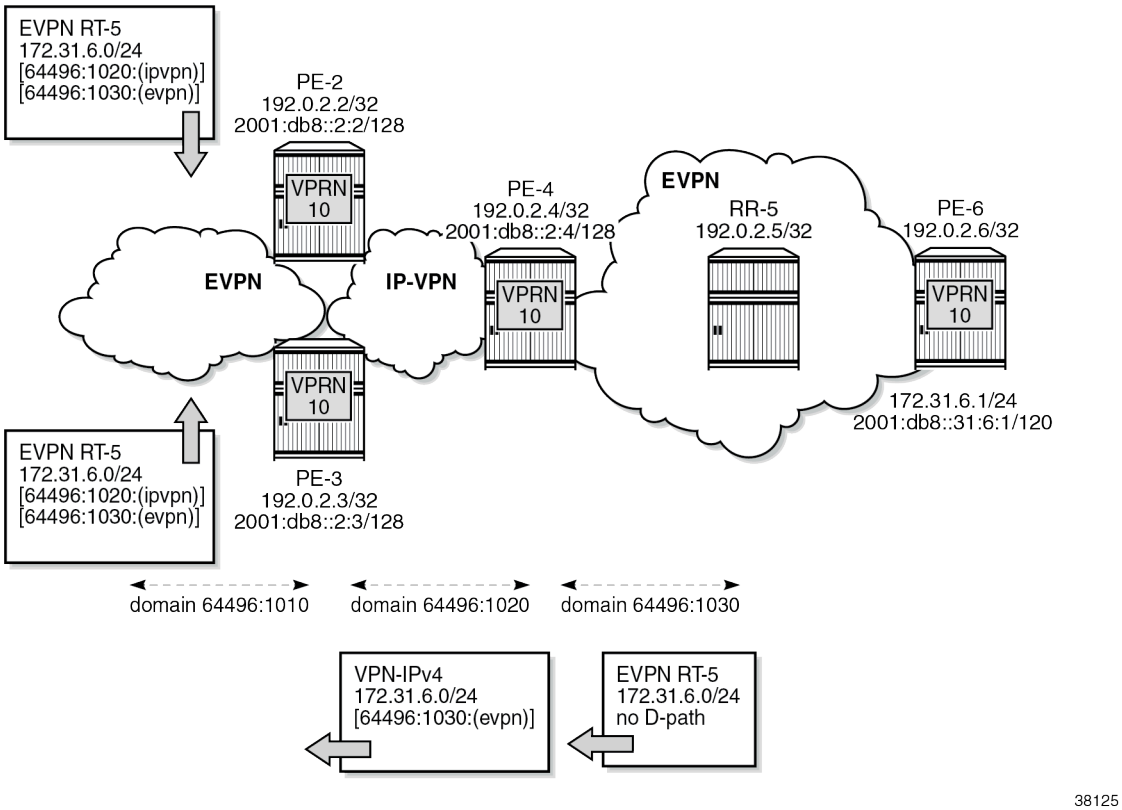

The following command on PE-2 shows that in the EVPN-IFL route for prefix 172.31.6.0/24 that was received from PE-3, a D-path loop has been detected in VPRN 10:

```
[/]
A:admin@PE-2# show router bgp routes evpn ip-prefix prefix 172.31.6.0/24 hunt
===============================================================================
 BGP Router ID:192.0.2.2        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN IP-Prefix Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : n/a
Nexthop        : 192.0.2.3
Path Id        : None
From           : 192.0.2.3
Res. Nexthop   : 192.168.23.2
Local Pref.    : 100                    Interface Name : int-PE-2-PE-3
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : 10
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Connector      : None
Community      : target:64496:10 bgp-tunnel-encap:MPLS
Cluster        : 192.0.2.5
Originator Id  : 192.0.2.6              Peer Router Id : 192.0.2.3
Flags          : Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
D-Path         : [64496:1020:(ipvpn)][64496:1030:(evpn)]
EVPN type      : IP-PREFIX
ESI            : ESI-0
Tag            : 0
Gateway Address: 00:00:00:00:00:00
Prefix         : 172.31.6.0/24
Route Dist.    : 192.0.2.3:10
MPLS Label     : LABEL 524283
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
Add Paths Send : Default
Last Modified  : 00h11m56s
DPath Loop VRFs:  10
---snip---
```

The preceding EVPN-IFL route from PE-3 for prefix 172.31.6.0/24 is not installed in the VPRN route table and is not forwarded to other PEs. The route table for VPRN 10 on PE-2 only has an IP-VPN route for prefix 172.31.6.0/24 with next hop PE-4:

```
[/]
A:admin@PE-2# show router 10 route-table

===============================================================================
Route Table (Service: 10)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
172.31.1.0/24                                 Remote  EVPN-IFL  00h12m46s  170
       192.0.2.1 (tunneled:SR-OSPF:524290)                          10
    172.31.6.0/24                                 Remote  BGP VPN   00h12m30s  170
       2001:db8:aaaa:104:7fff:9000:: (tunneled:SRV6)                20
172.31.7.0/24                                 Remote  BGP VPN   00h12m24s  170
       2001:db8:aaaa:104:7fff:9000:: (tunneled:SRV6)                20
-------------------------------------------------------------------------------
No. of Routes: 3
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

Domain IDs in R-VPLS BGP-EVPN MPLS and BGP-EVPN VXLAN instances

Loops can also be prevented in Layer 3 EVPN data center gateway (DC GW) scenarios where EVPN-IFF routes are translated into IP-VPN routes, and vice versa. Because redundant GWs are used, the scenario is subject to Layer 3 routing loops and the D-path attribute helps preventing these loops without the need for extra routing policies to tag or drop routes. The figure Figure 76: Example topology with R-VPLS shows a slightly modified example topology with R-VPLS with PE-2 and PE-3 acting as redundant DC GWs. PE-1 advertises an EVPN-IFF route for prefix 10.20.201.0/24 and PE-6 advertises an EVPN-IFF route for prefix 10.20.206.0/24.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 76: Example topology with R-VPLS

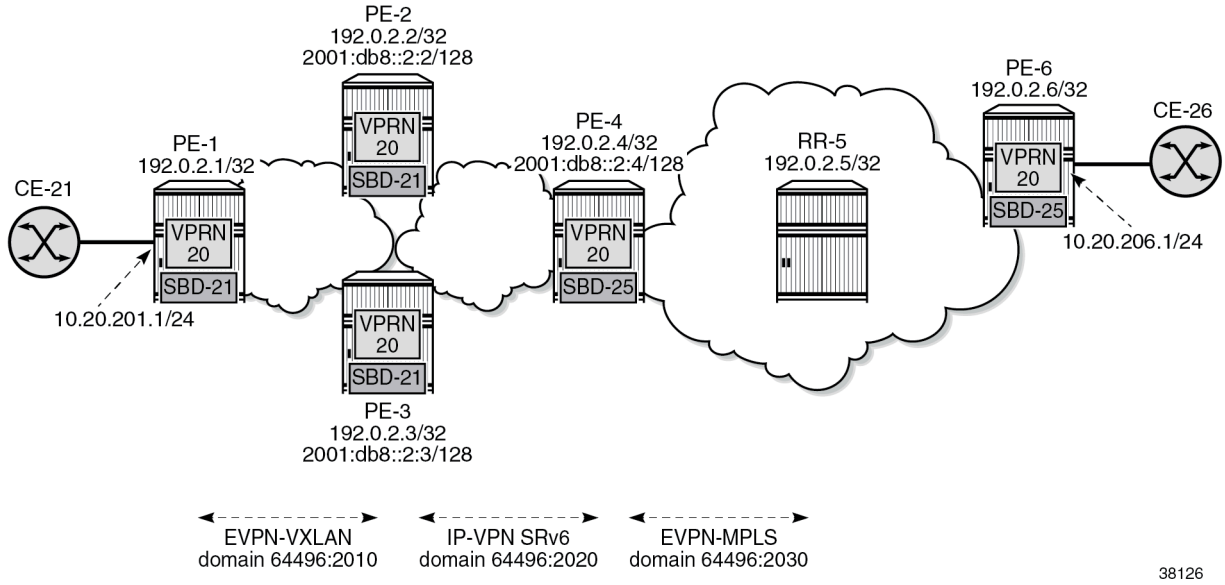

The service configuration on PE-1 does not include a domain ID, as follows:

```
# on PE-1:
configure {
    service {
        vpls "SBD-21" {
            admin-state enable
            service-id 21
            customer "1"
            vxlan {
                instance 1 {
                    vni 1
                }
            }
            routed-vpls {
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 21
                routes {
                    ip-prefix {
                        advertise true
                    }
                }
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
        }
        vprn "VPRN 20" {
            admin-state enable
            service-id 20
            customer "1"
            autonomous-system 64496
            interface "int-PE-1-CE-21" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                ipv4 {
                    primary {
                        address 10.20.201.1
                        prefix-length 24
                    }
                }
                sap 1/1/c5/1:20 {
                }
            }
            interface "int-SBD-21" {
                vpls "SBD-21" {
                    evpn-tunnel {
                    }
                }
            }
        }
```

On DC GW PE-2, domain ID 64496:2010 is configured in VPLS "SBD-21" whereas domain ID 64496:2020 is configured in VPRN 20. The configuration on DC GW PE-3 is similar.

```
# on PE-2:
configure {
    service {
        vpls "SBD-21" {
            admin-state enable
            service-id 21
            customer "1"
            vxlan {
                instance 1 {
                    vni 1
                }
            }
            routed-vpls {
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 21
                routes {
                    ip-prefix {
                        advertise true
                        domain-id "64496:2010"
                    }
                }
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
        }
        vprn "VPRN 20" {
            admin-state enable
            service-id 20
            customer "1"
            autonomous-system 64496
            segment-routing-v6 1 {
                locator "PE-2_loc" {                # on PE-3: "PE-3_loc"
                    function {
                        end-dt46 {
                        }
                    }
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            bgp-ipvpn {
                segment-routing-v6 1 {
                    admin-state enable
                    route-distinguisher "192.0.2.2:26"  # on PE-3; 192.0.2.3:26
                    source-address 2001:db8::2:2        # on PE-3: 2001:db8::2:3
                    domain-id "64496:2020"
                    vrf-target {
                        community "target:64496:20"
                    }
                    srv6 {
                        instance 1
                        default-locator "PE-2_loc"     # on PE-3: "PE-3_loc"
                    }
                }
            }
            interface "int-SBD-21" {
                vpls "SBD-21" {
                    evpn-tunnel {
                    }
                }
            }
        }
```

The service configuration examples for PE-1, PE-2, and PE-3 show how a loop is detected at the DC GWs in VPN-IPv4 routes for prefix 10.20.201.0/24 received from the other DC GW. The following command on DC GW PE-2 shows that a D-path loop is detected in VPRN 20 in a VPN-IPv4 route for prefix 10.20.201.0/24 received from DC GW PE-3:

```
[/]
A:admin@PE-2# show router bgp routes vpn-ipv4 rd 192.0.2.3:26 hunt
===============================================================================
 BGP Router ID:192.0.2.2        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv4 Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : 10.20.201.0/24
Nexthop        : 2001:db8::2:3
Route Dist.    : 192.0.2.3:26           VPN Label      : 524283
Path Id        : None
From           : 2001:db8::2:3
Res. Nexthop   : n/a
Local Pref.    : 100                    Interface Name : int-PE-2-PE-3
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : 10
Connector      : None
Community      : target:64496:20
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.3
Fwd Class      : None                   Priority       : None
Flags          : Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
D-Path         : [64496:2010:(evpn)]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Route Tag       : 0
Neighbor-AS     : n/a
Orig Validation: N/A
Source Class    : 0                       Dest Class     : 0
Add Paths Send : Default
Last Modified   : 00h00m51s
SRv6 TLV Type   : SRv6 L3 Service TLV (5)
SRv6 SubTLV     : SRv6 SID Information (1)
Sid             : 2001:db8:aaaa:103::
Full Sid        : 2001:db8:aaaa:103:7fff:b000::
Behavior        : End.DT46 (20)
SRv6 SubSubTLV : SRv6 SID Structure (1)
Loc-Block-Len   : 48                      Loc-Node-Len   : 16
Func-Len        : 20                      Arg-Len        : 0
Tpose-Len       : 20                      Tpose-offset   : 64
VPRN Imported   : None
DPath Loop VRFs:  20

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Routes : 1
===============================================================================*A:PE-2# show
 router bgp routes vpn-ipv4 rd 192.0.2.3:26 hunt
```

The figure Figure 77: Loop prevention between DC GW PE-2 and DC GW PE-3 shows that PE-1 sends an EVPN-IFF route for prefix 10.20.201.0/24 without D-path attribute to PE-2 and PE-3. Both PE-2 and PE-3 re-advertise prefix 10.20.201.0/24 as a VPN-IPv4 route with D-path attribute 64496:2010:(evpn). When PE-2 receives this VPN-IPv4 route from PE-3, it detects a loop based on the D-path attribute with domain segment 64496:2010:(evpn) and does not install the route in the VPRN route table. Likewise, PE-3 receives the VPN-IPv4 route from PE-2 and does not install it in the VPRN route table.

Figure 77: Loop prevention between DC GW PE-2 and DC GW PE-3

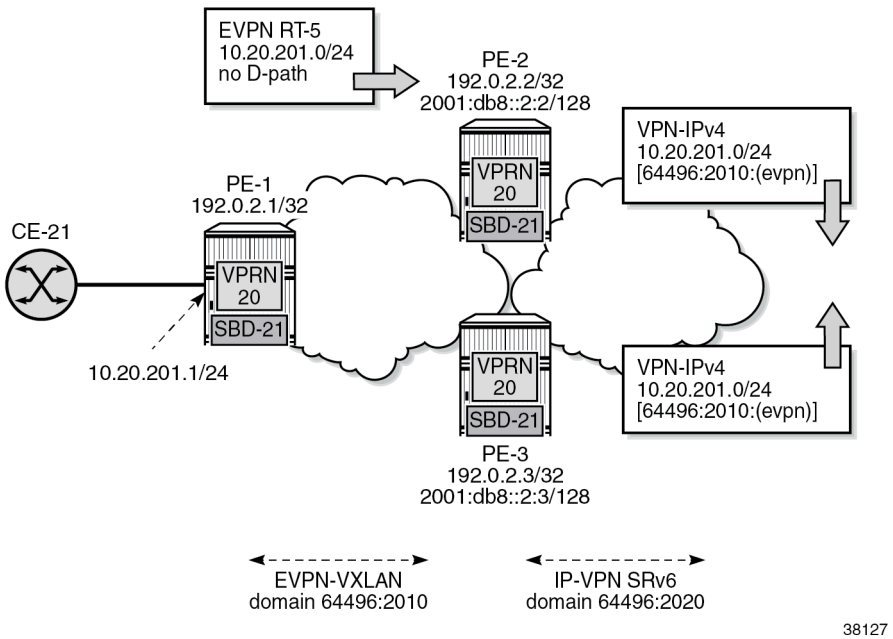

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

PE-2 does not use the VPN-IPv4 route for prefix 10.20.201.0/24 from PE-3. The VPRN route table on PE-2 contains the EVPN-IFF route received from PE-1 for prefix 10.20.201.0/24:

```
[/]
A:admin@PE-2# show router 20 route-table

===============================================================================
Route Table (Service: 20)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
10.20.201.0/24                                Remote  EVPN-IFF  00h01m59s  169
       int-SBD-21 (ET-02:0f:ff:ff:ff:52)                          0
10.20.206.0/24                                Remote  BGP VPN   00h01m43s  170
       2001:db8:aaaa:104:7fff:6000:: (tunneled:SRV6)               20
-------------------------------------------------------------------------------
No. of Routes: 2
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

Conclusion

The D-path attribute provides traceability for VPRN BGP routes and can be used for BGP best path selection. The D-path attribute for VPRN routes also helps preventing loops without the need for dedicated routing policies to tag and drop routes.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Dual EVPN-MPLS Instance VPLS Services

This chapter provides information about the dual EVPN-MPLS instance VPLS services.

Topics in this chapter include:

Applicability

The information and MD-CLI configuration in this chapter are based on SR OS Release 22.10.R1. Dual EVPN-MPLS instance in VPLS is supported in SR OS Release 21.10.R1 and later.

Overview

One of the scale issues that low-scale access nodes or leaf PEs face in high-scale architectures is the limited number of EVPN/IP-VPN next hops, tunnels, and service labels that they support.

The following solutions reduce the number of EVPN next hops exposed to the access nodes:

- inter-AS model B, as described in the Inter-AS VPRN Model B chapter
- next-hop-self route reflectors (NHS-RRs)

The figure Figure 78: Access nodes receive next hops from the NHS-RRs shows the NHS-RR solution reducing the number of EVPN next hops that are sent to the low-scale access nodes PE-1 and PE-2. Only the two NHS-RRs are exposed as next hops to PE-1.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 78: Access nodes receive next hops from the NHS-RRs

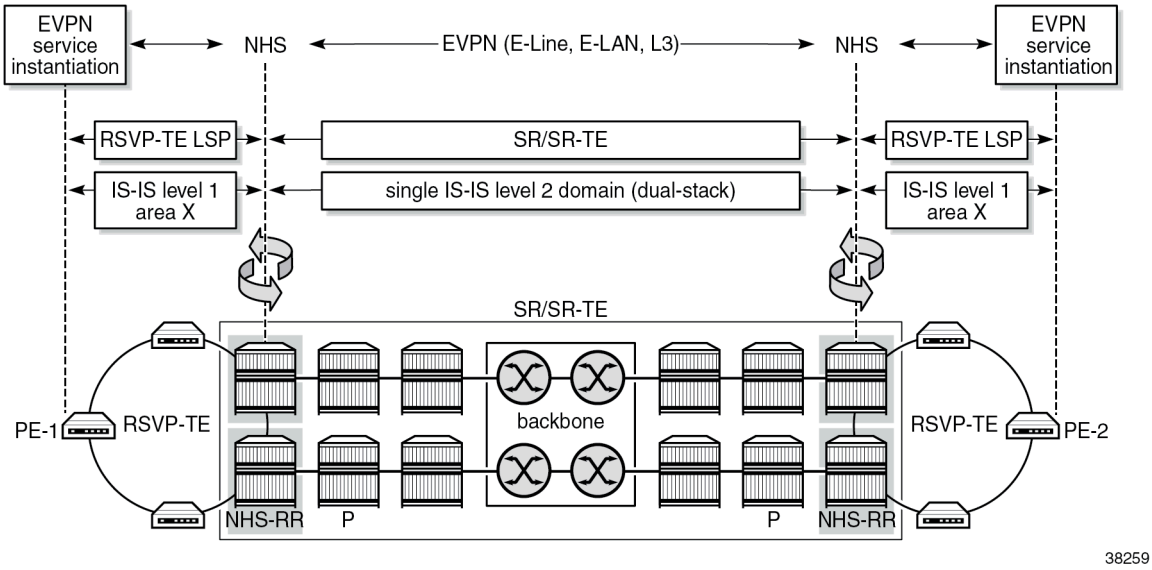

The number of EVPN next hops is reduced, but the number of service labels to be learned is not. PE-1 still learns one service label per remote PE for each service it is attached to. In case of EVPN E-LAN services and broadcast, unknown unicast, and multicast (BUM) traffic, the ingress PE still needs one copy of every BUM packet per egress PE that exists in the remote domains, even if all the BUM traffic goes through one of the two NHS-RRs (or ASBRs in the case of model B).

The following solutions reduce the number of service labels:

- VPRN services on the NHS-RRs with **allow-export-bgp-vpn** configured
- dual EVPN-MPLS instance VPLS services on the NHS-RRs

The **allow-export-bgp-vpn** command applies to VPRN services using EVPN-IFL, VPN-IPv4, and VPN-IPv6 families. Routes from the WAN are imported to the VPRN service and exported to the access nodes as new VPN-IP routes. The values of the service labels, route targets (RTs), and BGP next hops of the re-advertised routes are based on the configuration of the exporting VPRN.

The figure Figure 79: Access nodes receive one service label per service from each NHS-RR shows a dual EVPN-MPLS instance VPLS service on the NHS-RRs, which offers a similar solution for EVPN-VPLS services to the **allow-export-bgp-vpn** solution for VPRN services. EVPN-MPLS routes received from the WAN are imported to the network EVPN-MPLS instance and redistributed to the access EVPN-MPLS instance with a new route distinguisher (RD), next hop, service label, and possibly a new RT. The ingress PE learns only one service label for each NHS-RR per service, as opposed to one service label per remote PE that is attached to the same EVPN service. With this solution, the replication of BUM traffic is also optimized because the ingress PE sends a single copy of each BUM packet to the NHS-RR, as opposed to one copy per egress PE.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 79: Access nodes receive one service label per service from each NHS-RR

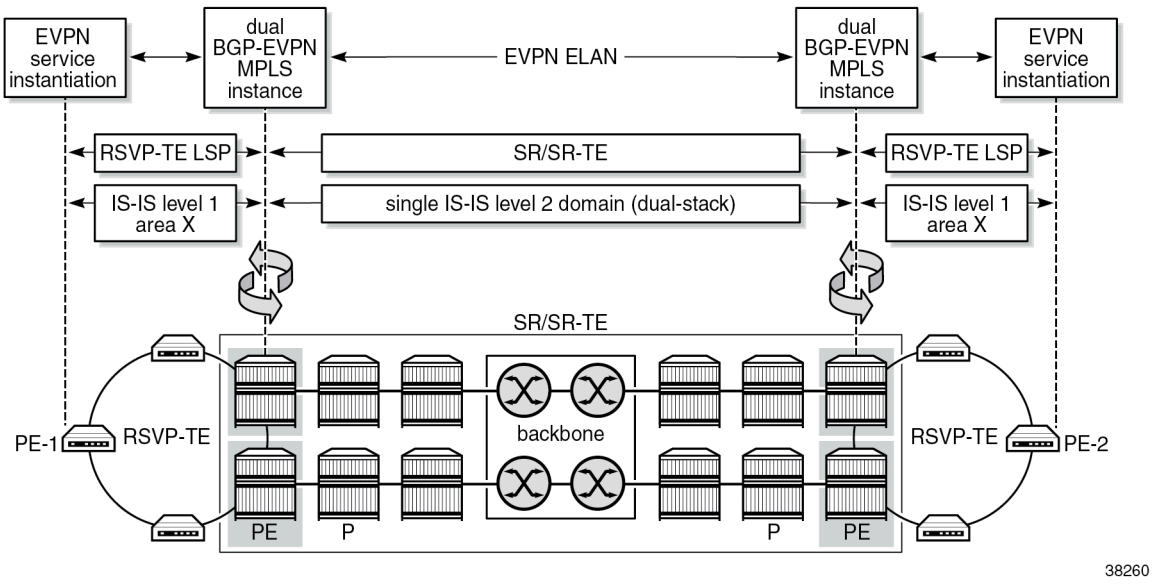

In the example, redundant NHS-RRs are used. Redundancy is handled via anycast multihoming, which implies that two or more PEs are configured with the same service parameters as part of the same redundancy group: identical route distinguishers and RTs per instance, and the same anycast IP address. The ingress PEs set up EVPN destinations to only one PE in the anycast group for a specific service. EVPN BUM destinations are not established between PEs in the same anycast group because the received anycast peer inclusive multicast Ethernet tag (IMET) routes have the same local originating IP address. In anycast multihoming scenarios, policies are required to prevent control-plane loops.

Configuration

The following scenarios are described in this section:

- dual EVPN-MPLS instance VPLS without multihoming
- dual EPVN-MPLS instance VPLS with anycast multihoming

Dual EVPN-MPLS instance VPLS without multihoming

The figure Figure 80: Example topology 1 shows EVPN-MPLS VPLS 1 configured on four PEs. PE-2 and PE-4 are EVPN gateways (GWs). RR-7 is the route reflector for PE-2 and PE-4 in the WAN network.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 80: Example topology 1

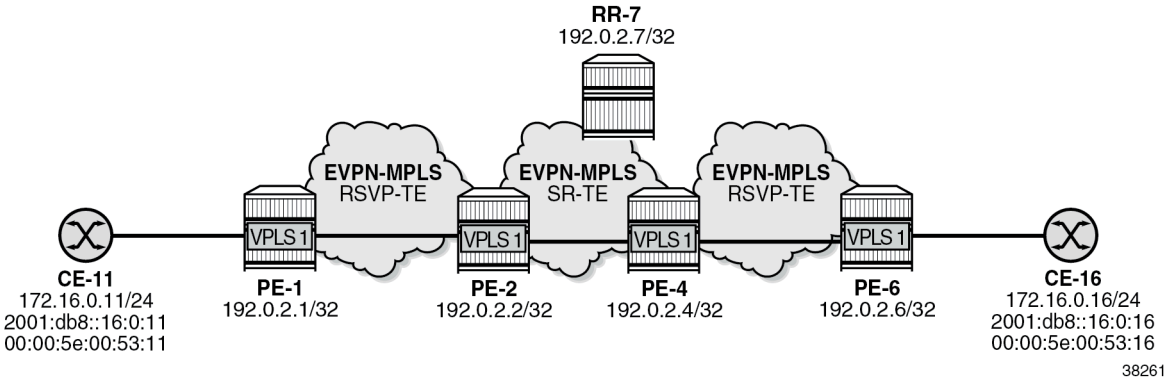

The initial configuration includes:

- cards, MDAs, ports
- router interfaces
- IS-IS level 1 between PE-1 and PE-2 and between PE-4 and PE-6
- IS-IS level 2 between PE-2, PE-4, and RR-7
- SR-TE tunnels between PE-2 and PE-4
- MPLS LSPs between PE-1 and PE-2 and between PE-4 and PE-6

BGP is configured on all nodes for the EVPN address family. PE-1 peers with the dual-homed EVPN GW PE-2. In a similar way, PE-6 peers with EVPN GW PE-4. The BGP configuration on PE-1 is as follows:

```
# on PE-1:
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            split-horizon true
            rapid-update {
                evpn true
            }
            group "access1" {
                peer-as 64496
                family {
                    evpn true
                }
            }
            neighbor "192.0.2.2" {
                group "access1"
            }
        }
```

EVPN GW PE-2 peers with PE-1 in BGP group "access1" and with RR-7 in BGP group "WAN":

```
# on PE-2:
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            split-horizon true
            rapid-update {
                evpn true
            }
            group "WAN" {
                next-hop-self true
                peer-as 64496
                family {
                    evpn true
                }
                export {
                    policy ["drop-tag-10"]
                }
            }
            group "access1" {
                next-hop-self true
                peer-as 64496
                family {
                    evpn true
                }
                cluster {
                    cluster-id 192.0.2.2
                }
                export {
                    policy ["drop-tag-20"]
                }
            }
            neighbor "192.0.2.1" {
                group "access1"
            }
            neighbor "192.0.2.7" {
                group "WAN"
            }
        }
```

The BGP configuration on PE-4 is similar. The export policies use tags to avoid loops in topologies with redundant EVPN GWs, as described in the section Dual EVPN-MPLS instance VPLS with anycast multihoming.

RR-7 peers with PE-2 and PE-4 in BGP group "WAN":

```
# on RR-7:
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            split-horizon true
            rapid-update {
                evpn true
            }
            group "WAN" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                peer-as 64496
                family {
                    evpn true
                }
                cluster {
                    cluster-id 192.0.2.7
                }
            }
            neighbor "192.0.2.2" {
                group "WAN"
            }
            neighbor "192.0.2.4" {
                group "WAN"
            }
        }
```

On PE-1, VPLS 1 is configured with a single EVPN-MPLS instance. The RD 192.0.2.1:1 for BGP 1 is auto-derived from the values for the IPv4 system address and the EVI. PE-1 imports and exports routes with RT 64496:101.

```
# on PE-1:
configure {
    service {
        vpls "VPLS-1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
                # route-distinguisher 192.0.2.1:1  # will be auto-derived
                route-target {
                    export "target:64496:101"
                    import "target:64496:101"
                }
            }
            bgp-evpn {
                evi 1
                mpls 1 {
                    admin-state enable
                    auto-bind-tunnel {
                        resolution filter
                        resolution-filter {
                            rsvp true
                        }
                    }
                }
            }
            sap 1/1/c10/1:1 {
            }
        }
```

On PE-2, VPLS 1 is configured with two EVPN-MPLS instances: instance 1 is configured with multihoming mode access and instance 2 with the (default) multihoming mode network, as follows:

```
# on PE-2:
configure {
    service {
        system {
            bgp-auto-rd-range {
                ip-address 192.0.2.2
                community-value {
                    start 2000
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    end 2999
                }
            }
        }
        vpls "VPLS-1" {
            admin-state enable
            description "dual BGP-EVPN MPLS instance VPLS"
            service-id 1
            customer "1"
            bgp 1 {
                # route-distinguisher 192.0.2.2:1    # will be auto-derived
                route-target {
                    export "target:64496:101"
                    import "target:64496:101"
                }
            }
            bgp 2 {
                route-distinguisher auto-rd
                route-target {
                    export "target:64496:100"
                    import "target:64496:100"
                }
            }
            bgp-evpn {
                evi 1
                mpls 1 {
                    admin-state enable
                    mh-mode access
                    auto-bind-tunnel {
                        resolution filter
                        resolution-filter {
                            rsvp true
                        }
                    }
                }
                mpls 2 {
                    admin-state enable
                    # mh-mode network               # default MH mode
                    auto-bind-tunnel {
                        resolution filter
                        resolution-filter {
                            sr-te true
                        }
                    }
                }
            }
        }
```

Note: The RD for BGP 1 can be auto-derived from the values for the IPv4 system address and the EVI, for example, 192.0.2.2:1 on PE-2. The RD for BGP 2 cannot be auto-derived from the values for the IPv4 system address and the EVI, because the RD for BGP 2 must be different from the RD for BGP 1, so it must be configured manually or with **auto-rd**.

On PE-4, the configuration is similar:

```
# on PE-4:
configure {
    service {
        system {
            bgp-auto-rd-range {
                ip-address 192.0.2.4
                community-value {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    start 2000
                    end 2999
                }
            }
        }
        vpls "VPLS-1" {
            admin-state enable
            description "dual BGP-EVPN MPLS instance VPLS"
            service-id 1
            customer "1"
            bgp 1 {
                # route-distinguisher 192.0.2.4:1    # will be auto-derived
                route-target {
                    export "target:64496:102"
                    import "target:64496:102"
                }
            }
            bgp 2 {
                route-distinguisher auto-rd     # different RD
                route-target {
                    export "target:64496:100"
                    import "target:64496:100"
                }
            }
            bgp-evpn {
                evi 1
                mpls 1 {
                    admin-state enable
                    mh-mode access
                    auto-bind-tunnel {
                        resolution filter
                        resolution-filter {
                            rsvp true
                        }
                    }
                }
                mpls 2 {
                    admin-state enable
                    # mh-mode network               # default MH mode
                    auto-bind-tunnel {
                        resolution filter
                        resolution-filter {
                            sr-te true
                        }
                    }
                }
            }
        }
```

The following command on PE-2 shows BGP instances 1 and 2 in VPLS 1. RD 192.0.2.2:1 for BGP instance 1 is auto-derived from the IPv4 system address and the EVI; the RD for BGP instance 2 is configured with **auto-rd** and has the value 192.0.2.2:2000. The RT values are configured.

```
[/]
A:admin@PE-2# show service id 1 bgp

===============================================================================
BGP Information
===============================================================================
Bgp Instance         : 1
Vsi-Import           : None
Vsi-Export           : None
Route Dist           : None
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Oper Route Dist      : 192.0.2.2:1
Oper RD Type         : derivedEvi
Rte-Target Import    : 64496:101           Rte-Target Export: 64496:101
Oper RT Imp Origin   : configured          Oper RT Import   : 64496:101
Oper RT Exp Origin   : configured          Oper RT Export   : 64496:101
ADV Service MTU      : -1

Bgp Instance         : 2
Vsi-Import           : None
Vsi-Export           : None
Route Dist           : auto-rd
Oper Route Dist      : 192.0.2.2:2000
Oper RD Type         : auto
Rte-Target Import    : 64496:100           Rte-Target Export: 64496:100
Oper RT Imp Origin   : configured          Oper RT Import   : 64496:100
Oper RT Exp Origin   : configured          Oper RT Export   : 64496:100
ADV Service MTU      : -1

PW-Template Id       : None
-------------------------------------------------------------------------------
===============================================================================
```

The following command on PE-2 shows EVPN destination 192.0.2.1 in EVPN-MPLS instance 1:

```
[/]
A:admin@PE-2# show service id 1 evpn-mpls instance 1

===============================================================================
BGP EVPN-MPLS Dest
===============================================================================
TEP Address                   Egr Label     Num.   Mcast Last Change
                              Transport:Tnl MACs         Sup BCast Domain
-------------------------------------------------------------------------------
192.0.2.1                     524286        0      bum   12/13/2022 09:56:36
                              rsvp:1                     No
-------------------------------------------------------------------------------
Number of entries : 1
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN-MPLS Ethernet Segment Dest
===============================================================================
Eth SegId                      Num. Macs              Last Change
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

The following command on PE-2 shows EVPN destination 192.0.2.4 in EVPN-MPLS instance 2:

```
[/]
A:admin@PE-2# show service id 1 evpn-mpls instance 2

===============================================================================
BGP EVPN-MPLS Dest
===============================================================================
TEP Address                   Egr Label     Num.   Mcast Last Change
                              Transport:Tnl MACs         Sup BCast Domain
-------------------------------------------------------------------------------
192.0.2.4                     524282        0      bum   12/13/2022 09:56:39
                              sr-te:655362               No
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Number of entries : 1
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN-MPLS Ethernet Segment Dest
===============================================================================
Eth SegId                       Num. Macs              Last Change
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

When traffic is sent between CE-11 and CE-16, MAC address 00:00:5e:00:53:11 of CE-11 is learned on the local SAP in VPLS 1 on PE-1 and MAC address 00:00:5e:00:53:16 of CE-16 is learned on the local SAP in VPLS 1 on PE-6. EVPN MAC routes are advertised to the BGP-EVPN peers.

The forwarding database (FDB) on PE-1 is as follows:

```
[/]
A:admin@PE-1# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
1          00:00:5e:00:53:11 sap:1/1/c10/1:1         L/0      12/13/22 10:04:14
1          00:00:5e:00:53:16 mpls-1:                 Evpn     12/13/22 10:04:14
                             192.0.2.2:524284
           rsvp:1
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

The FDB on PE-2 shows that an EVPN MAC route is received in EVPN-MPLS instance 1 for address 00:00:5e:00:53:11 whereas an EVPN MAC route is received in EVPN-MPLS instance 2 for address 00:00:5e:00:53:16.

```
[/]
A:admin@PE-2# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
1          00:00:5e:00:53:11 mpls-1:                 Evpn     12/13/22 10:04:14
                             192.0.2.1:524286
           rsvp:1
1          00:00:5e:00:53:16 mpls-2:                 Evpn     12/13/22 10:04:14
                             192.0.2.4:524282
           sr-te:655362
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
```

The following command shows the received EVPN-MAC routes on PE-2 for MAC address 00:00:5e:00:53:16. The route with RD 192.0.2.4:2000 is used:

```
[/]
A:admin@PE-2# show router bgp routes evpn mac mac-address 00:00:5e:00:53:16
===============================================================================
 BGP Router ID:192.0.2.2        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN MAC Routes
===============================================================================
Flag  Route Dist.          MacAddr           ESI
      Tag                  Mac Mobility      Label1
                           Ip Address
                           NextHop
-------------------------------------------------------------------------------
u*>i  192.0.2.4:2000       00:00:5e:00:53:16 ESI-0
      0                    Seq:0             LABEL 524282
                           n/a
                           192.0.2.4

*>i   192.0.2.6:1          00:00:5e:00:53:16 ESI-0
      0                    Seq:0             LABEL 524286
                           n/a
                           192.0.2.6

-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

Dual EVPN-MPLS instance VPLS with anycast multihoming

Figure 81: Example topology 2 shows example topology 2 with VPLS 2 configured on six PEs. PE-2 and PE-3 are redundant EVPN GWs with anycast address 23.23.23.23; PE-4 and PE-5 are redundant EVPN GWs with anycast address 45.45.45.45. RR-7 is the route reflector for PE-2, PE-3, PE-4, and PE-5 in the WAN network.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 81: Example topology 2

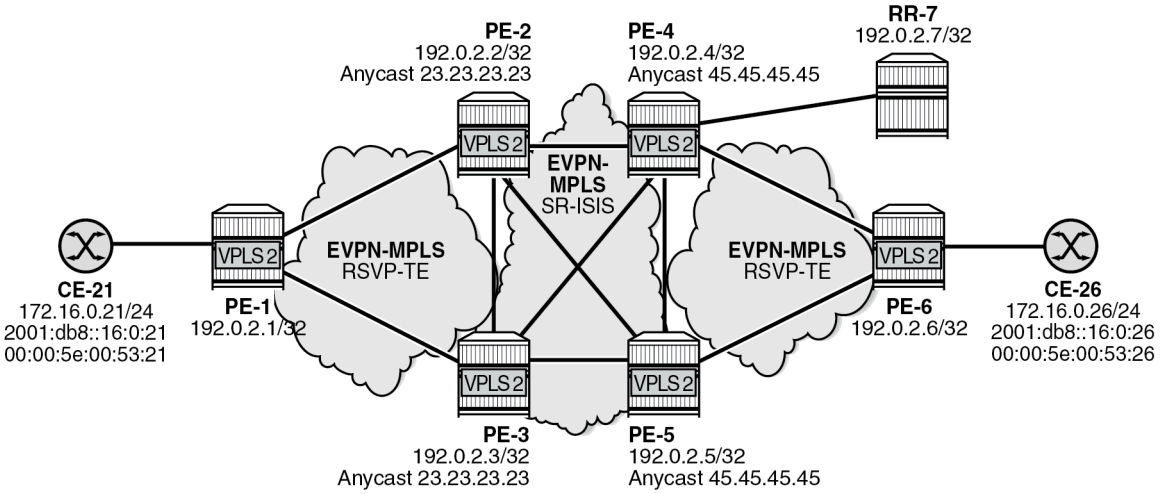

The initial configuration includes:

- cards, MDAs, ports
- router interfaces
- IS-IS level 1 between PE-1, PE-2, and PE-3
- IS-IS level 1 between PE-4, PE-5, and PE-6
- IS-IS level 2 between PE-2, PE-3, PE-4, PE-5, and RR-7
- SR-ISIS between PE-2, PE-3, PE-4, and PE-5
- MPLS LSPs between PE-1 and PE-2, between PE-1 and PE-3, between PE-4 and PE-6, and between PE-5 and PE-6

The BGP configuration on PE-1 and PE-6 is similar.

```
# on PE-1:
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            split-horizon true
            rapid-update {
                evpn true
            }
            group "access1" {
                peer-as 64496
                family {
                    evpn true
                }
            }
            neighbor "192.0.2.2" {        # on PE-6: 192.0.2.4
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                group "access1"
            }
            neighbor "192.0.2.3" {        # on PE-6: 192.0.2.5
                group "access1"
            }
        }
```

The BGP configuration on PE-3 is:

```
# on PE-3:
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            split-horizon true
            rapid-update {
                evpn true
            }
            group "WAN" {
                next-hop-self true
                peer-as 64496
                family {
                    evpn true
                }
                export {
                    policy ["drop-tag-10"]
                }
            }
            group "access1" {
                next-hop-self true
                peer-as 64496
                family {
                    evpn true
                }
                cluster {
                    cluster-id 192.0.2.3
                }
                export {
                    policy ["drop-tag-20"]
                }
            }
            neighbor "192.0.2.1" {
                group "access1"
            }
            neighbor "192.0.2.7" {
                group "WAN"
            }
        }
```

The BGP configuration on PE-2, PE-4, and PE-5 is similar.

On PE-1, VPLS 2 is configured with a single EVPN-MPLS instance. PE-1 imports and exports routes with RT 64496:501. The configuration is as follows:

```
# on PE-1:
configure {
    service {
        vpls "VPLS-2" {
            admin-state enable
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            service-id 2
            customer "1"
            bgp 1 {
                # route-distinguisher 192.0.2.1:2  # will be auto-derived
                route-target {
                    export "target:64496:501"
                    import "target:64496:501"
                }
            }
            bgp-evpn {
                evi 2
                mpls 1 {
                    admin-state enable
                    auto-bind-tunnel {
                        resolution filter
                        resolution-filter {
                            rsvp true
                        }
                    }
                }
            }
            sap 1/1/c10/1:2 {
            }
        }
```

On PE-2 and PE-3, the following policies are used in VPLS 2:

- Export policy "vsi-501-export" adds the communities "SOO-23" for the site of origin (SOO) and "RT64496:501" for the RT.
- Export policy "vsi-502-export" adds the communities "SOO-23" and "RT64496:502".
- Import policy "vsi-501-import" prevents loops based on the SOO and accepts routes with RT 64496:501.
- Import policy "vsi-502-import" prevent loops based on the SOO and accepts routes with RT 64496:502.

```
# on PE-2, PE-3:
configure {
    policy-options {
        community "RT64496:501" {
            member "target:64496:501" { }
        }
        community "RT64496:502" {
            member "target:64496:502" { }
        }
        community "SOO-23" {
            member "origin:23:23" { }
        }
        policy-statement "vsi-501-export" {
            default-action {
                action-type accept
                community {
                    add ["RT64496:501" "SOO-23"]
                }
            }
        }
        policy-statement "vsi-501-import" {
            entry 10 {
                from {
                    family [evpn]
                    community {
                        name "SOO-23"
                    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
                action {
                    action-type reject
                }
            }
            entry 20 {
                from {
                    family [evpn]
                    community {
                        name "RT64496:501"
                    }
                }
                action {
                    action-type accept
                }
            }
        }
        policy-statement "vsi-502-export" {
            default-action {
                action-type accept
                community {
                    add ["RT64496:502" "SOO-23"]
                }
            }
        }
        policy-statement "vsi-502-import" {
            entry 10 {
                from {
                    family [evpn]
                    community {
                        name "SOO-23"
                    }
                }
                action {
                    action-type reject
                }
            }
            entry 20 {
                from {
                    family [evpn]
                    community {
                        name "RT64496:502"
                    }
                }
                action {
                    action-type accept
                }
            }
        }
```

On PE-2 and PE-3, VPLS 2 is configured with two EVPN-MPLS instances: instance 1 is configured with multihoming mode access and instance 2 with multihoming mode network. For redundancy, anycast multihoming is configured with anycast address 23.23.23.23 and identical RDs and RTs for the same instance. The RD for BGP 1 is 192.0.2.23:2 and the RD for BGP 2 is 192.0.2.32:2. The **default-route-tag 10** command is configured for service instance 1, while **default-route-tag 20** is configured for service instance 2. These route tags are used in the BGP peer export policies to differentiate the different routes.On PE-2 and PE-3, VPLS 2 is configured as follows:

```
# on PE-2, PE-3:
configure {
    service {
        vpls "VPLS-2" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            admin-state enable
            description "dual BGP-EVPN MPLS instance VPLS"
            service-id 2
            customer "1"
            bgp 1 {
                route-distinguisher "192.0.2.23:2"
                vsi-import ["vsi-501-import"]
                vsi-export ["vsi-501-export"]
            }
            bgp 2 {
                route-distinguisher "192.0.2.32:2"
                vsi-import ["vsi-502-import"]
                vsi-export ["vsi-502-export"]
            }
            bgp-evpn {
                evi 2
                incl-mcast-orig-ip 23.23.23.23
                mpls 1 {
                    admin-state enable
                    default-route-tag 0xa       # default route tag 10
                    mh-mode access
                    auto-bind-tunnel {
                        resolution filter
                        resolution-filter {
                            rsvp true
                        }
                    }
                }
                mpls 2 {
                    admin-state enable
                    default-route-tag 0x14       # default route tag 20
                    auto-bind-tunnel {
                        resolution filter
                        resolution-filter {
                            sr-isis true
                        }
                    }
                }
            }
        }
```

Note: For anycast multihoming, the RDs must be identical, so all RDs are configured manually.

In datacenter GWs (DC GWs) with EVPN-VXLAN and EVPN-MPLS instances, route policies can match on the encapsulation type VXLAN or MPLS. In DC GWs with two EVPN-MPLS instances, the default route tag is used instead. The default route tag prevents a MAC/IP route that is installed in instance 1 (access) from being readvertised back to the access peers. In a similar way, MAC/IP routes installed in instance 2 are not readvertised back to peers in instance 2. On PE-2 and PE-3, the BGP peer export policy "drop-tag-10" drops routes with tag 10 and is configured in BGP group "WAN" with neighbor RR-7; BGP peer export policy "drop-tag-20" drops routes with tag 20 and is configured in BGP group "access1" with neighbor PE-1.

```
# on PE-2, PE-3:
configure {
    policy-options {
        policy-statement "drop-tag-10" {
            description "used as export policy toward WAN BGP peers"
            entry 10 {
                from {
                    tag 10
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
                action {
                    action-type reject
                }
            }
            default-action {
                action-type accept
            }
        }
        policy-statement "drop-tag-20" {
            description "used as export policy toward DC BGP peers"
            entry 10 {
                from {
                    tag 20
                }
                action {
                    action-type reject
                }
            }
            default-action {
                action-type accept
            }
        }
        info
    }
    router "Base" {
        bgp {
            group "WAN" {
                export {
                    policy ["drop-tag-10"]
                }
            }
            group "access1" {
                export {
                    policy ["drop-tag-20"]
                }
            }
```

The figure Figure 82: Export policies on PE-2 drop routes based on tag shows an incoming EVPN MAC route on PE-2 for CE-21's MAC address 00:00:5e:00:53:21. PE-2 receives the EVPN MAC route with RT target:64496:501 from PE-1 (BGP-EVPN peer 192.0.2.1). On PE-2, BGP 1 in VPLS 2 imports routes with this RT and the MAC address is installed in the FDB. The EVPN MAC route is redistributed to BGP 2 where the communities "RT64496:502" and "SOO-23", as well as internal tag 20, are added to the route. When PE-2's BGP process sends an EVPN MAC route with tag 20 to BGP peer PE-1, the BGP export policy "drop-tag-20" drops the route, preventing PE-2 from re-advertising the EVPN MAC route back to the access peer 192.0.2.1. PE-2 can only send the EVPN MAC route to WAN neighbor 192.0.2.7 because the BGP export policy toward the WAN only drops the routes with tag 10, not the ones with tag 20.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 82: Export policies on PE-2 drop routes based on tag

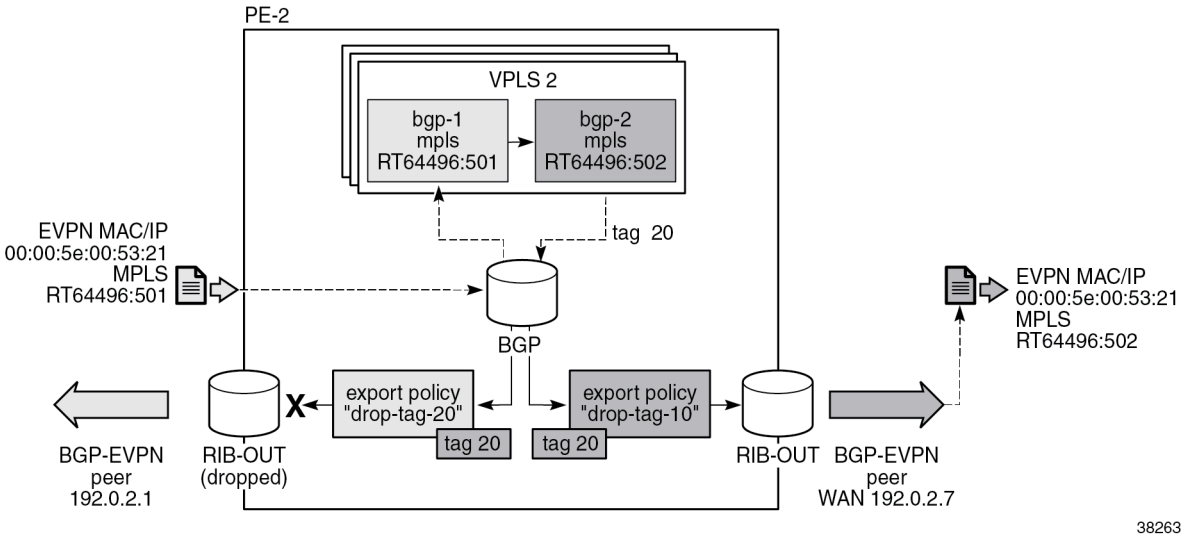

For completeness, the configuration on PE-4 and PE-5 is as follows:

```
# on PE-4, PE-5:
configure {
    policy-options {
        community "RT64496:502" {
            member "target:64496:502" { }
        }
        community "RT64496:503" {
            member "target:64496:503" { }
        }
        community "SOO-45" {
            member "origin:45:45" { }
        }
        policy-statement "drop-tag-20" {
            description "used as export policy toward DC BGP peers"
            entry 10 {
                from {
                    tag 20
                }
                action {
                    action-type reject
                }
            }
            default-action {
                action-type accept
            }
        }
        policy-statement "drop-tag-30" {
            description "used as export policy toward WAN BGP peers"
            entry 10 {
                from {
                    tag 30
                }
                action {
                    action-type reject
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            default-action {
                action-type accept
            }
        }
        policy-statement "vsi-502-export" {
            default-action {
                action-type accept
                community {
                    add ["RT64496:502" "SOO-45"]
                }
            }
        }
        policy-statement "vsi-502-import" {
            entry 10 {
                from {
                    family [evpn]
                    community {
                        name "SOO-45"
                    }
                }
                action {
                    action-type reject
                }
            }
            entry 20 {
                from {
                    family [evpn]
                    community {
                        name "RT64496:502"
                    }
                }
                action {
                    action-type accept
                }
            }
        }
        policy-statement "vsi-503-export" {
            default-action {
                action-type accept
                community {
                    add ["RT64496:503" "SOO-45"]
                }
            }
        }
        policy-statement "vsi-503-import" {
            entry 10 {
                from {
                    family [evpn]
                    community {
                        name "SOO-45"
                    }
                }
                action {
                    action-type reject
                }
            }
            entry 20 {
                from {
                    family [evpn]
                    community {
                        name "RT64496:503"
                    }
                }
                action {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    action-type accept
                }
            }
        }
    }
    service {
        vpls "VPLS-2" {
            admin-state enable
            description "dual BGP-EVPN MPLS instance VPLS"
            service-id 2
            customer "1"
            bgp 1 {
                route-distinguisher "192.0.2.45:2"
                vsi-import ["vsi-503-import"]
                vsi-export ["vsi-503-export"]
            }
            bgp 2 {
                route-distinguisher "192.0.2.54:2"
                vsi-import ["vsi-502-import"]
                vsi-export ["vsi-502-export"]
            }
            bgp-evpn {
                evi 2
                incl-mcast-orig-ip 45.45.45.45
                mpls 1 {
                    admin-state enable
                    default-route-tag 0x1e      # default route tag 30
                    mh-mode access
                    auto-bind-tunnel {
                        resolution filter
                        resolution-filter {
                            rsvp true
                        }
                    }
                }
                mpls 2 {
                    admin-state enable
                    default-route-tag 0x14      # default route tag 20
                    auto-bind-tunnel {
                        resolution filter
                        resolution-filter {
                            sr-isis true
                        }
                    }
                }
            }
        }
```

The following command on PE-2 shows BGP instances 1 and 2 in VPLS 2. RD 192.0.2.23:2 is configured in BGP instance 1; RD 192.0.2.32:2 is configured in BGP instance 2. The RTs are defined by virtual switching instance (VSI) policies.

```
[/]
A:admin@PE-2# show service id 2 bgp

===============================================================================
BGP Information
===============================================================================
Bgp Instance         : 1
Vsi-Import           : vsi-501-import
Vsi-Export           : vsi-501-export
Route Dist           : 192.0.2.23:2
Oper Route Dist      : 192.0.2.23:2
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Oper RD Type          : configured
Rte-Target Import     : None                Rte-Target Export: None
Oper RT Imp Origin    : vsi                 Oper RT Import   : Policy Based
Oper RT Exp Origin    : vsi                 Oper RT Export   : Policy Based
ADV Service MTU       : -1

Bgp Instance          : 2
Vsi-Import            : vsi-502-import
Vsi-Export            : vsi-502-export
Route Dist            : 192.0.2.32:2
Oper Route Dist       : 192.0.2.32:2
Oper RD Type          : configured
Rte-Target Import     : None                Rte-Target Export: None
Oper RT Imp Origin    : vsi                 Oper RT Import   : Policy Based
Oper RT Exp Origin    : vsi                 Oper RT Export   : Policy Based
ADV Service MTU       : -1

PW-Template Id        : None
-------------------------------------------------------------------------------
===============================================================================
```

The following command shows that EVPN destination 192.0.2.1 is reachable via an RSVP tunnel and EVPN destination 192.0.2.4 via an SR-ISIS tunnel. In EVPN-MPLS instance 2 of VPLS 2 on PE-2, the EVPN destination 192.0.2.4 is reachable via an SR-ISIS tunnel:

```
[/]
A:admin@PE-2# show service id 2 evpn-mpls

===============================================================================
BGP EVPN-MPLS Dest
===============================================================================
TEP Address                    Egr Label     Num.   Mcast Last Change
                               Transport:Tnl MACs         Sup BCast Domain
-------------------------------------------------------------------------------
192.0.2.1                      524284        1      bum   12/13/2022 08:53:25
                               rsvp:1                     No
192.0.2.4                      524278        1      bum   12/13/2022 08:53:50
                               isis:524291                No
-------------------------------------------------------------------------------
Number of entries : 2
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN-MPLS Ethernet Segment Dest
===============================================================================
Eth SegId                       Num. Macs              Last Change
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

When traffic is sent between CE-21 and CE-26, the FDB in PE-1 shows that traffic toward MAC address 00:00:5e:00:53:26 is sent via RSVP tunnel 1 toward PE-2:

```
[/]
A:admin@PE-1# show service id 2 fdb detail

===============================================================================
Forwarding Database, Service 2
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
2          00:00:5e:00:53:21 sap:1/1/c10/1:2         L/90     12/13/22 10:17:06
2          00:00:5e:00:53:26 mpls-1:                 Evpn     12/13/22 10:17:32
                             192.0.2.2:524280
           rsvp:1
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

The following command on PE-1 shows that only the EVPN MAC route received from PE-2 is used, not the one from PE-3 in the same anycast group. This is due to the best path selection done by BGP for the two routes, which have the same route key:

```
[/]
A:admin@PE-1# show router bgp routes evpn mac mac-address 00:00:5e:00:53:26
===============================================================================
 BGP Router ID:192.0.2.1        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN MAC Routes
===============================================================================
Flag  Route Dist.         MacAddr           ESI
      Tag                 Mac Mobility      Label1
                          Ip Address
                          NextHop
-------------------------------------------------------------------------------
u*>i  192.0.2.23:2        00:00:5e:00:53:26 ESI-0
      0                   Seq:0             LABEL 524280
                          n/a
                          192.0.2.2

*>i   192.0.2.23:2        00:00:5e:00:53:26 ESI-0
      0                   Seq:0             LABEL 524282
                          n/a
                          192.0.2.3

-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

The FDB for VPLS 2 on PE-2 shows that MAC address 00:00:5e:00:53:21 can be reached using EVPN-MPLS instance 1 whereas MAC address 00:00:5e:00:53:26 can be reached using EVPN-MPLS instance 2:

```
[/]
A:admin@PE-2# show service id 2 fdb detail

===============================================================================
Forwarding Database, Service 2
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
2          00:00:5e:00:53:21 mpls-1:                 Evpn     12/13/22 10:17:20
                             192.0.2.1:524284
           rsvp:1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
2          00:00:5e:00:53:26 mpls-2:                 Evpn     12/13/22 10:17:32
                             192.0.2.4:524278
           isis:524291
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

The FDB for VPLS 2 on PE-4 is as follows:

```
[/]
A:admin@PE-4# show service id 2 fdb detail

===============================================================================
Forwarding Database, Service 2
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
2          00:00:5e:00:53:21 mpls-2:                 Evpn     12/13/22 10:17:28
                             192.0.2.2:524279
           isis:524290
2          00:00:5e:00:53:26 mpls-1:                 Evpn     12/13/22 10:17:32
                             192.0.2.6:524284
           rsvp:1
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

The FDB for VPLS 2 on PE-6 is as follows:

```
[/]
A:admin@PE-6# show service id 2 fdb detail

===============================================================================
Forwarding Database, Service 2
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
2          00:00:5e:00:53:21 mpls-1:                 Evpn     12/13/22 10:17:34
                             192.0.2.4:524279
           rsvp:1
2          00:00:5e:00:53:26 sap:1/1/c10/1:2         L/60     12/13/22 10:17:32
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

Conclusion

Dual-instance EVPN-MPLS GWs reduce the number of service labels to be learned at the access nodes, and optimizes the replication of BUM traffic from the access nodes.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

EVPN E-LAN services with SRv6 transport

This chapter provides information about SRv6 support for distributed EVPN-enabled VPLS Layer 2 multipoint overlay services.

Topics in this chapter include:

Applicability

The information and configuration in this chapter are based on SR OS Release 22.10.R1. SRv6 support for distributed EVPN-enabled VPLS Layer 2 multipoint overlay services is supported on FP-based platforms with FP4-based network ports in SR OS Release 22.7.R1 and later.

Overview

On FP-based platforms with FP4-based network ports, SR OS provides SRv6 support for distributed EVPN-enabled VPLS Layer 2 multipoint overlay services. SRv6 tunnels carry EVPN data between the PEs on which the EVPN service is provisioned. As usual in EVPN services, a full mesh of SRv6 tunnels is set up among all PEs that participate in the EVPN-enabled VPLS service. This supports the flooding of Broadcast, Unknown unicast, or Multicast (BUM) traffic to all remote destinations in the service, while ensuring that the PEs receive the traffic without looping or duplication of frames. Two or more routers may participate in a single EVPN-enabled VPLS service; a single router may participate in multiple EVPN-enabled VPLS services. The PE routers attached to an EVPN-enabled VPLS service with SRv6 transport use SRv6 End.DT2U behavior to terminate and forward unicast traffic, and SRv6 End.DT2M behavior to terminate and forward BUM traffic.

An SRv6 L2 Service TLV, which is carried in a BGP Prefix-SID attribute, signals the SRv6 Service SID for the End.DT2U or End.DT2M behavior for an EVPN-enabled VPLS Layer 2 overlay service, as per RFC 9252. The SRv6 Service SID is equivalent to an MPLS label for EVPN service routes in RFC 7432.

When a PE is attached to an EVPN-enabled VPLS service with SRv6 transport, the PE advertises its originating IP address in an Inclusive Multicast Ethernet Tag (IMET) route (also known as an EVPN type 3 route), along with the service attributes and the SRv6 SID corresponding to the End.DT2M behavior for the service. A remote PE attached to the same EVPN-enabled VPLS service imports the IMET route based on the import route target and adds an SRv6 destination entry to its flooding list for the EVPN-enabled VPLS service. In this way, all PEs that participate in an EVPN-enabled VPLS service learn about each other.

As in any other type of EVPN-enabled VPLS service, a PE learns the MAC address of a locally connected CE, either via data plane MAC learning or static provisioning. In the case of data plane MAC learning, a PE learns the source MAC address from data frames that it receives from the CE and adds a temporary entry for it in a VPLS forwarding database (FDB), which, on each PE, is private for each EVPN-enabled VPLS service.

A local MAC address is advertised in an EVPN MAC/IP advertisement route (EVPN type 2 route) for the EVPN-enabled VPLS service, along with the service parameters and an SRv6 SID corresponding to the

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

End.DT2U behavior for the service. A remote PE that imports the EVPN MAC/IP advertisement route adds an entry for the advertised MAC addresses to the FDB, pointing at an SRv6 destination based on the received SRv6 SID. In this way, remote PEs that participate in an EVPN-enabled VPLS service with SRv6 transport learn how to unicast return traffic to the remote (source) MAC address.

The **locator** command in the **service vpls <service-name> segment-routing-v6 <instance>** context configures the SRv6 locator that the PE uses to terminate SRv6 traffic for the EVPN-enabled VPLS service.

The base SRv6 configuration is as described in the "SRv6 Encapsulation in the Base Routing Instance" chapter in the Segment Routing and PCE volume in the *7450 ESS, 7750 SR, and 7950 XRS MD-CLI Advanced Configuration Guide - Part I*.

Configuration

Figure 83: Example topology shows the example topology with three PE routers. The SRv6-enabled network that it represents comprises PE-1, PE-2, and PE-3 in the control and data planes, and a BGP route reflector RR-4 in the control plane only. The SRv6-enabled network has only IPv6 addresses and interfaces. IS-IS and BGP are configured on all routers. The system interfaces have also an IPv4 address, from which a unique router-id is automatically derived for IS-IS and BGP respectively.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 83: Example topology

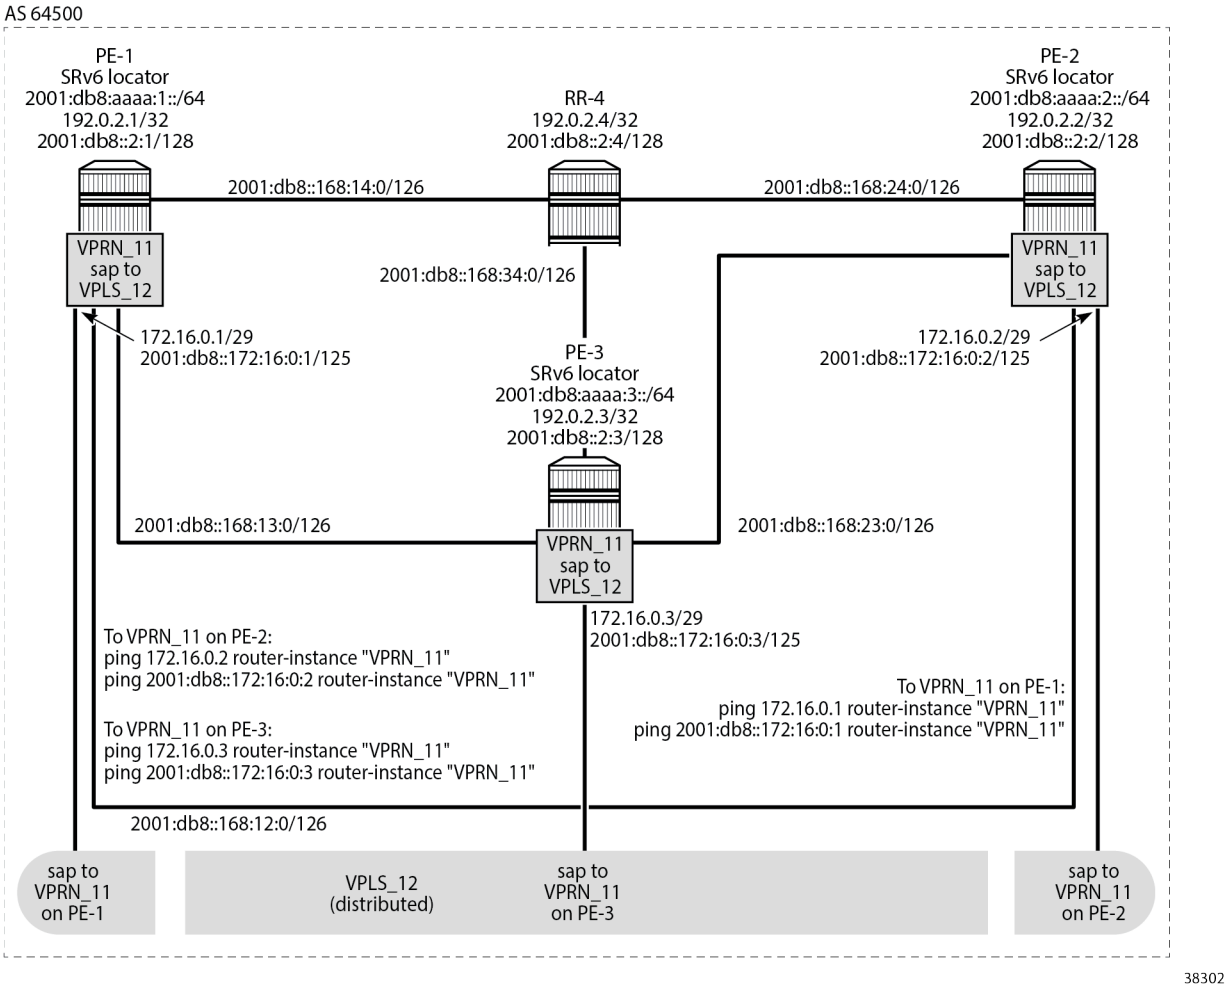

For the traffic of data frames from the EVPN-enabled VPLS service on a local PE to the same EVPN-enabled VPLS service on a remote PE, the local PE acts as the SRv6 ingress PE node, while the remote PE acts as the SRv6 egress PE node. SRv6 and forwarding port extensions (FPE) are configured only on the PE routers.

The **ping** commands between IPv4 and IPv6 interface addresses in the EVPN-enabled VPLS service simulate IPv4 and IPv6 data traffic respectively.

Configure the router

This configuration includes:

- on PE-1, PE-2, PE-3, and RR-4:
  - ports, IPv6-only interfaces, and system interfaces
  - IS-IS:
    - level 2 capability with wide metrics (for the 128-bit identifiers)

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- native IPv6 routing
- the **traffic-engineering** and **traffic-engineering-options** commands, as a best practice to advertise the router capability within the autonomous system (AS)

– BGP, with internal group "gr_v6_internal" that includes:

- the EVPN family
- BGP neighbor system IPv6 addresses

- on PE-1, PE-2, and PE-3, port cross-connect (PXC), using internal loopbacks on an FP4 MAC chip, as described in the "Segment Routing over IPv6" chapter in the Segment Routing and PCE volume in the *7450 ESS, 7750 SR, and 7950 XRS MD-CLI Advanced Configuration Guide - Part I*

The following example configuration applies for PE-1. A similar configuration applies for PE-2, PE-3, and RR-4. RR-4 has PE-1, PE-2 and PE-3 as BGP neighbors in a cluster.

```
A:admin@PE-1# configure {
    port 1/1/c2/1 {
        ethernet {
            mode hybrid
        }
        admin-state enable
    }
    port 1/1/c3/1 {
        ethernet {
            mode hybrid
        }
        admin-state enable
    }
    port 1/1/c4/1 {
        ethernet {
            mode hybrid
        }
        admin-state enable
    }
    port 1/1/c1/1 {
        ethernet {
            mode hybrid
        }
        admin-state enable
    }
    port 1/1/c1/2 {
        ethernet {
            mode hybrid
        }
        admin-state enable
    }
    router "Base" {
        interface "int-PE-1-PE-2" {
            description "interface between PE-1 and PE-2"
            port 1/1/c2/1:1000
            ipv6 {
                address 2001:db8::168:12:1 {
                    prefix-length 126
                }
            }
        }
        interface "int-PE-1-PE-3" {
            description "interface between PE-1 and PE-3"
            port 1/1/c3/1:1000
            ipv6 {
                address 2001:db8::168:13:1 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                prefix-length 126
            }
        }
    }
    interface "int-PE-1-RR-4" {
        description "interface between PE-1 and RR-4"
        port 1/1/c4/1:1000
        ipv6 {
            address 2001:db8::168:14:1 {
                prefix-length 126
            }
        }
    }
    interface "system" {
        ipv4 {
            primary {
                address 192.0.2.1
                prefix-length 32
            }
        }
        description "system interface of PE-1"
        ipv6 {
            address 2001:db8::2:1 {
                prefix-length 128
            }
        }
    }
    autonomous-system 64500
    isis 0 {
        level-capability 2
        area-address [49.0001]
        traffic-engineering true
        traffic-engineering-options {
            ipv6 true
            application-link-attributes {
            }
        }
        advertise-router-capability as
        ipv6-routing native
        level 2 {
            wide-metrics-only true
        }
        interface "system" {
            passive true
        }
        interface "int-PE-1-PE-2" {
            interface-type point-to-point
        }
        interface "int-PE-1-PE-3" {
            interface-type point-to-point
        }
        interface "int-PE-1-RR-4" {
            interface-type point-to-point
        }
        admin-state enable
    }
    bgp {
        rapid-withdrawal true
        split-horizon true
        rapid-update  {
            evpn true
        }
        group "gr_v6_internal" {
            description "internal bgp group on PE-1"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                family {
                    evpn true
                }
                peer-as 64500
            }
            neighbor "2001:db8::2:4" {
                group "gr_v6_internal"
            }
        }
    }
}
```

Configure the VPRNs to simulate CEs

On each PE, the VPRN configuration includes an IPv4 address and an IPv6 address for an interface from the local VPRN to the EVPN-enabled VPLS service. These IPv4 and IPv6 addresses must be in the same address range on all PEs, because the same EVPN-enabled VPLS service is provisioned on each PE. Each interface to the (local) EVPN-enabled VPLS service also includes a SAP.

The VPRNs are introduced only to simulate CEs from where the **ping** commands can be launched.

The following example configuration applies for VPRN 11 on PE-1. A similar configuration applies for VPRN 11 on PE-2 and for VPRN 11 on PE-3.

```
A:admin@PE-1# configure {
    service {
        vprn "VPRN_11" {
            service-id 11
            customer "1"
            description "CE_1"
            interface "local" {
                mac 00:00:5e:00:53:01
                ipv4 {
                    primary {
                        address 172.16.0.1
                        prefix-length 29
                    }
                }
                ipv6 {
                    address 2001:db8::172:16:0:1 {
                        prefix-length 125
                    }
                }
                sap 1/1/c1/2:11 {
                }
            }
            admin-state enable
        }
    }
}
```

For example, VPRN 11 on PE-2 has the following interface, with corresponding IPv4 and IPv6 addresses. Similar output applies for VPRN 11 on PE-1 and for VPRN 11 on PE-3.

```
A:admin@PE-2# show router 11 interface

===============================================================================
Interface Table (Service: 11)
===============================================================================
Interface-Name                   Adm       Opr(v4/v6)  Mode    Port/SapId
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
   IP-Address                                                     PfxState
-------------------------------------------------------------------------------
local                           Up        Up/Up       VPRN     1/1/c2/2:11
   172.16.0.2/29                                                  n/a
   2001:db8::172:16:0:2/125                                       PREFERRED
   fe80::200:22ff:fe22:2222/64                                    PREFERRED
-------------------------------------------------------------------------------
Interfaces : 1
===============================================================================
```

VPRN 11 on PE-1 has the following IPv4 and IPv6 routes. Similar output applies for VPRN 11 on PE-2 and for VPRN 11 on PE-3.

For IPv4:

```
A:admin@PE-1# show router 11 route-table

===============================================================================
Route Table (Service: 11)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
172.16.0.0/29                                 Local   Local     00h01m43s  0
       local                                                        0
-------------------------------------------------------------------------------
No. of Routes: 1
---snip---
===============================================================================
```

For IPv6:

```
A:admin@PE-1# show router 11 route-table ipv6

===============================================================================
IPv6 Route Table (Service: 11)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
2001:db8::172:16:0:0/125                      Local   Local     00h01m42s  0
       local                                                        0
-------------------------------------------------------------------------------
No. of Routes: 1
---snip---
===============================================================================
```

VPRN 11 on PE-1 has one locally learned MAC address for the locally configured interface. Similar output applies for VPRN 11 on PE-2 and for VPRN 11 on PE-3.

```
A:admin@PE-1# show router 11 arp summary

===============================================================
ARP Table Summary (Service: 11)
===============================================================
Local ARP Entries    : 1
---snip---
Dynamic ARP Entries  : 0
---snip---
---------------------------------------------------------------
No. of ARP Entries   : 1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================
```

The **show router 11 arp** command shows the association between the IP address and the MAC address, and the interface that the MAC address belongs to. The MAC address for the local interface to the EVPN-enabled VPLS service corresponds with that of the SAP that is configured for it in VPRN 11. Because the interface is statically configured, the association between the IP address and the MAC address does not expire. Similar output applies for PE-2 and for PE-3.

```
A:admin@PE-1# show router 11 arp

===============================================================================
ARP Table (Service: 11)
===============================================================================
IP Address      MAC Address       Expiry    Type   Interface
-------------------------------------------------------------------------------
172.16.0.1      00:00:5e:00:53:01 00h00m00s Oth[I] local
-------------------------------------------------------------------------------
No. of ARP Entries: 1
===============================================================================
```

Configure data path support, FPE, and SRv6

Configure data path support (PXC) and FPE identically on PE-1, PE2, and PE-3.

```
A:admin@PE-1# configure {
    card 1 {
        mda 1 {
            xconnect {
                mac 1 {
                    loopback 1 {
                    }
                    loopback 2 {
                    }
                }
            }
        }
    }
    port-xc {
        pxc 1 {
            port 1/1/m1/1
            admin-state enable
        }
        pxc 2 {
            port 1/1/m1/2
            admin-state enable
        }
    }
    port pxc-1.a {
        admin-state enable
    }
    port pxc-1.b {
        admin-state enable
    }
    port pxc-2.a {
        admin-state enable
    }
    port pxc-2.b {
        admin-state enable
    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    port 1/1/m1/1 {
        admin-state enable
    }
    port 1/1/m1/2 {
        admin-state enable
    }
    fwd-path-ext {
        fpe 1 {
            path {
                pxc 1
            }
            application {
                srv6 {
                    type origination
                }
            }
        }
        fpe 2 {
            path {
                pxc 2
            }
            application {
                srv6 {
                    type termination
                }
            }
        }
    }
}
```

Configure the SRv6 locator *"PE-1_loc_VPLS"* with **ip-prefix** *2001:db8:aaaa:1::/64* in the **router Base segment-routing segment-routing-v6** context on PE-1 and similar on PE-2, with **ip-prefix** *2001:db8:aaaa:2::/64* for SRv6 locator *"PE-2_loc_VPLS"*, and on PE-3, with **ip-prefix** *2001:db8:aaaa:3::/64* for SRv6 locator *"PE-3_loc_VPLS"*.

```
A:admin@PE-1# configure {
    router "Base" {
        segment-routing {
            segment-routing-v6 {
                source-address 2001:db8::2:1
                locator "PE-1_loc_VPLS" {
                    block-length 48
                    prefix {
                        ip-prefix 2001:db8:aaaa:1::/64
                    }
                    admin-state enable
                }
            }
        }
    }
}
```

Use FPE 1 as the SRv6 origination FPE and FPE 2 as the SRv6 termination FPE on PE-1, and similar on PE-2 for SRv6 locator *"PE-2_loc_VPLS"*, and on PE-3 for SRv6 locator *"PE-3_loc_VPLS"*. For more information, see the "Segment Routing over IPv6" chapter in the Segment Routing and PCE volume in the *7450 ESS, 7750 SR, and 7950 XRS MD-CLI Advanced Configuration Guide - Part I*.

```
configure {
    router "Base" {
        segment-routing {
            segment-routing-v6 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                origination-fpe [1]
                locator "PE-1_loc_VPLS" {
                    termination-fpe [2]
                    admin-state enable
                }
            }
        }
    }
}
```

Advertise the SRv6 locator *"PE-1_loc_VPLS"* in IS-IS while ensuring level 2 capability on PE-1, and similar on PE-2 for SRv6 locator *"PE-2_loc_VPLS"*, and on PE-3 for SRv6 locator *"PE-3_loc_VPLS"*.

```
A:admin@PE-1# configure {
    router "Base" {
        isis 0 {
            segment-routing-v6 {
                locator "PE-1_loc_VPLS" {
                    level-capability 2
                }
                admin-state enable
            }
        }
    }
}
```

Verify the IS-IS data base on PE-1 with the **show router isis 0 database detail** command. The output of this command (shortened here for PE-1, PE-3 and RR-4) provides information about each IS-IS-enabled router. For each uniquely identified IS-IS-enabled router, the SRv6 information indicates:

- the IS-IS-advertised router capabilities
- the IS-IS topology details
- the IPv4 and IPv6 reachability details
- the advertised SRv6 locator TLV
- the advertised configured SRv6 End SID and SRv6 End-X SIDs

```
A:admin@PE-1# show router isis 0 database detail

===============================================================================
Rtr Base ISIS Instance 0 Database (detail)
===============================================================================

Displaying Level 1 database
-------------------------------------------------------------------------------
Level (1) LSP Count : 0

Displaying Level 2 database
-------------------------------------------------------------------------------
LSP ID    : PE-1.00-00                                  Level     : L2
---snip---

-------------------------------------------------------------------------------
LSP ID    : PE-2.00-00                                  Level     : L2
Sequence  : 0xa                   Checksum  : 0xf34c   Lifetime  : 1147
Version   : 1                     Pkt Type  : 20       Pkt Ver   : 1
Attributes: L1L2                  Max Area  : 3        Alloc Len : 432
SYS ID    : 1920.0000.2002        SysID Len : 6        Used Len  : 432

TLVs :
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
  Area Addresses:
    Area Address : (3) 49.0001
  Supp Protocols:
    Protocols     : IPv4
    Protocols     : IPv6
  IS-Hostname   : PE-2
  Router ID   :
    Router ID   : 192.0.2.2
  TE Router ID v6  :
    Router ID   : 2001:db8::2:2
  Router Cap : 192.0.2.2, D:0, S:0
    TE Node Cap : B E M  P
    SRv6 Cap: 0x0000
    SR Alg: metric based SPF
    Node MSD Cap: BMI : 0 SRH-MAX-SL : 10 SRH-MAX-END-POP : 9 SRH-MAX-H-ENCAPS : 3 SRH-MAX-END-
D : 9
  I/F Addresses :
    I/F Address   : 192.0.2.2
  I/F Addresses IPv6 :
    IPv6 Address    : 2001:db8::2:2
    IPv6 Address    : 2001:db8::168:12:2
    IPv6 Address    : 2001:db8::168:23:1
    IPv6 Address    : 2001:db8::168:24:1
  TE IS Nbrs   :
    Nbr   : PE-1.00
    Default Metric  : 10
    Sub TLV Len     : 36
    IPv6 Addr : 2001:db8::168:12:2
    Nbr IPv6  : 2001:db8::168:12:1
  TE IS Nbrs   :
    Nbr   : PE-3.00
    Default Metric  : 10
    Sub TLV Len     : 36
    IPv6 Addr : 2001:db8::168:23:1
    Nbr IPv6  : 2001:db8::168:23:2
  TE IS Nbrs   :
    Nbr   : RR-4.00
    Default Metric  : 10
    Sub TLV Len     : 18
    IPv6 Addr : 2001:db8::168:24:1
  TE IP Reach   :
    Default Metric  : 0
    Control Info:    , prefLen 32
    Prefix   : 192.0.2.2
  IPv6 Reach:
    Metric: ( I ) 0
    Prefix   : 2001:db8::2:2/128
    Metric: ( I ) 10
    Prefix   : 2001:db8::168:12:0/126
    Metric: ( I ) 10
    Prefix   : 2001:db8::168:23:0/126
    Metric: ( I ) 10
    Prefix   : 2001:db8::168:24:0/126
    Metric: ( I ) 0
    Prefix   : 2001:db8:aaaa:2::/64
  SRv6 Locator  :
    MT ID : 0
    Metric: ( ) 0 Algo:0
    Prefix   : 2001:db8:aaaa:2::/64

-------------------------------------------------------------------------------
LSP ID    : PE-3.00-00                                Level     : L2
---snip---
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
LSP ID    : RR-4.00-00                                  Level     : L2
---snip---

Level (2) LSP Count : 4
-------------------------------------------------------------------------------
---snip---
===============================================================================
```

PE-1 learns the remote SRv6 locators that PE-2 and PE-3 advertise and installs a route for them in the IPv6 routing table. This route uses an SRv6 tunnel. Similar output applies for PE-2 and for PE-3.

```
A:admin@PE-1# show router route-table ipv6

===============================================================================
IPv6 Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
---snip---
2001:db8::2:2/128                             Remote  ISIS      00h12m04s  18
       fe80::60e:1ff:fe01:1-"int-PE-1-PE-2"                          10
2001:db8::2:3/128                             Remote  ISIS      00h12m04s  18
       fe80::612:1ff:fe01:1-"int-PE-1-PE-3"                          10
---snip---
2001:db8:aaaa:1::/64                          Local   SRV6      00h13m24s  3
       fe80::201-"_tmnx_fpe_2.a"                                     0
2001:db8:aaaa:2::/64                          Remote  ISIS      00h11m52s  18
       2001:db8:aaaa:2::/64 (tunneled:SRV6-ISIS)                     10
2001:db8:aaaa:3::/64                          Remote  ISIS      00h11m38s  18
       2001:db8:aaaa:3::/64 (tunneled:SRV6-ISIS)                     10
-------------------------------------------------------------------------------
No. of Routes: 13
---snip---
===============================================================================
```

Next to its own local locator prefix, PE-1 also learns the remote locator prefixes that PE-2 and PE-3 advertise. Similar output applies for PE-2 and for PE-3.

```
A:admin@PE-1# show router isis 0 segment-routing-v6 locator

===============================================================================
Rtr Base ISIS Instance 0 SRv6 Locator Table
===============================================================================
Prefix                              AdvRtr                        MT     Lvl/Typ
 AttributeFlags                      Tag                           Flags  Algo
-------------------------------------------------------------------------------
2001:db8:aaaa:1::/64                PE-1                          0      2/Int.
  -                                   0                             -      0
2001:db8:aaaa:2::/64                PE-2                          0      2/Int.
  -                                   0                             -      0
2001:db8:aaaa:3::/64                PE-3                          0      2/Int.
  -                                   0                             -      0
-------------------------------------------------------------------------------
No. of Locators: 3
-------------------------------------------------------------------------------
---snip---
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

From PE-1, the remote locator prefix 2001:db8:aaaa:2::/64 is routable via a next hop using the "int-PE-1-PE-2" interface. Similar output applies for the remote locator prefix 2001:db8:aaaa:3::/64 using the "int-PE-1-PE-3" interface. Similar output applies from PE-2 and from PE-3.

```
A:admin@PE-1# show router isis 0 routes

===============================================================================
Rtr Base ISIS Instance 0 Route Table
===============================================================================
Prefix[Flags]                     Metric     Lvl/Typ     Ver.  SysID/Hostname
  NextHop                                                MT     AdminTag/SID[F]
-------------------------------------------------------------------------------
---snip---
2001:db8::2:2/128                 10         2/Int.      12    PE-2
   fe80::60e:1ff:fe01:1-"int-PE-1-PE-2"                    0       0
2001:db8::2:3/128                 10         2/Int.      12    PE-3
---snip---
2001:db8:aaaa:1::/64              0          2/Int.      16    PE-1
   ::                                                      0       0
2001:db8:aaaa:2::/64              10         2/Int.      13    PE-2
   fe80::60e:1ff:fe01:1-"int-PE-1-PE-2"                    0       0
2001:db8:aaaa:3::/64              10         2/Int.      15    PE-3
   fe80::612:1ff:fe01:1-"int-PE-1-PE-3"                    0       0
-------------------------------------------------------------------------------
No. of Routes: 14 (14 paths)
-------------------------------------------------------------------------------
---snip---
===============================================================================
```

PE-1 transports IPv4 and IPv6 data to the remote SRv6 locator prefixes in an SRv6 encapsulated tunnel. For each SRv6 locator prefix destination, PE-1 sets up a different SRv6 tunnel with its specific label (TunnelId). Similar output applies for PE-2 and for PE-3.

```
A:admin@PE-1# show router tunnel-table ipv6

===============================================================================
IPv6 Tunnel Table (Router: Base)
===============================================================================
Destination                                Owner     Encap TunnelId  Pref
Nexthop                                    Color           Metric
-------------------------------------------------------------------------------
2001:db8:aaaa:2::/64                       srv6-isis SRV6  524289    0
  fe80::60e:1ff:fe01:1-"int-PE-1-PE-2"                       10
2001:db8:aaaa:3::/64                       srv6-isis SRV6  524290    0
  fe80::612:1ff:fe01:1-"int-PE-1-PE-3"                       10
-------------------------------------------------------------------------------
---snip---
===============================================================================
```

The **show router fp-tunnel-table 1 ipv6** command in PE-1 shows the local endpoints of the SRv6 tunnels in PE-1. Similar output applies for the local endpoints of the SRv6 tunnels in PE-2 and for the local endpoints of the SRv6 tunnels in PE-3.

```
A:admin@PE-1# show router fp-tunnel-table 1 ipv6

===============================================================================
IPv6 Tunnel Table Display
---snip---
===============================================================================
Destination                                Protocol         Tunnel-ID
  Lbl/SID
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    NextHop                                                 Intf/Tunnel
  Lbl/SID (backup)
    NextHop   (backup)
-------------------------------------------------------------------------------
2001:db8:aaaa:2::/64                        SRV6              524289
  -
    fe80::60e:1ff:fe01:1-"int-PE-1-PE-2"                    1/1/c2/1:1000
2001:db8:aaaa:3::/64                        SRV6              524290
  -
    fe80::612:1ff:fe01:1-"int-PE-1-PE-3"                    1/1/c3/1:1000
-------------------------------------------------------------------------------
Total Entries : 2
-------------------------------------------------------------------------------
===============================================================================
```

Verify data traffic

At this point, verify that IPv4 and IPv6 data traffic is not possible between the local VPRN 11 on PE-1 and the remote VPRN 11 on PE-2 and PE-3. PE-1 is not aware of the remote MAC addresses that are associated with IPv4 address 172.16.0.2 and IPv4 address 172.16.0.3 (or IPv6 address 2001:db8::172:16:0:2 and IPv6 address 2001:db8::172:16:0:3), because only interfaces that are locally connected to the EVPN-enabled VPLS service on PE-1 reply on the ARP request. Perform a similar verification for IPv4 and IPv6 data traffic between the local VPRN 11 on PE-2 and the remote VPRN 11 on PE-1 and PE-3, and for IPv4 and IPv6 data traffic between the local VPRN 11 on PE-3 and the remote VPRN 11 on PE-1 and PE-2.

For example, for IPv4 data traffic to the remote VPRN 11 on PE-2:

```
A:admin@PE-1# ping 172.16.0.2 router-instance "VPRN_11"
PING 172.16.0.2 56 data bytes
...   ...   ...   ...   . Request timed out. icmp_seq=1.
Request timed out. icmp_seq=2.
---snip---
---- 172.16.0.2 PING Statistics ----
5 packets transmitted, 0 packets received, 100% packet loss
```

For example, for IPv6 data traffic to the remote VPRN 11 on PE-2:

```
A:admin@PE-1# ping 2001:db8::172:16:0:2 router-instance "VPRN_11"
PING 2001:db8::172:16:0:2 56 data bytes
...   ...   . 112 bytes from 2001:db8::172:16:0:1 Address unreachable
VR CLS   LEN NXT HLIM SRC
 6  00    64  58   64 2001:db8::172:16:0:1
                      DST
                      2001:db8::172:16:0:2
ICMP6: Echo request
---snip---
---- 2001:db8::172:16:0:2 PING Statistics ----
5 packets transmitted, 5 packets bounced, 0 packets received, 100% packet loss
```

Configure the EVPN- and SRv6-enabled VPLS service on PE-1, PE-2, and PE-3

On each PE, this configuration includes a SAP to the local VPRN.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

On PE-1, create an SRv6 instance *1* for the EVPN-enabled VPLS service. Use the SRv6 locator *"PE-1_loc_VPLS"* from the **router Base segment-routing segment-routing-v6** context in the **service vpls "VPLS_12" segment-routing-v6 1** context and configure End.DT2U and End.DT2M behavior for it.

Use the configured SRv6 locator *"PE-1_loc_VPLS"* as the default locator in the **service vpls "VPLS_12" bgp-evpn segment-routing-v6 1** context. In the **service vpls "VPLS_12" bgp-evpn segment-routing-v6 1 locator "PE-1_loc_VPLS"** context, use the unique PE-1 system IPv6 address as the route next hop. This configuration can be verified with the **show service id 12 bgp** command (not shown). Perform a similar configuration on PE-2 (and PE-3), with the configured SRv6 locator *"PE-2_loc_VPLS"* (*"PE-3_loc_VPLS"*) as the default locator, and the PE-2 (PE-3) system IPv6 address as route next hop.

```
A:admin@PE-1# configure {
    service {
        vpls "VPLS_12" {
            service-id 12
            customer "1"
            description "VPLS_12 on PE-1"
            segment-routing-v6 1 {
                locator "PE-1_loc_VPLS" {
                    function {
                        end-dt2u {
                        }
                        end-dt2m {
                        }
                    }
                }
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 1
                segment-routing-v6 1 {
                    srv6 {
                        instance 1
                        default-locator "PE-1_loc_VPLS"
                    }
                    route-next-hop {
                        system-ipv6
                    }
                    admin-state enable
                }
            }
            sap 1/1/c1/1:11 {
                description "sap to VPRN_11 on PE-1"
            }
            admin-state enable
        }
    }
}
```

The **show service id 12 fdb expiry** command shows that MAC learning and MAC aging are enabled. For example, the VPLS FDB entries that are locally learned expire after 300 seconds.

```
A:admin@PE-1# show service id 12 fdb expiry

===============================================================================
Forwarding Database, Service 12
===============================================================================
---snip---
Table Size        : 250                Allocated Count   : 0
Total In Use      : 0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Learned Count     : 0                   Static Count      : 0
---snip---
BGP EVPN Count    : 0                   EVPN Static Cnt   : 0
---snip---
Remote Age        : 900                 Local Age         : 300
---snip---
Mac Learning      : Enabled             Discard Unknown   : Disabled
Mac Aging         : Enabled             Relearn Only      : False
---snip---
===============================================================================
```

The **show service id 12 bgp-evpn** command shows how BGP EVPN behavior is configured. MAC advertisement for EVPN MAC/IP advertisement routes (for **ping** commands) and inclusive multicast advertisement for EVPN IMET routes (for flooding and BUM traffic) are enabled. The next hop corresponds with the local system IPv6 address. The route resolution uses the route table of the VPRN that has a local interface to the EVPN-enabled VPLS service. Similar output applies for PE-2 and for PE-3.

```
A:admin@PE-1# show service id 12 bgp-evpn

===============================================================================
BGP EVPN Table
===============================================================================
EVI                : 1
Creation Origin    : manual

MAC/IP Routes
MAC Advertisement  : Enabled             Unknown MAC Route  : Disabled
CFM MAC Advertise  : Disabled

Multicast Routes
Sel Mcast Advert   : Disabled
Ing Rep Inc McastAd: Enabled
---snip---
===============================================================================
Segment Routing v6 Instance 1 Service 12
===============================================================================
Admin State             : Enabled
Srv6 Instance           : 1
Default Locator         : PE-1_loc_VPLS

Oper Group              : (Not Specified)
Default Route Tag       : 0x0
Source Address          : (Not Specified)
ECMP                    : 1
Force Vlan VC Fwd       : disabled
Next Hop Type           : system-ipv6
Evi 3-byte Auto-RT      : disabled
Route Resolution        : route-table
Force QinQ VC Fwd       : none
MH Mode                 : network
Rest Prot Src Mac       : disabled
Split Horizon Group     : n/a
===============================================================================
```

The configuration of the SRv6 End.DT2U and End.DT2M behavior for the SRv6 locator that is used in the EVPN-enabled VPLS service results in corresponding SRv6 full SIDs. For example, the **show service id 12 segment-routing-v6 instance 1** command on PE-2 shows them. For the SRv6 End.DT2U behavior, the SRv6 function is 524288 (0x80000) and the corresponding SRv6 full SID is 2001:db8::aaaa:2:8000::.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

For the SRv6 End.DT2M behavior, the SRv6 function is 524287 (0x7ffff) and the corresponding SRv6 full SID is 2001:db8::aaaa:2:7fff:f000::. Similar output applies for PE-1 and for PE-3.

```
A:admin@PE-2# show service id 12 segment-routing-v6 instance 1

===============================================================================
Segment Routing v6 Instance 1 Service 12
===============================================================================
Locator
  Type         Function  SID                                        Status
-------------------------------------------------------------------------------
PE-2_loc_VPLS
  End.DT2U       *524288 2001:db8:aaaa:2:8000::                     ok
  End.DT2M       *524287 2001:db8:aaaa:2:7fff:f000::                ok
===============================================================================
Legend: * - System allocated
```

The **show router segment-routing-v6 local-sid** command shows that the SRv6 local SIDs belong to the VPLS context. Similar output applies for PE-1 and for PE-3.

```
A:admin@PE-2# show router segment-routing-v6 local-sid

===============================================================================
Segment Routing v6 Local SIDs
===============================================================================
SID                                             Type           Function
  Locator
  Context
-------------------------------------------------------------------------------
2001:db8:aaaa:2:7fff:f000::                     End.DT2M       524287
  PE-2_loc_VPLS
  SvcId: 12 Name: VPLS_12
2001:db8:aaaa:2:8000::                          End.DT2U       524288
  PE-2_loc_VPLS
  SvcId: 12 Name: VPLS_12
-------------------------------------------------------------------------------
SIDs : 2
-------------------------------------------------------------------------------
===============================================================================
```

Enabling the SRv6 End.DT2M behavior allows the exchange of EVPN IMET BGP update messages for the EVPN family. The **show log log-id <log-id>** command on PE-1 shows the BGP update message that PE-1 receives from PE-2, via the RR. It indicates the remote source address (orig_addr: 2001:db8::2:2), and the route distinguisher (RD: 192.0.2.2:1), tag (tag: 0), route target (Extended Community: target:64500:1), and next hop (Global NextHop 2001:db8::2:2) that PE-1 must use while sending IPv4 or IPv6 data traffic to PE-2. In addition, it indicates the Provicer Multicast Service Interface (PMSI) information about tunnel type (Tunnel-type Ingress Replication), MPLS label (MPLS Label 8388592 (0x7fffff)), and tunnel endpoint (Tunnel-Endpoint 2001:db8::2:2). Finally, it indicates that PE-1 must send the frames to the SRv6 locator (SRv6 SID: 2001:db8:aaaa:2::) with End.DT2M behavior (Behavior: 0x18 (24)). Similar output applies for the BGP update that PE-1 receives from PE-3, via the RR. PE-1 advertises a similar BGP update message to the RR, which forwards it to PE-2 and PE-3 (not shown here). PE-2 and PE-3 receive and advertise similar BGP update messages.

```
A:admin@PE-1# show log log-id "log_2"

===============================================================================
Event Log 1 log-name log_2
===============================================================================
Description : (Not Specified)
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Memory Log contents  [size=100   next event=4  (not wrapped)]
---snip---
2 2023/01/04 16:58:01.781 CET MINOR: DEBUG #2001 Base Peer 1: 2001:db8::2:4
"Peer 1: 2001:db8::2:4: UPDATE
Peer 1: 2001:db8::2:4 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 159
    Flag: 0x90 Type: 14 Len: 52 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 16 Global NextHop 2001:db8::2:2
        Type: EVPN-INCL-MCAST Len: 29 RD: 192.0.2.2:1, tag: 0, orig_addr len: 128, orig_addr:
 2001:db8::2:2
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0x80 Type: 9 Len: 4 Originator ID: 192.0.2.2
    Flag: 0x80 Type: 10 Len: 4 Cluster ID:
        4.4.4.4
    Flag: 0xc0 Type: 16 Len: 8 Extended Community:
        target:64500:1
    Flag: 0xc0 Type: 22 Len: 21 PMSI:
        Tunnel-type Ingress Replication (6)
        Flags: (0x0)[Type: None BM: 0 U: 0 Leaf: not required]
        MPLS Label 8388592
        Tunnel-Endpoint 2001:db8::2:2
    Flag: 0xc0 Type: 40 Len: 37 Prefix-SID-attr:
       SRv6 Services TLV (37 bytes):-
           Type: SRV6 L2 Service TLV (6)
           Length: 34 bytes, Reserved: 0x0
         SRv6 Service Information Sub-TLV (33 bytes)
            Type: 1 Len: 30 Rsvd1: 0x0
            SRv6 SID: 2001:db8:aaaa:2::
            SID Flags: 0x0 Endpoint Behavior: 0x18 Rsvd2: 0x0
            SRv6 SID Sub-Sub-TLV
               Type: 1 Len: 6
               BL:48 NL:16 FL:20 AL:0 TL:20 TO:64
"
---snip---
```

The reception of the EVPN IMET BGP update messages triggers PE-1 to install learned inclusive multicast routes as shown with the **show router bgp neighbor <ip-address> received-routes evpn** command. Because PE-1 receives EVPN IMET BGP update messages from PE-2 and from PE-3 with different route distinguishers, PE-1 installs a learned inclusive multicast route for each one of them. Similar output applies for PE-2 and for PE-3. The BGP EVPN inclusive multicast routes that are received, can also be displayed with the **show router bgp routes evpn incl-mcast** command.

```
A:admin@PE-1# show router bgp neighbor 2001:db8::2:4 received-routes evpn
===============================================================================
 BGP Router ID:192.0.2.1        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
---snip---
===============================================================================
BGP EVPN Inclusive-Mcast Routes
===============================================================================
Flag  Route Dist.         OrigAddr
      Tag                 NextHop
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
u*>i  192.0.2.2:1         2001:db8::2:2
      0                   2001:db8::2:2

u*>i  192.0.2.3:1         2001:db8::2:3
      0                   2001:db8::2:3

-------------------------------------------------------------------------------
Routes : 2
===============================================================================
---snip---
=======================================================
```

The **show router bgp neighbor <ip-address> advertised-routes evpn** command on PE-2 shows the local inclusive multicast routes on PE-2. PE-2 advertises them to its BGP neighbors. Similar output applies for PE-1 and for PE-3.

```
A:admin@PE-2# show router bgp neighbor 2001:db8::2:4 advertised-routes evpn
===============================================================================
 BGP Router ID:192.0.2.2        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
---snip---
===============================================================================
BGP EVPN Inclusive-Mcast Routes
===============================================================================
Flag  Route Dist.         OrigAddr
      Tag                 NextHop
-------------------------------------------------------------------------------
i     192.0.2.2:1         2001:db8::2:2
      0                   2001:db8::2:2

-------------------------------------------------------------------------------
Routes : 1
===============================================================================
---snip---
=======================================================
```

From the received EVPN IMET BGP update messages, PE-1 learns the SRv6 tunnel endpoints for multicast traffic, as shown with the **show service id 12 segment-routing-v6 instance 1 destinations** command. The segment ID (SRv6 SID) corresponds with the expected End.DT2M behavior on PE-2 and PE-3 respectively. Similar output applies for PE-2 and for PE-3.

```
A:admin@PE-1# show service id 12 segment-routing-v6 instance 1 destinations

===============================================================================
TEP, SID
===============================================================================
Instance  TEP Address               Segment Id                SupBcasDom Num
 Mcast                                                                   MACs
-------------------------------------------------------------------------------
1         2001:db8::2:2             2001:db8:aaaa:2:7fff:f000:  No        0
                                    :
 BUM
1         2001:db8::2:3             2001:db8:aaaa:3:7fff:f000:  No        0
                                    :
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
 BUM
------------------------------------------------------------------------------
Number of TEP, SID: 2
------------------------------------------------------------------------------
==============================================================================
---snip---
==============================================================================
```

The list of next hops for the EVPN family can be shown with the **show router bgp next-hop evpn** command. For each next hop, the details can be shown. The **show router bgp next-hop evpn 2001:db8::2:2 detail** command on PE-1 shows the details on PE-1 for next hop 2001:db8::2:2. It indicates that IPv4 and IPv6 data for the EVPN family uses the SRv6 tunnel for locator 2001:db8:aaaa:2::/64 and is sent to the next hop 2001:db8::2:2 via the resolved next hop fe80::60e:1ff:fe01:1, which corresponds with the "int-PE-1-PE-2" interface on PE-1. Similar output applies on PE-1 for next hop 2001:db8::2:3. Similar output applies for PE-2 and for PE-3.

```
A:admin@PE-1# show router bgp next-hop evpn 2001:db8::2:2 detail
==============================================================================
 BGP Router ID:192.0.2.1        AS:64500       Local AS:64500
==============================================================================

==============================================================================
BGP VPN Next Hop
==============================================================================
------------------------------------------------------------------------------
VPN Next Hop          : 2001:db8::2:2
Autobind              : gre/rtm
Labels                : --
User-labels           : 1
Admin-tag-policy      : --
Strict-tunnel-tagging : N
Color                 : --
Locator               : 2001:db8:aaaa:2::/64
Created               : 00h02m13s
Last-modified         : 00h02m13s
------------------------------------------------------------------------------
Resolving Prefix : 2001:db8::2:2/128
Preference       : 18                  Metric           : 10
Reference Count  : 1                   Owner            : GRE
Fib Programmed   : Y
Resolved Next Hop: fe80::60e:1ff:fe01:1
Egress Label     : n/a                 TunnelId         : 4294967293
Locator State    : Resolved
------------------------------------------------------------------------------
------------------------------------------------------------------------------
Next Hops : 1
==============================================================================
```

The **show router bgp routes evpn incl-mcast hunt** command shows a consolidated view on the inclusive multicast routes for the EVPN family. On PE-1, in the RIB In Entries section, it shows for each learned next hop how PE-1 must handle the BUM traffic destined for it. In the RIB Out Entries section, it shows for each local next hop how PE-1 expects the remote routers to handle BUM traffic destined for it. Similar output applies for PE-2 and for PE-3.

```
A:admin@PE-1# show router bgp routes evpn incl-mcast hunt
==============================================================================
 BGP Router ID:192.0.2.1        AS:64500       Local AS:64500
==============================================================================
---snip---
==============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
BGP EVPN Inclusive-Mcast Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : n/a
Nexthop        : 2001:db8::2:2
Path Id        : None
From           : 2001:db8::2:4
Res. Nexthop   : fe80::60e:1ff:fe01:1
Local Pref.    : 100                    Interface Name : int-PE-1-PE-2
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : 10
Connector      : None
Community      : target:64500:1
Cluster        : 4.4.4.4
Originator Id  : 192.0.2.2              Peer Router Id : 192.0.2.4
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
EVPN type      : INCL-MCAST
Tag            : 0
Originator IP  : 2001:db8::2:2
Route Dist.    : 192.0.2.2:1
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
Add Paths Send : Default
Last Modified  : 00h02m13s
SRv6 TLV Type  : SRv6 L2 Service TLV (6)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:aaaa:2::
Full Sid       : 2001:db8:aaaa:2:7fff:f000::
Behavior       : End.DT2M (24)
SRv6 SubSubTLV : SRv6 SID Structure (1)
Loc-Block-Len  : 48                     Loc-Node-Len   : 16
Func-Len       : 20                     Arg-Len        : 0
Tpose-Len      : 20                     Tpose-offset   : 64
-------------------------------------------------------------------------------
PMSI Tunnel Attributes :
Tunnel-type    : Ingress Replication
Flags          : Type: RNVE(0) BM: 0 U: 0 Leaf: not required
MPLS Label     : 8388592
Tunnel-Endpoint: 2001:db8::2:2
-------------------------------------------------------------------------------

Network        : n/a
Nexthop        : 2001:db8::2:3
---snip---

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
Network        : n/a
Nexthop        : 2001:db8::2:1
Path Id        : None
To             : 2001:db8::2:4
Res. Nexthop   : n/a
Local Pref.    : 100                    Interface Name : NotAvailable
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : n/a
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Connector      : None
Community      : target:64500:1
Cluster        : No Cluster Members
Originator Id  : None                  Peer Router Id : 192.0.2.4
Origin         : IGP
AS-Path        : No As-Path
EVPN type      : INCL-MCAST
Tag            : 0
Originator IP  : 2001:db8::2:1
Route Dist.    : 192.0.2.1:1
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                     Dest Class     : 0
SRv6 TLV Type  : SRv6 L2 Service TLV (6)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:aaaa:1::
Full Sid       : 2001:db8:aaaa:1:7fff:f000::
Behavior       : End.DT2M (24)
SRv6 SubSubTLV : SRv6 SID Structure (1)
Loc-Block-Len  : 48                    Loc-Node-Len   : 16
Func-Len       : 20                    Arg-Len        : 0
Tpose-Len      : 20                    Tpose-offset   : 64
-------------------------------------------------------------------------------
PMSI Tunnel Attributes :
Tunnel-type    : Ingress Replication
Flags          : Type: RNVE(0) BM: 0 U: 0 Leaf: not required
MPLS Label     : 8388592
Tunnel-Endpoint: 2001:db8::2:1
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
Routes : 3
===============================================================================
```

Verify data traffic

At this point, verify that IPv4 and IPv6 data traffic is possible between the local VPRN 11 on PE-1 and the remote VPRN 11 on PE-2 and PE-3. Perform a similar verification for IPv4 and IPv6 data traffic between the local VPRN 11 on PE-2 and the remote VPRN 11 on PE-1 and PE-3, and for IPv4 and IPv6 data traffic between the local VPRN 11 on PE-3 and the remote VPRN 11 on PE-1 and PE-2.

For example, for IPv4 data traffic to the remote VPRN 11 on PE-2:

```
A:admin@PE-1# ping 172.16.0.2 router-instance "VPRN_11"
PING 172.16.0.2 56 data bytes
64 bytes from 172.16.0.2: icmp_seq=1 ttl=64 time=7.34ms.
---snip---
---- 172.16.0.2 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 1.90ms, avg = 3.34ms, max = 7.34ms, stddev = 2.02ms
```

For example, for IPv6 data traffic to the remote VPRN 11 on PE-2:

```
A:admin@PE-1# ping 2001:db8::172:16:0:2 router-instance "VPRN_11"
PING 2001:db8::172:16:0:2 56 data bytes
64 bytes from 2001:db8::172:16:0:2 icmp_seq=1 hlim=64 time=5.50ms.
---snip---
---- 2001:db8::172:16:0:2 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
round-trip min = 1.76ms, avg = 2.79ms, max = 5.50ms, stddev = 1.37ms
```

When the SRv6 End.DT2U behavior is enabled, the sending of IPv4 or IPv6 data traffic triggers the exchange of EVPN MAC/IP BGP update messages for the EVPN family. The **show log log-id <log-id>** command on PE-1 shows the BGP update message that PE-1 receives from PE-2, via the RR. It indicates the remote MAC address (mac: 00:00:5e:00:53:02), and the route distinguisher (RD: 192.0.2.2:1), ESI (ESI: ESI-0), tag (tag: 0), label (label1: 8388608 (0x800000)), route target (Extended Community: target:64500:1), and next hop (Global NextHop 2001:db8::2:2) that PE-1 must use while sending IPv4 or IPv6 data traffic to PE-2. In addition, it indicates that PE-1 must send the frames to the SRv6 locator (SRv6 SID: 2001:db8:aaaa:2::) with End.DT2U behavior (Behavior: 0x17 (23)). PE-1 derives the SRv6 full SID that is needed for this (2001:db8:aaaa:2:8000::). Similar output applies for the BGP update that PE-1 receives from PE-3, via the RR. PE-1 advertises a similar BGP update message to the RR, which forwards it to PE-2 and PE-3 (not shown here). PE-2 and PE-3 receive and advertise similar BGP update messages.

```
A:admin@PE-1# show log log-id "log_2"

===============================================================================
Event Log 1 log-name log_2
===============================================================================
Description : (Not Specified)
Memory Log contents  [size=100   next event=4  (not wrapped)]
---snip---
2 2023/01/04 17:01:49.282 CET MINOR: DEBUG #2001 Base Peer 1: 2001:db8::2:4
"Peer 1: 2001:db8::2:4: UPDATE
Peer 1: 2001:db8::2:4 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 139
    Flag: 0x90 Type: 14 Len: 56 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 16 Global NextHop 2001:db8::2:2
        Type: EVPN-MAC Len: 33 RD: 192.0.2.2:1 ESI: ESI-0, tag: 0, mac len: 48 mac:
 00:00:5e:00:53:02, IP len: 0, IP: NULL, label1: 8388608
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0x80 Type: 9 Len: 4 Originator ID: 192.0.2.2
    Flag: 0x80 Type: 10 Len: 4 Cluster ID:
        4.4.4.4
    Flag: 0xc0 Type: 16 Len: 8 Extended Community:
        target:64500:1
    Flag: 0xc0 Type: 40 Len: 37 Prefix-SID-attr:
       SRv6 Services TLV (37 bytes):-
           Type: SRV6 L2 Service TLV (6)
           Length: 34 bytes, Reserved: 0x0
         SRv6 Service Information Sub-TLV (33 bytes)
            Type: 1 Len: 30 Rsvd1: 0x0
            SRv6 SID: 2001:db8:aaaa:2::
            SID Flags: 0x0 Endpoint Behavior: 0x17 Rsvd2: 0x0
            SRv6 SID Sub-Sub-TLV
               Type: 1 Len: 6
               BL:48 NL:16 FL:20 AL:0 TL:20 TO:64
"
---snip---
```

The reception of the EVPN MAC/IP BGP update messages triggers PE-1 to install learned MAC routes, as shown with the **show router bgp neighbor <ip-address> received-routes evpn** command. In contrast to the learned inclusive multicast routes, the learned MAC routes expire in accordance with the configuration that is shown with the **show service id 12 fdb expiry** command. PE-1 installs a learned MAC/IP route for each of the remote CEs. PE-1 derives the SRv6 function (524288) from the received label field. The earlier

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

installed inclusive multicast routes remain in place (not shown). The BGP EVPN MAC/IP advertisement routes that are received, can also be displayed with the **show router bgp routes evpn mac** command. Similar output applies for PE-2 and for PE-3.

```
A:admin@PE-1# show router bgp neighbor 2001:db8::2:4 received-routes evpn
===============================================================================
 BGP Router ID:192.0.2.1        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
---snip---
===============================================================================
BGP EVPN MAC Routes
===============================================================================
Flag  Route Dist.          MacAddr           ESI
      Tag                  Mac Mobility      Label1
                           Ip Address
                           NextHop
-------------------------------------------------------------------------------
u*>i  192.0.2.2:1          00:00:5e:00:53:02 ESI-0
      0                    Seq:0             LABEL 524288
                           n/a
                           2001:db8::2:2

u*>i  192.0.2.3:1          00:00:5e:00:53:03 ESI-0
      0                    Seq:0             LABEL 524288
                           n/a
                           2001:db8::2:3

-------------------------------------------------------------------------------
Routes : 2
===============================================================================

===============================================================================
BGP EVPN Inclusive-Mcast Routes
===============================================================================
---snip---
-------------------------------------------------------------------------------
Routes : 2
===============================================================================
---snip---
===============================================================
```

The **show router bgp neighbor <ip-address> advertised-routes evpn** command on PE-2 shows the local MAC routes on PE-2. PE-2 advertises them to its BGP neighbors. Similar output applies for PE-1 and for PE-3.

```
A:admin@PE-2# show router bgp neighbor 2001:db8::2:4 advertised-routes evpn
===============================================================================
 BGP Router ID:192.0.2.2        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
---snip---
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
BGP EVPN MAC Routes
===============================================================================
Flag  Route Dist.          MacAddr           ESI
      Tag                  Mac Mobility      Label1
                           Ip Address
                           NextHop
-------------------------------------------------------------------------------
i     192.0.2.2:1          00:00:5e:00:53:02 ESI-0
      0                    Seq:0             524288
                           n/a
                           2001:db8::2:2

-------------------------------------------------------------------------------
Routes : 1
===============================================================================

===============================================================================
BGP EVPN Inclusive-Mcast Routes
===============================================================================
Flag  Route Dist.          OrigAddr
      Tag                  NextHop
-------------------------------------------------------------------------------
i     192.0.2.2:1          2001:db8::2:2
      0                    2001:db8::2:2

-------------------------------------------------------------------------------
Routes : 1
===============================================================================
---snip---
===============================================================
```

From the received EVPN MAC/IP BGP update messages, PE-1 learns the SRv6 tunnel endpoints for unicast traffic, as shown with the **show service id 12 segment-routing-v6 instance 1 destinations** command. The segment ID (SRv6 SID) corresponds with the expected End.DT2U behavior on PE-2 and PE-3 respectively. The earlier learned SRv6 tunnel endpoints for BUM traffic remain in place. Similar output applies for PE-2 and for PE-3.

```
A:admin@PE-1# show service id 12 segment-routing-v6 instance 1 destinations

===============================================================================
TEP, SID
===============================================================================
Instance  TEP Address               Segment Id                SupBcasDom Num
 Mcast                                                                   MACs
-------------------------------------------------------------------------------
1         2001:db8::2:2             2001:db8:aaaa:2:7fff:f000:  No        0
                                    :
 BUM
1         2001:db8::2:2             2001:db8:aaaa:2:8000::      No        1
 -
1         2001:db8::2:3             2001:db8:aaaa:3:7fff:f000:  No        0
                                    :
 BUM
1         2001:db8::2:3             2001:db8:aaaa:3:8000::      No        1
 -
-------------------------------------------------------------------------------
Number of TEP, SID: 4
-------------------------------------------------------------------------------
===============================================================================
---snip---
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The **show service id 12 fdb expiry** command on PE-1 shows that PE-1 learns one MAC address locally, while PE-1 learns two remote MAC addresses via BGP EVPN.

```
A:admin@PE-1# show service id 12 fdb expiry

===============================================================================
Forwarding Database, Service 12
===============================================================================
---snip---
Table Size        : 250                 Allocated Count   : 3
Total In Use      : 3
Learned Count     : 1                   Static Count      : 0
---snip---
BGP EVPN Count    : 2                   EVPN Static Cnt   : 0
---snip---
Remote Age        : 900                 Local Age         : 300
---snip---
Mac Learning      : Enabled             Discard Unknown   : Disabled
Mac Aging         : Enabled             Relearn Only      : False
---snip---
===============================================================================
```

The locally learned MAC address belongs to the originator of the **ping** commands in the VPRN 11 context on PE-1, while the BGP EVPN learned MAC addresses belong to the destinations for those **ping** commands, which are in the VPRN 11 context on PE-2 and in the VPRN 11 context on PE-3 respectively. The Transport:Tnl-Id (for example 2001:db8:aaaa:2:8000::) indicates that PE-1 transports frames to the destination (on or connected to PE-2) via the SRv6 full SID to PE-2 for the End.DT2U behavior. The VPLS FDB entries that PE-1 learns locally expire after 300 seconds. The removal of a locally learned entry from the local VPLS FDB triggers the removal of the corresponding BGP EVPN learned entries in the remote VPLS FDBs. Similar output applies for the **ping** commands in the VPRN 11 context on PE-2 and for the **ping** commands in the VPRN 11 context on PE-3.

```
A:admin@PE-1# show service id 12 fdb detail

===============================================================================
Forwarding Database, Service 12
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
12         00:00:5e:00:53:01 sap:1/1/c1/1:11         L/30     01/04/23 17:01:49
12         00:00:5e:00:53:02 srv6-1:                 Evpn     01/04/23 17:01:49
                             2001:db8::2:2
           2001:db8:aaaa:2:8000::
12         00:00:5e:00:53:03 srv6-1:                 Evpn     01/04/23 17:01:58
                             2001:db8::2:3
           2001:db8:aaaa:3:8000::
-------------------------------------------------------------------------------
No. of MAC Entries: 3
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

Next to the locally learned MAC address for the locally configured interface, VPRN 11 on PE-1 has two dynamically learned MAC addresses, one for each of the BGP EVPN learned MAC addresses. Similar output applies for VPRN 11 on PE-2 and for VPRN 11 on PE-3.

```
A:admin@PE-1# show router 11 arp summary
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===========================================================
ARP Table Summary (Service: 11)
===========================================================
Local ARP Entries    : 1
---snip---
Dynamic ARP Entries  : 2
---snip---
-----------------------------------------------------------
No. of ARP Entries   : 3
===========================================================
```

The **show router 11 arp** command on PE-1 shows the association between the IP address and the MAC address, and the interface that the MAC address belongs to. The MAC address for the remote interface to the EVPN-enabled VPLS service corresponds with that of the SAP that is configured for it in VPRN 11 on PE-2 and VPRN 11 on PE-3. The association between the IP address and the MAC address for dynamically learned remote MAC addresses expires after 4 hours. Similar output applies for PE-2 and for PE-3.

```
A:admin@PE-1# show router 11 arp

===============================================================================
ARP Table (Service: 11)
===============================================================================
IP Address      MAC Address       Expiry    Type   Interface
-------------------------------------------------------------------------------
172.16.0.1      00:00:5e:00:53:01 00h00m00s Oth[I] local
172.16.0.2      00:00:5e:00:53:02 03h58m56s Dyn[I] local
172.16.0.3      00:00:5e:00:53:03 03h58m51s Dyn[I] local
-------------------------------------------------------------------------------
No. of ARP Entries: 3
===============================================================================
```

The **show router bgp routes evpn mac hunt** command shows a consolidated view on the MAC routes for the EVPN family. On PE-1, in the RIB In Entries section, it shows for each learned next hop how PE-1 must handle the IPv4 and IPv6 unicast data destined for it and where PE-1 must send it to. In the RIB Out Entries section, it shows for each local next hop how PE-1 expects the remote routers to handle the IPv4 and IPv6 unicast data destined for it and where PE-1 expects that data. Similar output applies for PE-2 and for PE-3.

```
A:admin@PE-1# show router bgp routes evpn mac hunt
===============================================================================
 BGP Router ID:192.0.2.1        AS:64500       Local AS:64500
===============================================================================
---snip---
===============================================================================
BGP EVPN MAC Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : n/a
Nexthop        : 2001:db8::2:2
Path Id        : None
From           : 2001:db8::2:4
Res. Nexthop   : fe80::60e:1ff:fe01:1
Local Pref.    : 100                    Interface Name : int-PE-1-PE-2
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : 10
Connector      : None
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Community      : target:64500:1
Cluster        : 4.4.4.4
Originator Id  : 192.0.2.2              Peer Router Id : 192.0.2.4
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
EVPN type      : MAC
ESI            : ESI-0
Tag            : 0
IP Address     : n/a
Route Dist.    : 192.0.2.2:1
Mac Address    : 00:00:5e:00:53:02
MPLS Label1    : LABEL 524288           MPLS Label2    : n/a
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
Add Paths Send : Default
Last Modified  : 00h01m02s
SRv6 TLV Type  : SRv6 L2 Service TLV (6)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:aaaa:2::
Full Sid       : 2001:db8:aaaa:2:8000::
Behavior       : End.DT2U (23)
SRv6 SubSubTLV : SRv6 SID Structure (1)
Loc-Block-Len  : 48                     Loc-Node-Len   : 16
Func-Len       : 20                     Arg-Len        : 0
Tpose-Len      : 20                     Tpose-offset   : 64

Network        : n/a
Nexthop        : 2001:db8::2:3
---snip---

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
Network        : n/a
Nexthop        : 2001:db8::2:1
Path Id        : None
To             : 2001:db8::2:4
Res. Nexthop   : n/a
Local Pref.    : 100                    Interface Name : NotAvailable
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : n/a
Connector      : None
Community      : target:64500:1
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.4
Origin         : IGP
AS-Path        : No As-Path
EVPN type      : MAC
ESI            : ESI-0
Tag            : 0
IP Address     : n/a
Route Dist.    : 192.0.2.1:1
Mac Address    : 00:00:5e:00:53:01
MPLS Label1    : 524288                 MPLS Label2    : n/a
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
SRv6 TLV Type  : SRv6 L2 Service TLV (6)
SRv6 SubTLV    : SRv6 SID Information (1)
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Sid             : 2001:db8:aaaa:1::
Full Sid        : 2001:db8:aaaa:1:8000::
Behavior        : End.DT2U (23)
SRv6 SubSubTLV : SRv6 SID Structure (1)
Loc-Block-Len  : 48                    Loc-Node-Len   : 16
Func-Len       : 20                    Arg-Len        : 0
Tpose-Len      : 20                    Tpose-offset   : 64

-------------------------------------------------------------------------------
Routes : 3
===============================================================================
```

Conclusion

Distributed EVPN-enabled VPLS services can be transported over SRv6 tunnels that are automatically set up between PEs. PEs attached to the same EVPN-enabled VPLS service exchange EVPN IMET routes and MAC/IP advertisement routes that contain the SRv6 SIDs. Those SRv6 SIDs are required so that PEs can create SRv6 destinations to send unicast and BUM traffic to the other PEs in the service.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

EVPN ESI Type 1

This chapter provides information about EVPN ESI Type 1.

Topics in this chapter include:

- Applicability
- Overview
- Configuration
- Conclusion

Applicability

The information and configuration in this chapter are based on SR OS Release 22.5.R1.

Overview

In SR OS releases earlier than 21.5.R1, the 10-byte Ethernet Segment Identifier (ESI) can only be configured manually; the auto-derived EVPN ESI type 1 (as per RFC 7432) is supported in SR OS Release 21.5.R1 and later. The **auto-esi** command is used to configure the ESI mode.

```
*[ex:/configure service system bgp evpn ethernet-segment "ESI-23"]
A:admin@PE-2# auto-esi ?

 auto-esi <keyword>
 <keyword>  - (none|type-1)
 Default    - none

    EVPN Ethernet segment auto-ESI type

    Warning: Modifying this element toggles
    'configure service system bgp evpn ethernet-segment "ESI-23" admin-state'
    automatically for the new value to take effect.
```

The default **auto-esi** value is **none**, which forces the user to configure the 10-byte ESI manually. When **type-1** is configured, a manual ESI cannot be configured and the ESI is auto-derived, as per RFC 7432.

ESI type 1 is auto-derived from the CE's Link Aggregation Control Protocol (LACP) system MAC address and port key. Figure 84: ESI type 1 example shows an example of ESI type 1 for LACP system MAC address 00:00:5e:00:53:00 and administrative key 257 (= 0x0101).

Figure 84: ESI type 1 example

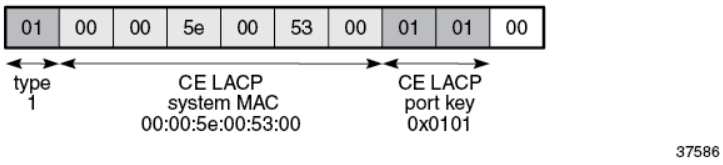

RFC 7432, section "Ethernet Segment", defines ESI type 1 as follows:

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- Type 0x01 (byte 0)
- CE LACP system MAC address (bytes 1 through 6); for example, 00:00:5e:00:53:00
- CE LACP port key (bytes 7 and 8); for example, 0x0101
- 0x00 (byte 9 must be zero)

As per RFC 7432, this mechanism can only be used if the ESIs are unique, so the CE LACP system MAC and LACP port key combinations must be unique in the network.

Figure 85: ESI auto-configuration example shows the example where CE-1 has LACP system MAC address 00:00:5e:00:53:00 and LACP port key 257 (= 0x0101). CE-1 sends Link Aggregation Control Protocol Data Units (LACPDUs) to PE-2 and PE-3 with these values. Both PE-2 and PE-3 use ESI 01:00:00:5e:00:53:00:01:01:00 in ES "ESI-23". This applies both to all-active and to single-active ESs.

Figure 85: ESI auto-configuration example

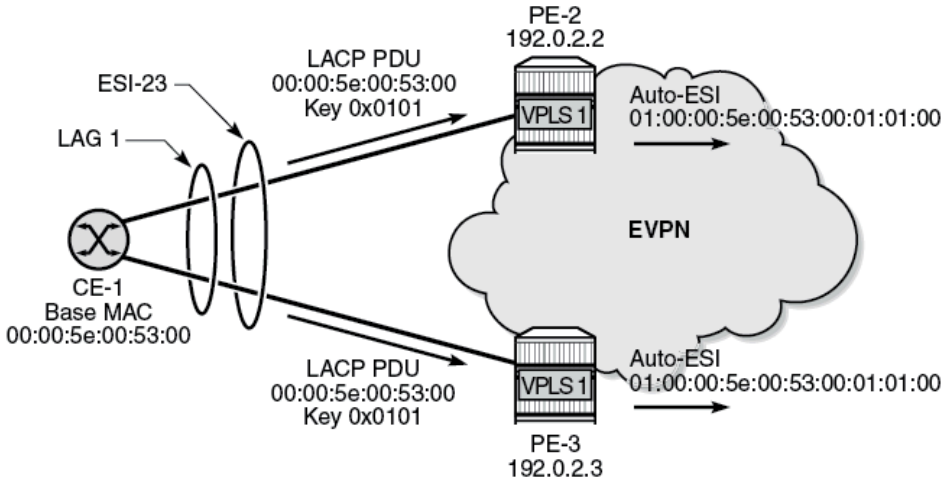

The CE treats both PE-2 and PE-3 as the same switch. This allows the CE to aggregate links that are attached to different PEs in the same bundle.

When the ES LAG goes operationally down, due to the ports going down or LACP going down or standby, the previously auto-derived ESI is retained. However, when the LACP information on the CE is changed, such as a different LACP port key, the ES goes down and a new ESI will be generated.

The all-active ES "AA-ESI-23" with ESI type 1 is configured as follows:

```
# on PE-2, PE-3:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "AA-ESI-23" {
                        admin-state enable
                        multi-homing-mode all-active
                        auto-esi type-1
                        association {
                            lag "lag-1" {
                            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        }
                    }
```

The following restrictions apply for ESI type 1:

- ESI type 1 is only supported on non-virtual (regular) ESs. The following error message is raised when attempting to configure **auto-esi type-1** for a virtual ES:

```
*[ex:/configure service system bgp evpn ethernet-segment "vESI-23"]
A:admin@PE-2# commit
MINOR: SVCMGR #1003: configure service system bgp evpn ethernet-segment "vESI-23" auto-esi -
 Inconsistent value - not supported along with virtual ethernet-segment
```

- ESI type 1 is not supported in ESs with associations other than LAG:

```
*[ex:/configure service system bgp evpn ethernet-segment "ESI-23" association port 1/2/1]
A:admin@PE-2# commit
MINOR: SVCMGR #1003: configure service system bgp evpn ethernet-segment "ESI-23" auto-esi
 - Inconsistent value - not supported with association - configure service system bgp evpn
 ethernet-segment "ESI-23"


*[ex:/configure service system bgp evpn ethernet-segment "ESI-23" association sdp 24]
A:admin@PE-2# commit
MINOR: SVCMGR #1003: configure service system bgp evpn ethernet-segment "ESI-23" auto-esi
 - Inconsistent value - not supported with association - configure service system bgp evpn
 ethernet-segment "ESI-23"
```

- An ES with ESI type 1 can only be enabled if the LAG has LACP enabled:

```
*[ex:/configure service system bgp evpn ethernet-segment "ESI-23" association lag "lag-4"]
A:admin@PE-2# commit
MINOR: MGMT_CORE #4001: configure service system bgp evpn ethernet-segment "ESI-23" auto-esi
 - lacp needs to be enabled on lag for auto-esi type 1 - configure lag "lag-4" lacp
```

- ESI type 1 is allowed with all-active and single-active ESs. When used in single-active mode, the CE must use a single LAG to connect to the multi-homed PEs.
- It is not possible to manually configure an ESI when **auto-esi type-1** is configured:

```
*[ex:/configure service system bgp evpn ethernet-segment "ESI-23"]
A:admin@PE-2# auto-esi type-1

*[ex:/configure service system bgp evpn ethernet-segment "ESI-23"]
A:admin@PE-2# esi 01:00:00:00:00:23:00:00:00:01

*[ex:/configure service system bgp evpn ethernet-segment "ESI-23"]
A:admin@PE-2# commit
MINOR: SVCMGR #1003: configure service system bgp evpn ethernet-segment "ESI-23" auto-esi -
 Inconsistent value - not supported along with esi configuration
```

- An ES with a manually configured ESI cannot be created with the same ESI value as the auto-derived ESI type 1 in another ES.

```
*[ex:/configure service system bgp evpn ethernet-segment "AA-ESI-23-5"]
A:admin@PE-2# esi 01:00:00:5e:00:53:00:01:01:00

*[ex:/configure service system bgp evpn ethernet-segment "AA-ESI-23-5"]
A:admin@PE-2# commit
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
MINOR: SVCMGR #8047: configure service system bgp evpn ethernet-segment "AA-ESI-23-5" -
 Ethernet segment id is not valid - ESI already in use by another ethernet segment
```

- If an ES with manual ESI is active and another ES is configured with an auto-derived ESI with the same value as the manual ESI, the auto-ESI value is deleted, and a log event is added to log "99":

```
# in log "99":
110 2022/05/25 15:28:44.361 CEST MINOR: SVCMGR #2610 Base
"The Auto Ethernet segment identifier type-1 has been deleted for Ethernet Segment AA-ESI-23
 because the new ID 01:00:00:5e:00:53:00:01:01:00 conflicts with ES AA-ESI-23-5"
```

Configuration

In this section, ESI type 1 is configured in the following use cases:

- ESI type 1 in all-active ESs
- ESI type 1 in single-active ESs

Figure 86: Example topology shows the example topology with four PEs and two CEs. CE-1 is connected via LAG 1 to the all-active ES "AA-ESI-23" on PE-2 and PE-3; CE-6 is connected via LAG-2 to the single-active ES "SA-ESI-45" on PE-4 and PE-5. In this example, an EVPN-MPLS VPLS is configured, but other services are also supported.

Figure 86: Example topology

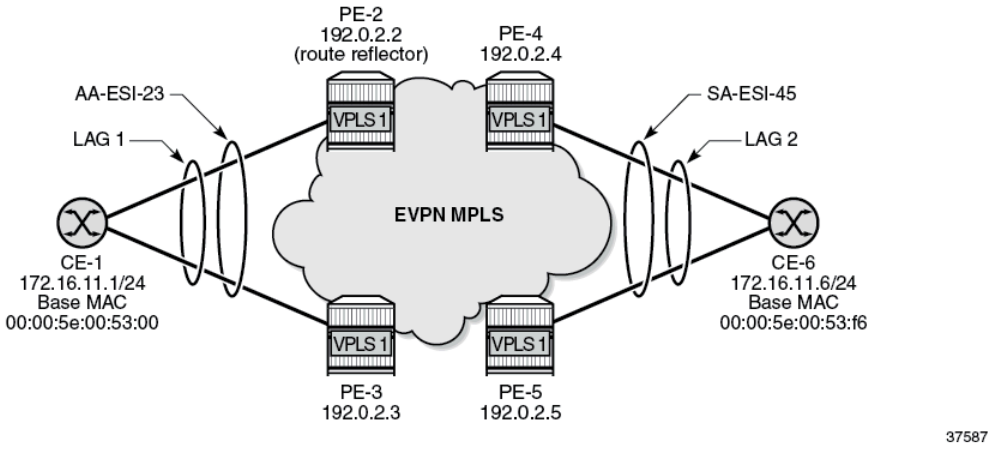

The initial configuration includes:

- cards, MDAs, ports
- on PEs: router interfaces, IS-IS, LDP

On the PEs, BGP is configured for the EVPN address family. PE-2 acts as the route reflector with the following configuration:

```
# on PE-2:
configure {
    router "Base" {
        autonomous-system 64500
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            rapid-update {
                evpn true
            }
            group "internal" {
                peer-as 64500
                family {
                    evpn true
                }
                cluster {
                    cluster-id 1.1.1.1
                }
            }
            neighbor "192.0.2.3" {
                group "internal"
            }
            neighbor "192.0.2.4" {
                group "internal"
            }
            neighbor "192.0.2.5" {
                group "internal"
            }
        }
```

On CE-1, LAG 1 is configured with LACP enabled and administrative key 257, as follows:

```
# on CE-1:
configure {
    lag "lag-1" {
        admin-state enable
        mode hybrid
        max-ports 64
        lacp {
            mode active
            administrative-key 257
        }
        port 1/1/1 {
        }
        port 1/1/2 {
        }
```

The LACP system MAC address of CE-1 can be retrieved with the following command:

```
[/]
A:admin@CE-1# show chassis | match MAC
  Base MAC address                  : 00:00:5e:00:53:00
```

ESI type 1 in all-active ESs

On PE-2 and PE-3, the all-active ES "AA-ESI-23" is configured with **auto-esi type-1** and LAG 1:

```
# on PE-2, PE-3:
configure {
    service {
        system {
            bgp {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                evpn {
                    ethernet-segment "AA-ESI-23" {
                        admin-state enable
                        multi-homing-mode all-active
                        auto-esi type-1
                        association {
                            lag "lag-1" {
                            }
                        }
                    }
```

The EVPN-MPLS VPLS 1 is configured as follows:

```
# on PE-2, PE-3:
configure {
    service {
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
            }
            bgp-evpn {
                evi 1
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    ecmp 2
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            sap lag-1:1 {
            }
        }
```

The operational ESI on PE-2 is 01:00:00:5e:00:53:00:01:01:00 for CE LACP system MAC address 00:00:5e:00:53:00 and administrative key 0x0101, as can be verified with the following command:

```
[/]
A:admin@PE-2# show service system bgp-evpn ethernet-segment name "AA-ESI-23"

===============================================================================
Service Ethernet Segment
===============================================================================
Name                    : AA-ESI-23
Eth Seg Type            : None
Admin State             : Enabled            Oper State         : Up
ESI                     : auto-esi
Oper ESI                : 01:00:00:5e:00:53:00:01:01:00
Auto-ESI Type           : Type 1
AC DF Capability        : Include
Multi-homing            : allActive          Oper Multi-homing  : allActive
ES SHG Label            : 524279
Source BMAC LSB         : None
Lag Id                  : 1
ES Activation Timer     : 3 secs (default)
Oper Group              : (Not Specified)
Svc Carving             : auto               Oper Svc Carving   : auto
Cfg Range Type          : primary
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

This output is slightly different for a manually configured ES, as follows:

```
# on PE-2, PE-3:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "AA-ESI-23-5" {
                        admin-state enable
                        esi 01:00:00:00:00:23:05:00:00:01
                        multi-homing-mode all-active
                        association {
                            lag "lag-5" {
                            }
                        }
                    }
                }
            }
```

```
[/]
A:admin@PE-2# show service system bgp-evpn ethernet-segment name "AA-ESI-23-5"

===============================================================================
Service Ethernet Segment
===============================================================================
Name                    : AA-ESI-23-5
Eth Seg Type            : None
Admin State             : Enabled           Oper State         : Up
ESI                     : 01:00:00:00:00:23:05:00:00:01
Oper ESI                : 01:00:00:00:00:23:05:00:00:01
Auto-ESI Type           : None
AC DF Capability        : Include
Multi-homing            : allActive         Oper Multi-homing  : allActive
ES SHG Label            : 524278
Source BMAC LSB         : None
Lag Id                  : 5
ES Activation Timer     : 3 secs (default)
Oper Group              : (Not Specified)
Svc Carving             : auto              Oper Svc Carving   : auto
Cfg Range Type          : primary
===============================================================================
```

ESI type 1 in single-active ESs

CE-6 is connected via LAG 2 to the single-active ES "SA-ESI-45" on PE-4 and PE-5. An ES operational group and LAG monitor operational group is required in this use case.

On CE-6, LAG 2 is configured with LACP enabled and administrative key 32768 (= 0x8000), as follows:

```
# on CE-6:
configure {
    lag "lag-2" {
        admin-state enable
        mode hybrid
        max-ports 64
        lacp {
            mode active
            administrative-key 32768
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        port 1/1/1 {
        }
        port 1/1/2 {
        }
    }
```

The LACP system MAC address of CE-6 is the following:

```
[/]
A:admin@CE-6# show chassis | match MAC
  Base MAC address                  : 00:00:5e:00:53:f6
```

On PE-4 and PE-5, operational group "op-grp-2" is configured and assigned to single-active ES "SA-ESI-45".

Note: When an operational group is associated to an ES, the hold timers for the operational group must be zero (the default value for the group down timer).

LAG 2 monitors this operational group. The configuration is as follows:

```
# on PE-4:
configure {
    lag "lag-2" {
        admin-state enable
        encap-type dot1q
        mode access
        monitor-oper-group "op-grp-2"
        max-ports 64
        lacp {
            mode active
            system-id 00:00:00:00:45:02
            administrative-key 1
        }
        port 1/1/1 {
        }
    }
    service {
        oper-group "op-grp-2" {
            hold-time {
             ## down    # default 0
                up 0
            }
        }
        system {
            bgp {
                evpn {
                    ethernet-segment "SA-ESI-45" {
                        admin-state enable
                        multi-homing-mode single-active
                        oper-group "op-grp-2"
                        auto-esi type-1
                        ac-df-capability exclude
                            service-carving-mode manual     # required for oper-group
                            manual {
                                preference {
                                    mode non-revertive
                                    value 200
                                }
                            }
                        }
                        association {
                            lag "lag-2" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        }
                    }
                }
            }
        }
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
            }
            bgp-evpn {
                evi 1
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    ecmp 2
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            sap lag-2:1 {
            }
        }
```

The following command on Designated Forwarder (DF) PE-4 shows that the operational ESI is 01:00:00:5e:00:53:f6:80:00:00:

```
# [/]
A:admin@PE-4# show service system bgp-evpn ethernet-segment name "SA-ESI-45" all

===============================================================================
Service Ethernet Segment
===============================================================================
Name                    : SA-ESI-45
Eth Seg Type            : None
Admin State             : Enabled            Oper State         : Up
ESI                     : auto-esi
Oper ESI                : 01:00:00:5e:00:53:f6:80:00:00
Auto-ESI Type           : Type 1
AC DF Capability        : Exclude
Multi-homing            : singleActive       Oper Multi-homing  : singleActive
ES SHG Label            : 524281
Source BMAC LSB         : None
Lag                     : lag-2
ES Activation Timer     : 3 secs (default)
Oper Group              : op-grp-2
Svc Carving             : manual             Oper Svc Carving   : manual
Cfg Range Type          : lowest-pref

-------------------------------------------------------------------------------
DF Pref Election Information
-------------------------------------------------------------------------------
Preference     Preference     Last Admin Change        Oper Pref      Do No
Mode           Value                                   Value          Preempt
-------------------------------------------------------------------------------
non-revertive  200            05/25/2022 15:14:19      200            Enabled
-------------------------------------------------------------------------------
EVI Ranges: <none>
ISID Ranges: <none>
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
EVI Information
===============================================================================
EVI                 SvcId               Actv Timer Rem      DF
-------------------------------------------------------------------------------
1                   1                   0                   yes
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================

-------------------------------------------------------------------------------
DF Candidate list
-------------------------------------------------------------------------------
EVI                                     DF Address
-------------------------------------------------------------------------------
1                                       192.0.2.4
1                                       192.0.2.5
-------------------------------------------------------------------------------
Number of entries: 2
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
---snip---
```

The operational ESI on Non-Designated Forwarder (NDF) PE-5 is the same as for PE-4.

The operational status of the operational group "op-grp-2" on DF PE-4 is up, while it is down on NDF PE-5 where the ES is inactive, as follows:

```
[/]
A:admin@PE-4# show service oper-group "op-grp-2"

===============================================================================
Service Oper Group Information
===============================================================================
Oper Group       : op-grp-2
Creation Origin  : manual                          Oper Status: up
Hold DownTime    : 0 secs                          Hold UpTime: 0 secs
Members          : 1                               Monitoring : 1
===============================================================================

[/]
A:admin@PE-5# show service oper-group "op-grp-2" detail

===============================================================================
Service Oper Group Information
===============================================================================
Oper Group       : op-grp-2
Creation Origin  : manual                          Oper Status: down
Hold DownTime    : 0 secs                          Hold UpTime: 0 secs
Members          : 1                               Monitoring : 1
===============================================================================

===============================================================================
Member Ethernet-Segment for OperGroup: op-grp-2
===============================================================================
Ethernet-Segment                        Status
-------------------------------------------------------------------------------
SA-ESI-45                               Inactive
-------------------------------------------------------------------------------
Ethernet-Segment Entries found: 1
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Monitoring LAG for OperGroup: op-grp-2
===============================================================================
Lag-id         Adm       Opr       Weighted  Threshold Up-Count  Act/Stdby
    name
-------------------------------------------------------------------------------
2              up        down      No        0         0         N/A
    lag-2
-------------------------------------------------------------------------------
LAG Entries found: 1
===============================================================================
port option not supported with monitoring
```

LAG 2 monitors the operational group "op-grp-2", so it follows the state of the ES "SA-ESI-45". On DF PE-4, LAG 2 is operationally up:

```
[/]
A:admin@PE-4# show lag "lag-2"

===============================================================================
Lag Data
===============================================================================
Lag-id         Adm     Opr     Weighted Threshold Up-Count MC Act/Stdby
    name
-------------------------------------------------------------------------------
2              up      up      No       0         1        N/A
    lag-2
===============================================================================
```

On NDF PE-5, LAG 2 is operationally down with reason operGroupDown:

```
[/]
A:admin@PE-5# show lag "lag-2" detail

===============================================================================
LAG Details
===============================================================================
Description        : N/A
-------------------------------------------------------------------------------
Details
-------------------------------------------------------------------------------
Lag-id              : 2                     Mode                : access
Lag-name            : lag-2
Adm                 : up                    Opr                 : down
Reason Down         : operGroupDown
Thres. Last Cleared : 05/25/2022 14:48:24   Thres. Exceeded Cnt  : 0
Dynamic Cost        : false                 Encap Type          : dot1q
Configured Address  : 02:1f:ff:00:01:42     Lag-IfIndex         : 1342177282
Hardware Address    : 02:1f:ff:00:01:42     Adapt Qos (access)  : distribute
Hold-time Down      : 0.0 sec               Port Type           : standard
Per-Link-Hash       : disabled
Include-Egr-Hash-Cfg: disabled              Forced              : -
Per FP Ing Queuing  : disabled              Per FP Egr Queuing  : disabled
Per FP SAP Instance : disabled
Access Bandwidth    : N/A                   Access Booking Factor: 100
Access Available BW : 0
Access Booked BW    : 0
LACP                : enabled               Mode                : active
LACP Transmit Intvl : fast                  LACP xmit stdby     : enabled
Selection Criteria  : highest-count         Slave-to-partner    : disabled
MUX control         : coupled
Subgrp hold time    : 0.0 sec               Remaining time      : 0.0 sec
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Subgrp selected     : 1                     Subgrp candidate     : -
Subgrp count        : 1
System Id           : 00:00:00:00:45:02     System Priority      : 32768
Admin Key           : 1                     Oper Key             : 1
Prtr System Id      : 00:00:5e:00:53:f6     Prtr System Priority : 32768
Prtr Oper Key       : 32768
Standby Signaling   : lacp
Port hashing        : port-speed            Port weight speed    : 0 gbps
Ports Up            : 0
Weights Up          : 0                     Hash-Weights Up      : 0
Monitor oper group  : op-grp-2
Oper group status   : down
Adaptive loadbal.   : disabled              Tolerance            : N/A

-------------------------------------------------------------------------------
Port-id        Adm     Act/Stdby Opr     Primary   Sub-group     Forced  Prio
-------------------------------------------------------------------------------
1/1/2          up      active    down    yes       1             -       32768

-------------------------------------------------------------------------------
Port-id        Role      Exp   Def   Dist  Col   Syn   Aggr  Timeout  Activity
-------------------------------------------------------------------------------
1/1/2          actor     No    No    No    No    No    Yes   Yes      Yes
1/1/2          partner   No    No    No    No    Yes   Yes   Yes      Yes
===============================================================================
```

When the LAG is operationally down, the SAP is operationally down. On DF PE-4, the SAP is up:

```
[/]
A:admin@PE-4# show service id 1 sap

===============================================================================
SAP(Summary), Service 1
===============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.  Egr.   Adm  Opr
                                           QoS   Fltr    QoS   Fltr
-------------------------------------------------------------------------------
lag-2:1                         1          1     none    1     none   Up   Up
-------------------------------------------------------------------------------
Number of SAPs : 1
-------------------------------------------------------------------------------
===============================================================================
```

On NDF PE-5, the SAP is operationally down:

```
[/]
A:admin@PE-5# show service id 1 sap lag-2:1

===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 1
SAP                : lag-2:1                 Encap            : q-tag
Description        : (Not Specified)
Admin State        : Up                      Oper State       : Down
Flags              : PortOperDown StandByForMHProtocol
Multi Svc Site     : None
Last Status Change : 05/25/2022 14:48:24
Last Mgmt Change   : 05/25/2022 15:14:27
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Auto-derived ESI changes when LACP port key on CE is modified

When the LAG goes operationally down due to ports going down or LACP going down, the auto-derived ESI is preserved. However, when the CE LACP configuration is changed— for example, with a different LACP port key—a new ESI is auto-derived.

In this example, the initial operational ESI on PE-4 is 01:00:00:5e:00:53:f6:80:00:00, as follows:

```
[/]
A:admin@PE-4# show service system bgp-evpn ethernet-segment name "SA-ESI-45" | match ESI
Name                    : SA-ESI-45
ESI                     : auto-esi
Oper ESI                : 01:00:00:5e:00:53:f6:80:00:00
Auto-ESI Type           : Type 1
```

On CE-6, the initial configuration of LAG 2 has LACP active with administrative key 32768:

```
[ex:/configure lag "lag-2"]
A:admin@CE-6# info
    admin-state enable
    mode hybrid
    max-ports 64
    lacp {
        mode active
        administrative-key 32768
    }
    port 1/1/1 {
    }
    port 1/1/2 {
    }
```

On CE-6, LAG 2 is reconfigured with administrative key 4095 (= 0x0fff), as follows:

```
# on CE-6:
configure {
    lag "lag-2" {
        admin-state enable
        mode hybrid
        max-ports 64
        lacp {
            mode active
            administrative-key 4095
        }
        port 1/1/1 {
        }
        port 1/1/2 {
        }
```

As a result, the operational ESI on PE-4 is 01:00:00:5e:00:53:f6:0f:ff:00, as follows:

```
[/]
A:admin@PE-4# show service system bgp-evpn ethernet-segment name "SA-ESI-45" | match ESI
Name                    : SA-ESI-45
ESI                     : auto-esi
Oper ESI                : 01:00:00:5e:00:53:f6:0f:ff:00
Auto-ESI Type           : Type 1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

When debugging is enabled for BGP updates, the following ES routes are seen: initially with ESI 01:00:00:5e:00:53:f6:**80:00**:00 and later with ESI 01:00:00:5e:00:53:f6:**0f:ff**:00, as follows:

```
24 2022/05/25 15:14:18.871 CEST MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 71
    Flag: 0x90 Type: 14 Len: 34 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.4
        Type: EVPN-ETH-SEG Len: 23 RD: 192.0.2.4:0 ESI: 01:00:00:5e:00:53:f6:80:00:00, IP-Len:
 4 Orig-IP-Addr: 192.0.2.4
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        df-election::DF-Type:Preference/DP:1/DF-Preference:200/AC:0
        target:00:00:5e:00:53:f6
"

---snip---

61 2022/05/25 15:23:01.331 CEST MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 71
    Flag: 0x90 Type: 14 Len: 34 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.4
        Type: EVPN-ETH-SEG Len: 23 RD: 192.0.2.4:0 ESI: 01:00:00:5e:00:53:f6:0f:ff:00, IP-Len:
 4 Orig-IP-Addr: 192.0.2.4
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        df-election::DF-Type:Preference/DP:1/DF-Preference:200/AC:0
        target:00:00:5e:00:53:f6
"
```

Conclusion

To simplify the configuration of single-active and all-active ESs with LAG association, ESI type 1 can be used to auto-derive the ESI from the CE's LACP system MAC address and LACP port key.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

EVPN for MPLS Tunnels

This chapter provides information about EVPN for MPLS tunnels.

Topics in this chapter include:

Applicability

This chapter was initially written for SR OS Release 13.0.R6, but the MD-CLI in the current edition corresponds to Release 21.2.R1. A prerequisite is to read the EVPN for VXLAN Tunnels (Layer 2) chapter.

Overview

EVPN-MPLS is standardized in RFC 7432, *BGP MPLS-Based Ethernet VPN*, as a Layer 2 VPN technology that can supplement VPLS for E-LAN services. Besides the optimizations introduced by EVPN, a significant number of service providers offering E-LAN services today are requesting EVPN for their multi-homing capabilities. EVPN supports all-active multi-homing (per-flow load-balancing multi-homing) as well as single-active multi-homing (per-service load-balancing multi-homing). In addition to those superior multi-homing capabilities, EVPN also provides a number of significant benefits, such as:

- IP-VPN-like operation and control for E-LAN services.
- Reduction and (in some cases) suppression of the Broadcast, Unknown unicast, and Multicast (BUM) traffic in the network.
- Simple provisioning and management.
- New set of tools to control the distribution of MAC addresses and Address Resolution Protocol (ARP) entries in the network.

The EVPN for Virtual eXtensible Local Area Network (VXLAN) tunnels (Layer 2) chapter focuses on the use of EVPN as a control plane for VXLAN tunnels, whereas this chapter provides configuration guidelines for EVPN when used for MPLS tunnels. Similar to EVPN-VXLAN services, VPLS services with EVPN for MPLS tunnels are referred to as EVPN-MPLS services.

As a reference, the EVPN route types and NLRIs (Network Layer Reachability Information messages) used by the EVPN family in RFC 7432 are shown in Figure 87: EVPN route types and NLRIs.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 87: EVPN route types and NLRIs

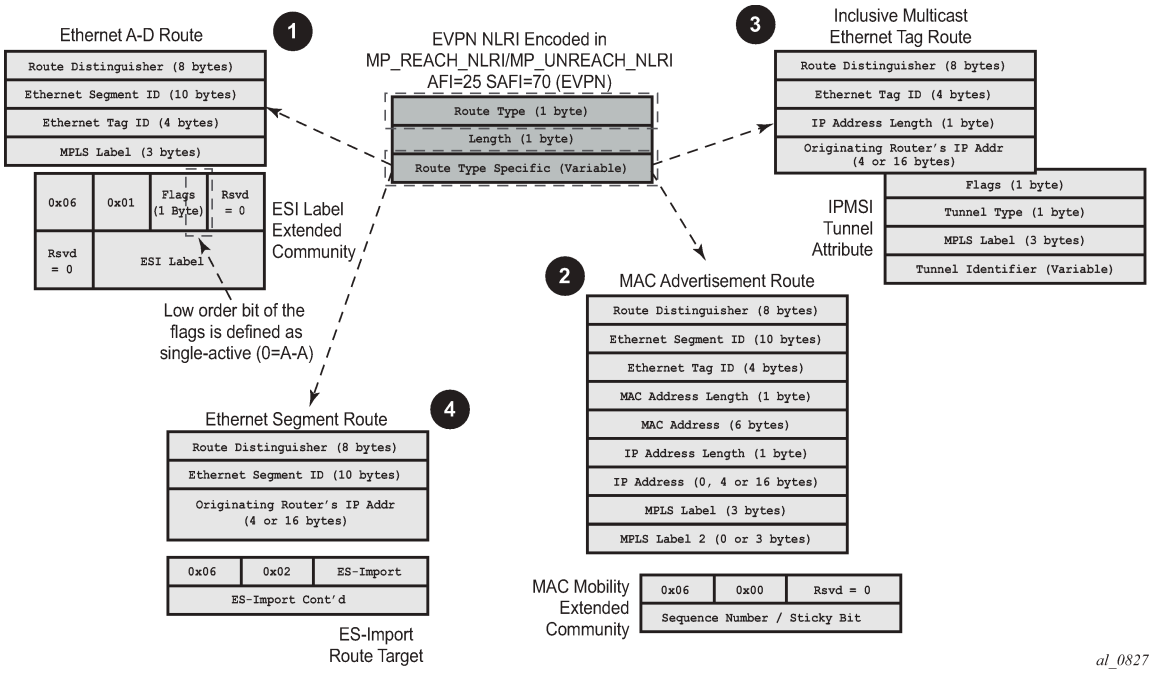

When no EVPN multi-homing is used in the network, only the base routes are used. Route types 2 and 3 are considered the base and mandatory routes:

- Route type 2 - MAC/IP route: This route advertises MAC addresses to be installed in the remote FDBs, or MAC/IP address pairs to be installed in the remote proxy-ARP/ND (Neighbor Discovery) tables.
- Route type 3 - Inclusive multicast route: This route advertises the multicast tree that the advertising PE intends to use for sending BUM traffic for an EVPN Instance (EVI). Ingress Replication, Point-to-multipoint multicast Label Distribution Protocol (P2MP mLDP), and composite tunnels are supported as tunnel types in route type 3 when BGP-EVPN MPLS is enabled. The ingress replication information, as well as the downstream MPLS label (for remote PEs to send BUM traffic to the advertising PE) are encoded in the Provider Multicast Service Interface Tunnel Attribute (PTA).

When EVPN multi-homing is used in an EVI, routes type 1 and 4 are used (where type 1 has two different purposes):

- Route type 1 - Auto-discovery per Ethernet segment (AD per ES) route: This route is advertised per ES from the PE, carries the Ethernet Segment Identifier (ESI) label (used for split-horizon) in multi-homing mode, and can affect procedures such as the Designated Forwarder (DF) election, as well as the aliasing/backup path/mass withdrawal on remote PEs.
- Route type 1 - Auto-discovery per EVPN instance (AD per-EVI) route: This route allows the remote PEs to provide aliasing and a backup path to the PEs part of the ES.
- Route type 4 - Ethernet Segment (ES) route: This route advertises a local configured ES. The exchange of this route can discover remote PEs that are part of the same ES and the DF election algorithm among them.

The AD per-EVI, MAC/IP, and inclusive multicast routes are considered service-level BGP-EVPN routes. Their RT/RD (Route-Target/Route-Distinguisher) are taken from the VPLS configuration.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The AD per-ES and the ES routes are considered base-level BGP-EVPN routes. However, their RT/RD are taken differently:

- The ES route RD is taken from the **service>system>bgp>evpn** configuration. The ES route RT is auto-derived from the Ethernet segment.
- The AD per-ES route RD is taken from the system level RD or service level RD. The RT extended community is taken from the service level RT or an RT set for the services defined on the Ethernet segment.

Configuration

This section describes the configuration of EVPN-MPLS for Layer 2 services on SR OS, as well as the available troubleshooting and show commands, and EVPN multi-homing.

Figure 88: EVPN-MPLS for VPLS services shows the topology used throughout this chapter. The network consists of a core with four EVPN PEs (PE-2, PE-3, PE-4, and PE-5) and two MTU devices that are dual-homed to the EVPN network. For MTU-1, all-active multi-homing is used, whereas MTU-6 is connected via single-active multi-homing to the EVPN network. Three CEs are connected to VPLS 1 in MTU-1, PE-3, and MTU-6 in order to test the connectivity.

Figure 88: EVPN-MPLS for VPLS services

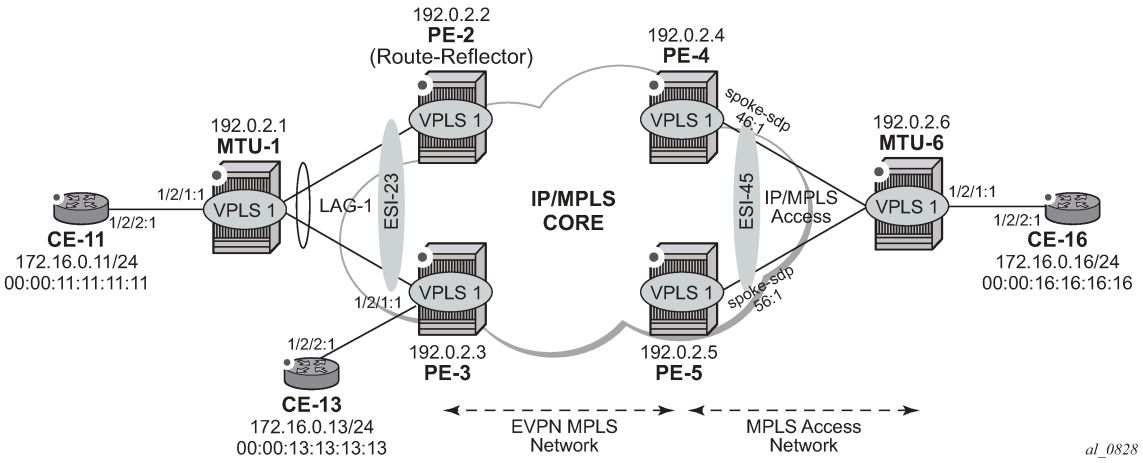

As part of the network infrastructure configuration, the following settings and protocols must be added to the configuration before starting with the EVPN-specific configuration for the services:

- The ports interconnecting the four PEs in the core are configured as network ports (or hybrid) and will have router network interfaces defined in them. The ports on PE-2/PE-3 connected to MTU-1 can be access or hybrid ports, whereas the ports on PE-4/PE-5 connected to MTU-6 can be network or hybrid ports. In case of hybrid ports, no LACP can be configured.
- The four PEs in the core (as well as MTU-6 in the access MPLS network) are running IS-IS and establishing point-to-point adjacencies for the exchange of the system IP addresses.
- LDP is used as the MPLS protocol to signal transport tunnel labels among PE-2, PE-3, PE-4, PE-5, and MTU-6. There is no LDP running between MTU-1 and the rest of the network, that is, MTU-1 is a pure Ethernet aggregation device.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- EVPN uses MP-BGP for exchanging reachability at service level. Therefore, BGP peering sessions must be established among the core PEs for the EVPN family. Although typically a separate router is used, in this chapter, PE-2 is used as BGP RR (route reflector) for EVPN routes. For example, the following output shows the configuration of BGP in the RR and one of the BGP clients. The relevant commands for EVPN are shown in bold.

  The configuration on the route reflector PE-2 is as follows:

```
# on RR PE-2:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            split-horizon true
            rapid-update {
                evpn true
            }
            group "internal" {
                peer-as 64500
                family {
                    evpn true
                }
                cluster {
                    cluster-id 1.1.1.1
                }
            }
            neighbor "192.0.2.3" {
                group "internal"
            }
            neighbor "192.0.2.4" {
                group "internal"
            }
            neighbor "192.0.2.5" {
                group "internal"
            }
        }
```

  The BGP configuration on the RR clients is as follows:

```
# on RR clients PE-3, PE-4, PE-5:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            split-horizon true
            rapid-update {
                evpn true
            }
            group "internal" {
                peer-as 64500
                family {
                    evpn true
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            neighbor "192.0.2.2" {
                group "internal"
            }
        }
```

Note:
The **def-recv-evpn-encap** command is not used in the preceding configuration because the default MPLS configuration is sufficient to have a correct interpretation of the received EVPN encapsulations.

The EVPN encapsulation type can be configured as MPLS or VXLAN in the general BGP configuration, the BGP group, or the BGP neighbor. For example, on PE-3, the EVPN encapsulation type for neighbor 192.0.2.2 can be configured using the following command:

```
[ex:configure router "Base" bgp neighbor "192.0.2.2"]
A:admin@PE-3# def-recv-evpn-encap

 def-recv-evpn-encap <keyword>
 <keyword> - (mpls|vxlan)

    Default EVPN encapsulation type
```

EVPN routes type 1 (auto-discovery per-EVI route), type 2 (MAC/IP route), type 3 (inclusive multicast route), and type 5 (IP-prefix route) are always sent with the RFC 5512, *the BGP Encapsulation Subsequent Address Family Identifier (SAFI) and the BGP Tunnel Encapsulation Attribute*, BGP encapsulation extended community that indicates the associated encapsulation of the route. Because the use of this extended community is not mandatory in RFC 7432, the **def-recv-evpn-encap** command indicates to the system what encapsulation is associated with routes received without any encapsulation. When interoperating with third-party EVPN vendors in mixed MPLS and EVPN-VXLAN networks, this command should be revised accordingly.

EVPN-MPLS configuration without multi-homing

After the base infrastructure (interfaces, IGP, LDP, BGP protocols) is configured, the service and EVPN can be enabled. When no multi-homing is used, the EVPN-MPLS configuration in a VPLS service looks similar to the configuration of EVPN-VXLAN for Layer 2, except for the commands related to the MPLS data plane. The following output shows the VPLS-1 configuration on PE-3 as an example:

```
# on PE-3:
configure {
    service {
        vpls "VPLS1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
            }
            bgp-evpn {
                evi 1
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    ecmp 2
                    auto-bind-tunnel {
                        resolution any
                    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
        }
        sap 1/2/1:1 {
        }
        sap lag-1:1 {
        }
    }
```

Where the following commands are relevant for a basic EVPN configuration:

- **bgp** enables the context for the BGP configuration relevant to the service. Two BGP instances can be configured, but only BGP instance 1 is configured in this example. If a manual (non-auto-derived) RD/RT, as well as import/export policies, are needed for the service, the commands in the **bgp** context must be configured. When **bgp-evpn** is enabled in a VPLS instance, other families are supported within the same service (bgp-ad, bgp-mh, bgp-vpls). This **bgp** context configures the common BGP parameters for all the BGP families in the service. Even if the general BGP parameters for the service are auto-derived (as in this example), the **bgp** context must be enabled.

```
[ex:configure service vpls "VPLS1"]
A:admin@PE-3# bgp ?

 [bgp-instance] <number>
 <number> - <1..2>

    BGP instance

[ex:configure service vpls "VPLS1"]
A:admin@PE-3# bgp 1 ?

 bgp

 apply-groups          - Apply a configuration group at this level
 apply-groups-exclude  - Exclude a configuration group at this level
 pw-template-binding   + Enter the pw-template-binding list instance
 route-distinguisher   - High-order 6 bytes that are used as string to compose
                         VSI-ID for use in NLRI
 route-target          + Enter the route-target context
 vsi-export            - VSI export policies
 vsi-import            - VSI import policies
```

- **bgp-evpn evi <1..65535>** — The EVPN instance or EVI is a 2-byte identifier used for the auto-derivation of the service RD, service RT, and for the service-carving algorithm when multi-homing is used. The EVI can be used for both **bgp-evpn vxlan** and **bgp-evpn mpls** when the user needs to auto-derive the RD and RT for the service. The auto-derivation is always based on:
  - RD system-ip:evi
  - RT autonomous-system:evi

  The configured and operating RD/RT values can be checked with the following show command (in this example, the evi value is 1):

```
[/]
A:admin@PE-3# show service id 1 bgp

===============================================================================
BGP Information
===============================================================================
Bgp Instance         : 1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Vsi-Import           : None
Vsi-Export           : None
Route Dist           : None
Oper Route Dist      : 192.0.2.3:1
Oper RD Type         : derivedEvi
Rte-Target Import    : None                Rte-Target Export: None
Oper RT Imp Origin   : derivedEvi          Oper RT Import   : 64500:1
Oper RT Exp Origin   : derivedEvi          Oper RT Export   : 64500:1

PW-Template Id       : None
-------------------------------------------------------------------------------
===============================================================================
```

Although not required for a basic BGP-EVPN MPLS configuration, some other parameters may be used at the **bgp-evpn** context level, when EVPN-MPLS services are deployed. Some examples are listed here:

- **bgp-evpn>routes>mac-ip>cfm-mac** must be enabled when eth-cfm is used across an EVPN-MPLS service among different PEs. If a Maintenance Endpoint (MEP) or Maintenance domain Intermediate Point (MIP) is configured in any of the SAP/SDP bindings in the VPLS and has to exchange eth-cfm packets with a remote MEP/MIP across the EVPN-MPLS core, this command must be enabled. In that way, the MEP/MIP MAC address can be advertised in EVPN (otherwise, the MEP/MIP MAC address would not be learned on remote EVPN-MPLS PEs and eth-cfm would not work correctly).
- **bgp-evpn>routes>mac-ip>advertise** and **bgp-evpn>mac-duplication** — See the EVPN for VXLAN Tunnels (Layer 2) chapter for a description of these two commands.
- **bgp-evpn>mpls <bgp-instance>** must be enabled.

When two BGP instances are added to a VPLS service, both BGP-EVPN MPLS and BGP-EVPN VXLAN can be configured at the same time in the service. A maximum of two BGP instances are supported in the same VPLS. In this chapter, only BGP instance 1 is used.

After the relevant **VPLS** parameters, **BGP** and **BGP-EVPN** attributes are added, the specific commands for **bgp-evpn mpls <bgp-instance>** can be configured as follows:

```
[ex:configure service vpls "VPLS1" bgp-evpn mpls 1]
A:admin@PE-3# info
    admin-state enable
    ingress-replication-bum-label true
    ecmp 2
    auto-bind-tunnel {
        resolution any
    }
```

- **ingress-replication-bum-label** controls whether the system will advertise different service labels for unicast and BUM traffic. If no EVPN multi-homing is configured in the network, this command can be disabled (**ingress-replication-bum-label false**) and the same MPLS label will be advertised for the unicast and BUM traffic for the VPLS instance. If EVPN multi-homing is configured in the PE, this command is strongly recommended to avoid potential transient issues. See the EVPN-MPLS multi-homing section.
- **ecmp** controls the number of remote PEs to which the local PE can load balance the unicast traffic. See the EVPN-MPLS multi-homing section.
- **auto-bind-tunnel** controls the resolution of EVPN destinations to MPLS transport tunnels. This command is also in VPRN services and works in the same way.
  - If the **auto-bind-tunnel resolution any** is configured, as in the example, EVPN destinations in the service are resolved based on the best tunnel in the Tunnel Table Manager (TTM). For instance, the following command shows the existing EVPN destinations for VPLS 1 in PE-3. The EVPN-MPLS

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

destination (Termination Endpoint (TEP) 192.0.2.2, label 524281) is resolved to LDP transport tunnel 65537 because the (best) LDP tunnel to 192.0.2.2 shown in the **show router tunnel-table** is LDP. If there was more than one tunnel type in the TTM to 192.0.2.2, the system would pick the lowest **Pref** (preference) tunnel.

```
[/]
A:admin@PE-3# show service id 1 evpn-mpls

===============================================================================
BGP EVPN-MPLS Dest
===============================================================================
TEP Address     Egr Label     Num. MACs   Mcast           Last Change
                Transport:Tnl                             Sup BCast Domain
-------------------------------------------------------------------------------
192.0.2.2       524281        0           bum             02/25/2021 16:46:38
                ldp:65537                                 No
192.0.2.4       524281        0           bum             02/25/2021 16:46:45
                ldp:65538                                 No
192.0.2.5       524281        0           bum             02/25/2021 16:46:52
                ldp:65539                                 No
-------------------------------------------------------------------------------
Number of entries : 3
-------------------------------------------------------------------------------
===============================================================================
---snip---
```

```
[/]
A:admin@PE-3# show router tunnel-table

===============================================================================
IPv4 Tunnel Table (Router: Base)
===============================================================================
Destination           Owner     Encap TunnelId  Pref   Nexthop        Metric
   Color
-------------------------------------------------------------------------------
192.0.2.2/32          ldp       MPLS  65537     9      192.168.23.1   10
192.0.2.4/32          ldp       MPLS  65538     9      192.168.34.2   10
192.0.2.5/32          ldp       MPLS  65539     9      192.168.35.2   10
192.0.2.6/32          ldp       MPLS  65540     9      192.168.34.2   20
-------------------------------------------------------------------------------
Flags: B = BGP or MPLS backup hop available
       L = Loop-Free Alternate (LFA) hop available
       E = Inactive best-external BGP route
       k = RIB-API or Forwarding Policy backup hop
===============================================================================
```

- If resolution is set to **any**, the following tunnel types are selected in order of preference: RSVP, LDP, Segment Routing, and BGP. The user can configure the preference of the segment-routing tunnel type in the TTM for a specific IGP instance.
- If one or more explicit tunnel types are specified using the resolution-filter option, then only these tunnel types will be selected again following the TTM preference.
- The user must set the resolution to filter to activate the list of tunnel-types configured under resolution-filter.

Although not shown in the **bgp-evpn mpls** basic configuration for PE-3, there are other parameters that can be modified:

```
[ex:/configure service vpls "VPLS1" bgp-evpn]
A:admin@PE-3# mpls 1 ?
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
mpls

admin-state           - Administrative state of BGP EVPN MPLS
apply-groups          - Apply a configuration group at this level
apply-groups-exclude  - Exclude a configuration group at this level
auto-bind-tunnel      + Enter the auto-bind-tunnel context
control-word          - Enable/disable setting the CW bit in the label message.
default-route-tag     - Default route tag
ecmp                  - Maximum ECMP routes information
entropy-label         - Enable/disable use of entropy-label.
fdb                   + Enter the fdb context
force-vc-forwarding   - VC forwarding action
ingress-replication-  - Use the same label as the one advertised for unicast traffic
 bum-label
oper-group            - Operational-Group identifier.
route-next-hop        + Enter the route-next-hop context
send-tunnel-encap     + Enter the send-tunnel-encap context
split-horizon-group   - Split horizon group
```

bgp <bgp>

defines the BGP instance: 1 or 2.

control-word

enables/disables the insertion of the control-word in the data path. The control-word is disabled by default and is not signaled in EVPN (based on RFC 7432) and has to be consistently configured in all the PEs in the network. The use of the **control-word** prevents packet reordering from happening in P routers that misinterpret the first nibble of the payload in the packets they receive. In some third-party EVPN vendors, the control-word is enabled by default, so it is recommended to enable it when interoperating with other vendors.

entropy-label

enables the use of entropy labels, as described in the *Entropy Label* chapter in the MPLS volume of the *7450 ESS, 7750 SR, and 7950 XRS Advanced Configuration Guide — Book I*.

force-vc-forwarding <vlan>

allows the system to preserve the VLAN ID and p-bits of the service-delimiting q-tag in a new tag added in the customer frame before sending it to the EVPN core. This command may be used with the **sap ingress qtag-manipulation** command: the configured translated VLAN ID will be sent to the EVPN binds, as opposed to the service-delimiting tag VLAN ID. If the ingress SAP/SDP-binding is null encapsulated, the output VLAN ID and p-bits will be zero.

fdb protected-src-mac-violation-action

is by default disabled. When enabled, all packets entering the object will be verified not to contain a protected source MAC address. In combination with the keyword discard, the packets that contain a protected MAC address will be discarded and an alarm is generated.

send-tunnel-encap

configures the encapsulation to be advertised with the EVPN routes for the service. The encapsulation is encoded in RFC 5512-based tunnel encapsulation extended communities. When configured in the **bgp-evpn>mpls** context, the supported options are none (delete send-tunnel-encap), mpls, mpls-over-udp, or both.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

admin-state

enables/disables the use of MPLS for EVPN. When enabled, a BGP route-refresh message is sent for the EVPN family.

split-horizon-group <*group-name*>

configures an explicit split-horizon-group (SHG) for all the EVPN destinations that can be shared with other SAP/SDP-bindings. See the VPLS to EVPN-MPLS integration section.

After **bgp-evpn mpls** is configured and enabled in the service, an inclusive multicast route is sent to the RR. The remote PEs receiving and importing that route will create an EVPN destination to the sending PE. An EVPN destination is identified by a TEP and MPLS label. Use the following show commands to view the service and the EVPN destinations created:

- **show service evpn-mpls**
- **show service id 1 evpn-mpls**
- **show service id 1 bgp-evpn**

An example of the output is shown for PE-2 when there is no traffic in the network. Therefore, only inclusive multicast routes have been exchanged among the four PEs.

```
[/]
A:admin@PE-2# show service evpn-mpls

===============================================================================
EVPN MPLS Tunnel Endpoints
===============================================================================
EvpnMplsTEP Address EVPN-MPLS Dest      ES Dest             ES BMac Dest
-------------------------------------------------------------------------------
192.0.2.3           1                   0                   0
192.0.2.4           1                   0                   0
192.0.2.5           1                   0                   0
-------------------------------------------------------------------------------
Number of EvpnMpls Tunnel Endpoints: 3
-------------------------------------------------------------------------------
===============================================================================
```

```
[/]
A:admin@PE-2# show service id 1 evpn-mpls

===============================================================================
BGP EVPN-MPLS Dest
===============================================================================
TEP Address     Egr Label     Num. MACs   Mcast           Last Change
                Transport:Tnl                             Sup BCast Domain
-------------------------------------------------------------------------------
192.0.2.3       524281        0           bum             02/25/2021 16:46:36
                ldp:65537                                 No
192.0.2.4       524281        0           bum             02/25/2021 16:46:45
                ldp:65538                                 No
192.0.2.5       524281        0           bum             02/25/2021 16:46:52
                ldp:65539                                 No
-------------------------------------------------------------------------------
Number of entries : 3
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN-MPLS Ethernet Segment Dest
===============================================================================
Eth SegId                       Num. Macs               Last Change
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================

===============================================================================
BGP EVPN-MPLS ES BMAC Dest
===============================================================================
ES BMAC Addr                            Last Change
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

```
[/]
A:admin@PE-2# show service id 1 bgp-evpn

===============================================================================
BGP EVPN Table
===============================================================================
MAC Advertisement  : Enabled            Unknown MAC Route  : Disabled
CFM MAC Advertise  : Disabled
Creation Origin    : manual
MAC Dup Detn Moves : 5                  MAC Dup Detn Window: 3
MAC Dup Detn Retry : 9                  Number of Dup MACs : 0
MAC Dup Detn BH    : Disabled
IP Route Advert    : Disabled
Sel Mcast Advert   : Disabled

EVI                : 1
Ing Rep Inc McastAd: Enabled
Accept IVPLS Flush : Disabled

-------------------------------------------------------------------------------
Detected Duplicate MAC Addresses             Time Detected
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN MPLS Information
===============================================================================
Admin Status       : Enabled            Bgp Instance       : 1
Force Vlan Fwding  : Disabled
Route NextHop Type : system-ipv4
Control Word       : Disabled
Max Ecmp Routes    : 2
Entropy Label      : Disabled
Default Route Tag  : none
Split Horizon Group: (Not Specified)
Ingress Rep BUM Lbl: Enabled
Ingress Ucast Lbl  : 524282             Ingress Mcast Lbl  : 524281
RestProtSrcMacAct  : none
Evpn Mpls Encap    : Enabled            Evpn MplsoUdp      : Disabled
Oper Group         :
===============================================================================

===============================================================================
BGP EVPN MPLS Auto Bind Tunnel Information
===============================================================================
Allow-Flex-Algo-Fallback  : false
Resolution                : any                Strict Tnl Tag   : false
Max Ecmp Routes           : 1
Bgp Instance              : 1
Filter Tunnel Types       : (Not Specified)
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
```

When traffic is generated, the PEs will start learning MAC addresses and advertising them in BGP so that the remote PEs learn those MAC addresses against EVPN destinations. For instance, when CE-13 sends traffic, PE-3 learns its MAC address and advertises it. The remote PEs (for instance, PE-2) will learn the MAC address and associate it with their EVPN destination to PE-3 (192.0.2.3:524282 in this example):

```
[/]
A:admin@PE-2# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
1          00:00:11:11:11:11 sap:lag-1:1             L/0      02/25/21 16:54:31
1          00:00:13:13:13:13 mpls:                   Evpn     02/25/21 16:54:31
                             192.0.2.3:524282
           ldp:65537
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

When the **ingress-replication-bum-label** is enabled in the PEs, the advertisement of MAC addresses will create new EVPN destinations, because the label is different from the one previously sent by the inclusive multicast route that created an EVPN destination. In the preceding example, when PE-3 advertises the CE-13 MAC address, PE-2 will create a new binding (see in the following output in bold) that shows one MAC address that is not Mcast (multicast) capable:

```
[/]
A:admin@PE-2# show service id 1 evpn-mpls

===============================================================================
BGP EVPN-MPLS Dest
===============================================================================
TEP Address     Egr Label     Num. MACs   Mcast           Last Change
                Transport:Tnl                           Sup BCast Domain
-------------------------------------------------------------------------------
192.0.2.3       524281        0           bum             02/25/2021 16:46:36
                ldp:65537                                 No
192.0.2.3       524282        1           none            02/25/2021 16:54:31
                ldp:65537                                 No
192.0.2.4       524281        0           bum             02/25/2021 16:46:45
                ldp:65538                                 No
192.0.2.5       524281        0           bum             02/25/2021 16:46:52
                ldp:65539                                 No
-------------------------------------------------------------------------------
Number of entries : 4
-------------------------------------------------------------------------------
===============================================================================
---snip---
```

When an EVPN-MPLS destination or MAC address is not created/installed correctly, the user may check the BGP-EVPN routes received and the routes kept in the RIB. The routes that the PE receives are shown

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

when **debug router bgp update** is enabled. These routes are shown even before any BGP processing is carried out.

```
# on PE-2:
30 2021/02/25 16:54:31.213 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 81
    Flag: 0x90 Type: 14 Len: 44 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.3
        Type: EVPN-MAC Len: 33 RD: 192.0.2.3:1 ESI: ESI-0, tag: 0, mac len: 48
                       mac: 00:00:13:13:13:13, IP len: 0, IP: NULL, label1: 8388512
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:1
        bgp-tunnel-encap:MPLS
"
```

```
[/]
A:admin@PE-2# show router bgp routes evpn mac
===============================================================================
 BGP Router ID:192.0.2.2        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN MAC Routes
===============================================================================
Flag  Route Dist.         MacAddr           ESI
      Tag                 Mac Mobility      Label1
                          Ip Address
                          NextHop
-------------------------------------------------------------------------------
u*>i  192.0.2.3:1         00:00:13:13:13:13 ESI-0
      0                   Seq:0             LABEL 524282
                          n/a
                          192.0.2.3

-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

If the route is successfully imported, it can be shown in the RIB (**show router bgp routes** commands). The route shown in the debug and the same route in a show command do not necessarily have the same label value. The reason for this expected mismatch is that the debug command shows the complete 24-bit field value because the route is shown before BGP can decide and decipher whether the label value is an MPLS label (high-order 20-bits of the label field) or a VNI (all 24 bits of the Label field for VXLAN). When the label in the debug command (8388512) is divided by 16 (2^4), the result is the MPLS label (524282), as follows: 8388512:16=524282.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

VPLS to EVPN-MPLS integration

The SR OS EVPN implementation supports RFC 8560, *(PBB-)EVPN Seamless Integration with (PBB-)EVPN*, so that EVPN-MPLS and VPLS can be integrated into the same network and within the same service.

The following behavior enables the integration of EVPN and SDP-bindings in the same VPLS network:

- Systems with EVPN endpoints and SDP-bindings to the same far-end bring down the SDP-bindings.
  - SR OS will allow the establishment of an EVPN destination and an SDP-binding to the same far-end but the SDP-binding will be kept operationally down. Only the EVPN endpoint will be operationally up. This is true for spoke-SDPs (manual and BGP-AD) and mesh-SDPs. It is also true between VXLAN and SDP-bindings.
  - If there is an EVPN endpoint to a specified far-end and a spoke-SDP establishment is attempted, the spoke-SDP will be set up but kept down with an operational flag indicating that there is an EVPN route to the same far-end.
  - If there is a spoke-SDP and a valid/used EVPN route arrives, the EVPN endpoint will be set up and the spoke-SDP will be brought down with an operational flag indicating that there is an EVPN route to the same far-end.
  - In the case of an SDP-binding and EVPN endpoint to different far-end IPs on the same remote PE, both links will be up. This can happen if the SDP-binding is terminated in an IPv6 address or IPv4 address different from the system address where the EVPN endpoint is terminated.

  The following example illustrates the preceding description. A spoke-SDP is added to the VPLS 1 configuration on PE-2:

```
# on PE-2:
configure {
    service {
        sdp 24 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.4
            }
        }
        vpls "VPLS1" {
            spoke-sdp 24:1 {
            }
        }
```

  The service configuration on PE-4 is as follows:

```
# on PE-4:
configure {
    service {
        sdp 42 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.2
            }
        }
        sdp 46 {
            admin-state enable
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.6
            }
        }
        vpls "VPLS1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
            }
            bgp-evpn {
                evi 1
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    ecmp 2
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            spoke-sdp 42:1 {
            }
            spoke-sdp 46:1 {
            }
        }
```

Spoke SDP 24:1 is operationally down, as can be verified as follows:

```
[/]
A:admin@PE-2# show service id 1 sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
24:1             Spok     192.0.2.4       Up      Down      524280    524279
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
```

Spoke SDP 24:1 is down because of an EVPN route conflict, as indicated by the flags:

```
[/]
A:admin@PE-2# show service id 1 sdp 24 detail | match Flag post-lines 1
Flags              : PWPeerFaultStatusBits
                     EvpnRouteConflict
```

- The user can add spoke-SDPs and all the EVPN-MPLS endpoints in the same SHG.
  - A CLI command exists in the **bgp-evpn>mpls** context so that the EVPN-MPLS endpoints can be added to an SHG.
  - The **bgp-evpn mpls split-horizon-group** must reference a user-configured SHG. User-configured SHGs can be configured within the service context.
  - The same group name can be associated with SAPs, spoke-SDPs, PW-templates, PW-template-bindings, and EVPN-MPLS endpoints.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

  - If the **split-horizon-group** command in **bgp-evpn>mpls** is not used, the default SHG (in which all the EVPN endpoints are) is still used, but it will not be possible to refer to it on SAPs/spoke-SDPs.
- The system disables the advertisement of MAC addresses learned on spoke- SDPs/SAPs that are part of an EVPN SHG.
  - When the SAPs or spoke-SDPs (manual or BGP-AD-discovered) are configured within the same SHG as the EVPN endpoints, MAC addresses will still be learned on them, but will not be advertised in EVPN.
  - The preceding statement is also true if proxy-ARP/ND is enabled and an IP-->MAC address pair is learned on a SAP/SDP-binding that belongs to the EVPN SHG.
  - The SAPs and/or spoke-SDPs added to an EVPN SHG should not be part of any EVPN multi-homed ES. If that happened, the PE would still advertise the AD per-EVI route for the SAP and/or spoke-SDP, attracting EVPN traffic that could not be forwarded to that SAP and/or SDP-binding.
  - Similar to the preceding statement, an SHG composed of SAPs/SDP-bindings used in a BGP-MH site should not be configured under **bgp-evpn>mpls>split-horizon-group**. This misconfiguration would prevent traffic being forwarded from the EVPN to the BGP-MH site, regardless of the DF/Non-DF state.

An example of a shared SHG configuration on PE-2 is as follows. Because the SAP and EVPN-MPLS are in the same SHG, no MAC addresses learned over SAP 1/2/1:2 will be advertised in EVPN (not even static MAC addresses).

```
# on PE-2:
configure {
    service {
        vpls "VPLS2" {
            admin-state enable
            service-id 2
            customer "1"
            bgp 1 {
            }
            bgp-evpn {
                evi 2
                mpls 1 {
                    admin-state enable
                    split-horizon-group "CORE"
                    ingress-replication-bum-label true
                    ecmp 2
                    auto-bind-tunnel
                        resolution any
                    }
                }
            }
            split-horizon-group "CORE" {
            }
            sap 1/2/1:2 {
                split-horizon-group "CORE"
            }
            sap lag-1:2 {
            }
```

EVPN-MPLS multi-homing

SR OS supports EVPN multi-homing as per RFC 7432.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The EVPN multi-homing implementation is based on the concept of the ES. An ES is a logical structure that can be defined in one or more PEs and identifies the CE (or access network) multi-homed to the EVPN PEs. An ES is associated with a port, LAG, or SDP object, and is shared by all the services defined on those objects.

Each ES has a unique identifier called ESI (Ethernet Segment Identifier) that is 10 bytes and is manually configured. The ESI is advertised in the control plane to all the PEs in an EVPN network; therefore, it is very important to ensure that the 10-byte ESI value is unique throughout the entire network. Single-homed CEs are assumed to be connected to an ES with ESI = 0 (single-homed ESs are not explicitly configured).

The ES is part of the base BGP-EVPN configuration and is not applied to any EVPN-MPLS service, by default. An ES can be shared by multiple services; the association of a specific SAP or spoke-SDP to an ES is automatically made when the SAP is defined in the same LAG or port configured in the ES, or when the spoke-SDP is defined in the same SDP configured in the ES. The following sections show the configuration of:

- an all-active multi-homing ES with a LAG associated with it
- a single-active multi-homing ES linked to an SDP

All-active multi-homing concepts

EVPN all-active multi-homing is built around three concepts: DF election, split-horizon (with an ESI-label), and aliasing, as shown in Figure 89: EVPN-MPLS all-active multi-homing concepts, from left to right.

Figure 89: EVPN-MPLS all-active multi-homing concepts

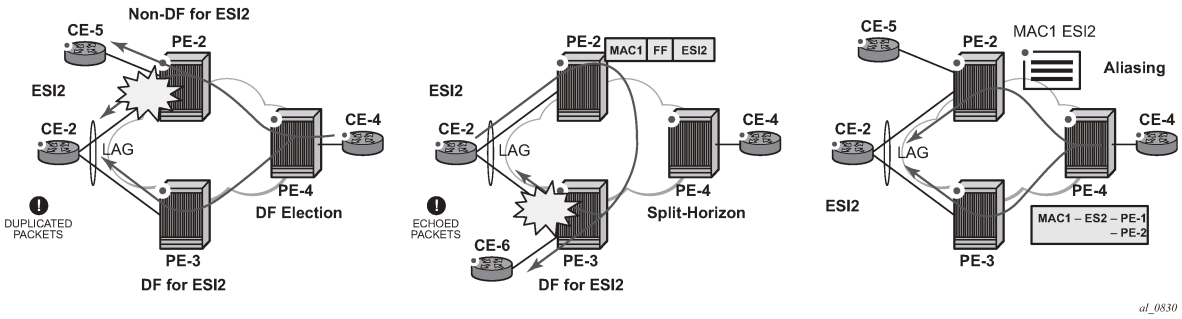

- With DF election, when PE-4 sends BUM traffic to the remote ES (CE-2), only one PE segment sends the BUM packets to the ES (PE-3 is the DF in the preceding example, and is elected to send BUM packets to CE-2). The non-DF, PE-2, removes the LAG SAP from the default multicast list (PE-2 does not bring CE-2 down, because it still needs to send upstream/downstream unicast traffic). PE-2 and PE-3 elect a DF for each service, based on the ES routes and the service-carving algorithm.
- With split-horizon, the PE part of the ES (PE-3 in the preceding example) identifies the BUM packets coming from the PE for the remote (PE-2), but within the same ES (ESI-2), and filters the packets so that they are not sent back to the ES, creating duplication. When PE-2 (non-DF) sends BUM traffic to PE-3 (DF), it uses a special MPLS label in the data path that PE-3 previously advertised for ESI-2 in an AD per-ES route. When PE-3 does an ingress lookup, it recognizes the ESI-label and filters the traffic (PE-3 still sends the BUM traffic to other SAPs/SDP-bindings).
- With aliasing, remote PEs that are not part of the ES can load-balance unicast traffic to all the PEs that are part of the ES, irrespective of from which PE a destination MAC address was learned. PE-4 will

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

create an EVPN destination to ESI-2 that will be resolved to the two next-hops: PE-2 and PE-3. Unicast load-balancing will happen as long as ECMP > 1 is enabled in PE-4.

Nokia recommends the use of **ingress-replication-bum-label** on the PEs that are part of an all-active ES. In an all-active multi-homing scenario, if a specified MAC address (for example, the CE-2 MAC address in the left-hand-side diagram), is not learned yet in a remote PE (for example, PE-4), but is known in the two PEs of the ES (for example, PE-2 and PE-3), the latter PEs might send duplicated packets to the CE.

This issue is solved by the use of **ingress-replication-bum-label** in PE-2 and PE-3. If configured, PE-2/PE-3 will know that the received packet is an unknown unicast packet; therefore, the Non-DF (PE-2) will not send the packet to CE-2 and there will not be duplication.

All-active multi-homing configuration

The all-active multi-homing configuration example is based on Figure 88: EVPN-MPLS for VPLS services.

MTU-1 is connected to the EVPN network using all-active multi-homing. According to RFC 7432, MTU-1 will be able to send traffic to both PEs for VPLS-1. Regular LAG load-balancing is used in MTU-1. Remote PEs such as PE-4 or PE-5 will be able to load-balance the unicast traffic to PE-2 and PE-3. PE-2 and PE-3 will discover that both are part of ESI-23 (due to the exchange of ES routes) and will elect a DF for VPLS-1. The non-DF for VPLS-1, in this case PE-2, will remove lag-1:1 from the VPLS-1 default multicast list. Also, when PE-2 and PE-3 send BUM traffic to each other, they will insert an ESI-label so that they can identify that the source of the BUM packet is ESI-23.

The following output shows the configuration of ESI-23 in PE-2 and PE-3, as well as the LAG interfaces for all-active multi-homing (see Figure 88: EVPN-MPLS for VPLS services). The configuration of LAG-1 in MTU-1 is also shown. Per RFC 7432, only a CE/MTU with a LAG can be connected to an all-active multi-homing ES. No other configuration is permitted on the CE for all-active multi-homing.

LAG 1 is configured on MTU-1, PE-2, and PE-3, as follows:

```
# on MTU-1:
configure {
    lag "lag-1" {
        admin-state enable
        encap-type dot1q
        mode access
        max-ports 64
        lacp {
            mode active
            administrative-key 32768
        }
        port 1/1/1 {
        }
        port 1/1/2 {
        }
    }
```

```
# on PE-2:
configure {
    lag "lag-1" {
        admin-state enable
        encap-type dot1q
        mode access
        max-ports 64
        lacp {
            mode active
            system-id 00:00:00:00:02:03
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            administrative-key 1
        }
        port 1/1/2 {
        }
    }
```

```
# on PE-3:
configure {
    lag "lag-1" {
        admin-state enable
        encap-type dot1q
        mode access
        max-ports 64
        lacp {
            mode active
            system-id 00:00:00:00:02:03
            administrative-key 1
        }
        port 1/1/1 {
        }
    }
```

Ethernet segment "ESI-23" is configured in the service **system bgp-evpn** context on PE-2 and PE-3, as follows:

```
# on PE-2, PE-3:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "ESI-23" {
                        admin-state enable
                        esi 01:00:00:00:00:23:00:00:00:01
                        multi-homing-mode all-active
                        df-election {
                            es-activation-timer 3
                        }
                        association {
                            lag "lag-1" {
                            }
                        }
                    }
                }
            }
        }
```

When configuring an ES, the following must be considered:

- Any EVPN parameter that is not specific to any particular VPLS service, and is common to all the EVIs, is configured in a base BGP-EVPN instance located at **config>service>system>bgp-evpn**. In this base instance, the following attributes may be configured:
  - **ethernet-segments**
  - the base BGP-EVPN instance **route-distinguisher** that will be used for the ES routes. If this **route-distinguisher** is not configured, by default a type-1 RD will be derived as system-ip:0. The ES route distinguisher can be configured using the following command:

```
[ex:configure service system bgp evpn]
A:admin@PE-2# route-distinguisher ?
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
route-distinguisher <string>
<string> - <0..255 characters>

   Route distinguisher for ES routes
```

- The ES must be configured with a name and can contain the following parameters when configured for all-active multi-homing:
  - **esi** — 10-byte identifier that represents the ES in the BGP control plane. The same ESI must be configured in all the PEs connected to the same CE/MTU (using a unique value that cannot be associated with any other CE/MTU/access network). RFC 7432 defines five different types of ESI. In SR OS, the **type** byte, as well as the other 9 bytes can be arbitrarily configured.
  - **multi-homing-mode all-active** — This command indicates that the ES is in all-active mode.
  - **association>lag** <*lag-id*> — The LAG connected to the CE/MTU must be added to the ES. In this example, LAG "lag-1" is added to ESI-23, on both PE-2 and PE-3. Although a different LAG-id may have been assigned to the same ES on PE-2 and PE-3, PE-2 and PE-3 must have the same configuration on the ES LAG; that is, encap-type. Also, if LACP is added (it is not mandatory), both PEs must have the same admin-key, system-id, and system-priority. MTU-1 will see PE-2 and PE-3 as a single LAG peer. For all-active multi-homing, only the **lag** option is accepted by the system; **port** or **sdp** are not accepted.
  - **admin-state** — This command controls the administrative state of the ES.
- The preceding parameters are the minimum necessary so that the ES can be activated. In addition to those parameters, there are a few more that the user can configure if requiring values different from the default ones:
  - **activation-timer** [**0..100**] can be configured at **redundancy>bgp-evpn-ethernet-segment>activation-timer** or at **service>system>bgp>evpn>ethernet-segment>df-election>es-activation-timer** level (the most specific value is used).

    The ES activation timer operation is as follows:
    - Upon reception of an ES, AD per-ES/EVI route update/withdrawal for a local ESI, the DF-candidate list of IP addresses is updated and the DF election algorithm is run without waiting for any timer.
    - If the result of the DF election requires the PE to be promoted from non-DF to DF, the ES activation timer will start, and only after its expiration will the PE add the SAP to the default multicast list. Transitions from non-DF to non-DF, or from DF to non-DF, are immediate and do not wait for any timer.
    - This use of an ES activation timer value minimizes the risks of loops and packet duplication due to transient multiple DFs.
    - The same ES activation timer must be configured in all the PEs that are part of the same ESI. The user must configure either a long timer to minimize the risks of loops/duplication, or **(es-)activation-timer**=0 to speed up the convergence for NDF to DF transitions. The default value is 3 seconds.
  - **service-carving-mode** — As defined in RFC 7432, service carving controls the distribution of DF/non-DF roles across the different services defined in an ES.

```
[ex:configure service system bgp evpn ethernet-segment "ESI-23" df-election]
A:admin@PE-2# service-carving-mode ?

 service-carving-mode <keyword>
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
 <keyword> - (auto|manual|off)
 Default   - auto

    Mode of service carving enabled per EVPN associated with this Ethernet segment
    entry
```

```
[ex:configure service system bgp evpn ethernet-segment "ESI-23" df-election]
A:admin@PE-2# manual ?

 manual

 evi                 + Enter the evi context
 isid                + Enter the isid context
 preference          + Enter the preference context
```

As shown above, **service-carving** has three different modes:

- **service-carving-mode auto** (default) — The DF election algorithm will run the function [V(evi) mod N(peers) = i(ordinal)] to know who the DF for a specified service and ESI is. In this example, ESI-23 is configured with mode **auto**; therefore, for VPLS-1 (with EVI-1), PE-3 will be elected as DF because evi(1) mod (2)peers = 1, and the ordinal 1 corresponds to the second lowest IP, PE-3. The algorithm takes the configured **evi** in the service; therefore, the **evi** is mandatory, and for the same service must match in all the PEs that are part of the ES. This guarantees that the election algorithm is consistent across all the PEs of the ESI.
- **service-carving-mode manual** — The user can manually decide for which **evi** identifiers the PE is DF**: service-carving-mode manual / manual evi** <*start*> **end** <*to*>. The PE will be non-DF for the non-specified EVIs. If **service-carving-mode manual** is configured, but no range is defined, all the services are considered to be non-DF. If a range is configured, but the **service-carving-mode** is not **manual**, the range has no effect. Only two PEs are supported when **service-carving-mode manual** is configured.
- **service-carving-mode off** — The lowest originator IP will win the election for a specified service and ES.
- Because the **evi** is used for the service carving algorithm, it must always be configured in a service with SAPs/SDP bindings created in an ES, regardless of the service-carving mode (service-carving off, auto, or manual).

Although not configured as part of the ES, the **config>redundancy>bgp-evpn-ethernet-segment>boot-timer** allows the necessary time for the control plane protocols to come up after the PE has rebooted, and before bringing up the ESs and running the DF algorithm. Some considerations about the boot timer:

- The boot timer should use a value long enough to allow the IOMs and BGP sessions to come up before exchanging ES routes and run the DF election for each EVI (it is 10 s, by default).
- The boot timer runs per EVI on the ESs in the system. While **system-up-time < boot-timer**, the system will not run the DF election for any EVI. When the boot timer expires, the DF election for the EVI is run and, if the system is elected DF for the EVI, the ES activation timer will start.
- The system will not advertise ES routes until the boot timer expires. This guarantees that the peer ES PEs do not run the DF election either, until the PE is ready to become the DF, if needed.
- The following show command displays the configured boot timer, as well as the remaining timer if the system is still in boot stage.

```


[/]
A:admin@PE-2# show redundancy bgp-evpn-multi-homing
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Redundancy BGP EVPN Multi-homing Information
===============================================================================
Boot-Timer              : 10 secs
Boot-Timer Remaining    : 0 secs
ES Activation Timer     : 3 secs
===============================================================================
```

After ESI-23 is configured in PE-2 and PE-3, the lag-1 SAPs in both PEs can be added to the VPLS-1 service. Until the ESI-23 is successfully enabled, the LAG SAPs will be kept down with a StandByForMHProtocol flag. This is illustrated in the following example for PE-2 where the LAG SAP is added and ESI-23 is disabled:

```
# on PE-2:
configure exclusive
    service {
        vpls "VPLS1" {
            sap lag-1:1 {
            }
        }
        system {
            bgp {
                evpn {
                    ethernet-segment "ESI-23" {
                        admin-state disable
                    }
                    commit
```

```
[/]
A:admin@PE-2# show service id 1 sap lag-1:1 detail | match "  Oper State"
Admin State        : Up                       Oper State        : Down
```

```
[/]
A:admin@PE-2# show service id 1 sap lag-1:1 detail | match Flag
Flags              : StandByForMHProtocol
```

ESI-23 is enabled and SAP lag-1:1 is operationally up, as follows:

```
# on PE-2:
configure exclusive
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "ESI-23" {
                        admin-state enable
                    }
                    commit
```

```
[/]
A:admin@PE-2# show log log-id "99"
---snip---

107 2021/02/25 17:06:03.759 CET MINOR: SVCMGR #2203 Base
"Status of SAP lag-1:1 in service 1 (customer 1) changed to admin=up oper=up flags="
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

All-active multi-homing operation

To confirm that all-active multi-homing is working correctly for ESI-23, the user can use the following commands:

- **show service system bgp-evpn** — Shows the RD is used for the ES route.
- **show service system bgp-evpn ethernet-segment** — Shows all the ESs configured in the PE and their admin/operational status.
- **show service system bgp-evpn ethernet-segment name ESI-23 evi evi-1 1** — Shows the DF candidate PEs for EVI 1 and whether the system is DF for EVI.
- **show service system bgp-evpn ethernet-segment name ESI-23 all** — Shows all the information related to a specific ESI.

The base BGP-EVPN information includes the RD:

```
[/]
A:admin@PE-2# show service system bgp-evpn

===============================================================================
System BGP EVPN Information
===============================================================================
Eth Seg Route Dist.               : <none>
Eth Seg Oper Route Dist.          : 192.0.2.2:0
Eth Seg Oper Route Dist Type      : default
Ad Per ES Route Target            : evi-rt
Leaf Label                        : 0
Mcast Leave Sync Prop             : 5
Attribute Uniform Prop            : Disabled
BGP Path Selection                : Disabled
===============================================================================
```

The following command shows the configured ESs in the PE and their status:

```
[/]
A:admin@PE-2# show service system bgp-evpn ethernet-segment

===============================================================================
Service Ethernet Segment
===============================================================================
Name                            ESI                             Admin     Oper
-------------------------------------------------------------------------------
ESI-23                          01:00:00:00:00:23:00:00:00:01 Enabled   Up
-------------------------------------------------------------------------------
Entries found: 1
===============================================================================
```

The following command shows that PE-2 is not the DF and the DF candidate PEs for EVI 1 are PE-2 and PE-3:

```
[/]
A:admin@PE-2# show service system bgp-evpn ethernet-segment name "ESI-23" evi evi-1 1

===============================================================================
EVI DF and Candidate List
===============================================================================
EVI           SvcId         Actv Timer Rem      DF  DF Last Change
-------------------------------------------------------------------------------
1             1             0                   no  02/25/2021 17:06:04
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
DF Candidates                          Time Added
-------------------------------------------------------------------------------
192.0.2.2                              02/25/2021 17:06:04
192.0.2.3                              02/25/2021 17:06:06
-------------------------------------------------------------------------------
Number of entries: 2
===============================================================================
```

The following command shows all information related to ESI-23 on PE-2:

```
[/]
A:admin@PE-2# show service system bgp-evpn ethernet-segment name "ESI-23" all

===============================================================================
Service Ethernet Segment
===============================================================================
Name                    : ESI-23
Eth Seg Type            : None
Admin State             : Enabled            Oper State         : Up
ESI                     : 01:00:00:00:00:23:00:00:00:01
Multi-homing            : allActive          Oper Multi-homing  : allActive
ES SHG Label            : 524279
Source BMAC LSB         : <none>
Lag                     : lag-1
ES Activation Timer     : 3 secs
Oper Group              : (Not Specified)
Svc Carving             : auto               Oper Svc Carving   : auto
Cfg Range Type          : primary
===============================================================================

===============================================================================
EVI Information
===============================================================================
EVI                 SvcId               Actv Timer Rem      DF
-------------------------------------------------------------------------------
1                   1                   0                   no
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================

-------------------------------------------------------------------------------
DF Candidate list
-------------------------------------------------------------------------------
EVI                                    DF Address
-------------------------------------------------------------------------------
1                                      192.0.2.2
1                                      192.0.2.3
-------------------------------------------------------------------------------
Number of entries: 2
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
---snip---
```

The following command shows all information related to ESI-23 on PE-3:

```
[/]
A:admin@PE-3# show service system bgp-evpn ethernet-segment name "ESI-23" all

===============================================================================
Service Ethernet Segment
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Name                    : ESI-23
Eth Seg Type            : None
Admin State             : Enabled            Oper State         : Up
ESI                     : 01:00:00:00:00:23:00:00:00:01
Multi-homing            : allActive          Oper Multi-homing  : allActive
ES SHG Label            : 524279
Source BMAC LSB         : <none>
Lag                     : lag-1
ES Activation Timer     : 3 secs
Oper Group              : (Not Specified)
Svc Carving             : auto               Oper Svc Carving   : auto
Cfg Range Type          : primary
===============================================================================

===============================================================================
EVI Information
===============================================================================
EVI                 SvcId               Actv Timer Rem      DF
-------------------------------------------------------------------------------
1                   1                   0                   yes
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================

-------------------------------------------------------------------------------
DF Candidate list
-------------------------------------------------------------------------------
EVI                                     DF Address
-------------------------------------------------------------------------------
1                                       192.0.2.2
1                                       192.0.2.3
-------------------------------------------------------------------------------
Number of entries: 2
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
---snip---
```

The preceding commands show the ESI-23 configuration on both PEs and the result of the DF election for EVI 1.

The following output shows the ES route received on PE-2:

```
# on PE-2:
63 2021/02/25 17:04:23.069 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 71
    Flag: 0x90 Type: 14 Len: 34 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.3
        Type: EVPN-ETH-SEG Len: 23 RD: 192.0.2.3:0
              ESI: 01:00:00:00:00:23:00:00:00:01, IP-Len: 4 Orig-IP-Addr: 192.0.2.3
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        df-election::DF-Type:Auto/DP:0/DF-Preference:0/AC:1
        target:00:00:00:00:23:00
"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The ES RT as shown as target:00:00:00:00:23:00 in the extended community is auto-derived from the ESI bytes 2 to 7 (with the type byte being byte 1). Only PE-2 and PE-3 generate this RT and therefore import each other's ES route.

The following message in log 99 on PE-3 shows the result of the DF election for EVI 1.

```
# log "99" on PE-3:
99 2021/02/25 17:04:27.467 CET MINOR: SVCMGR #2094 Base
"Ethernet Segment:ESI-23, EVI:1, Designated Forwarding state changed to:true"
```

The **show service system bgp-evpn ethernet-segment name ESI-23 all** command shows the ESI-label allocated to the PE: ES SHG Label 524282 in the CLI output for PE-3. In this example, this label is allocated by PE-3 for ESI-23 (a different one is allocated per ESI) and advertised in the AD per-ES route for ESI-23. The following output shows the AD per-ES and AD per-EVI (for evi 1) routes sent by PE-3 and received by PE-2.

- The AD per-ES route can be identified by the **MAX-ET** in the ethernet-tag field (as per RFC 7432) and carries the ESI-label as well as the multi-homing mode (all-active in this case) in the ESI-label extended community (see Figure 87: EVPN route types and NLRIs).

  The user can enable the aggregation of AD per-ES routes by using the following command: **configure service system bgp evpn ad-per-es-route route-target-type evi-route-target-set route-distinguisher-ip-address** *ip-address*. If enabled, a single AD per-ES route with the associated RD and a set of EVI route-targets will be advertised (to a maximum of 128). When there are more than 128 EVIs defined in the ES, more than one route will be sent by the system.

  ```
  [ex:/configure service system bgp evpn ad-per-es-route]
  A:admin@PE-2# route-distinguisher-ip-address ?

   route-distinguisher-ip-address <ipv4-address>
   <ipv4-address>  - <d.d.d.d>

      IP address for route distinguisher for EVPN AD-ES routes
  ```

  The following AD per-ES route is received on PE-2:

  ```
  # AD per-ES route received on PE-2:
  101 2021/02/25 17:06:06.205 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
  "Peer 1: 192.0.2.3: UPDATE
  Peer 1: 192.0.2.3 - Received BGP UPDATE:
      Withdrawn Length = 0
      Total Path Attr Length = 73
      Flag: 0x90 Type: 14 Len: 36 Multiprotocol Reachable NLRI:
          Address Family EVPN
          NextHop len 4 NextHop 192.0.2.3
          Type: EVPN-AD Len: 25 RD: 192.0.2.3:1 ESI: 01:00:00:00:00:23:00:00:00:01,
                        tag: MAX-ET Label: 0
      Flag: 0x40 Type: 1 Len: 1 Origin: 0
      Flag: 0x40 Type: 2 Len: 0 AS Path:
      Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
      Flag: 0xc0 Type: 16 Len: 16 Extended Community:
          target:64500:1
          esi-label:524279/All-Active
  "
  ```

- The AD per-EVI route has an eth-tag 0 and carries the service label in the NLRI.

```
# AD per-EVI route received on PE-2:
100 2021/02/25 17:06:06.204 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 73
    Flag: 0x90 Type: 14 Len: 36 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.3
        Type: EVPN-AD Len: 25 RD: 192.0.2.3:1 ESI: 01:00:00:00:00:23:00:00:00:01,
                      tag: 0 Label: 8388512
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:1
        bgp-tunnel-encap:MPLS
"
```

```
[/]
A:admin@PE-2# show router bgp routes evpn auto-disc esi 01:00:00:00:00:23:00:00:00:01
===============================================================================
 BGP Router ID:192.0.2.2        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Auto-Disc Routes
===============================================================================
Flag  Route Dist.         ESI                           NextHop
      Tag                                               Label
-------------------------------------------------------------------------------
u*>i  192.0.2.3:1         01:00:00:00:00:23:00:00:00:01 192.0.2.3
      0                                                 LABEL 524282

u*>i  192.0.2.3:1         01:00:00:00:00:23:00:00:00:01 192.0.2.3
      MAX-ET                                            LABEL 0

-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

```
[/]
A:admin@PE-2# show router bgp routes evpn auto-disc esi 01:00:00:00:00:23:00:00:00:01 hunt
---snip---
===============================================================================
BGP EVPN Auto-Disc Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : n/a
Nexthop        : 192.0.2.3
From           : 192.0.2.3
Res. Nexthop   : 192.168.23.2
---snip---
Community      : target:64500:1 bgp-tunnel-encap:MPLS
---snip---
EVPN type      : AUTO-DISC
ESI            : 01:00:00:00:00:23:00:00:00:01
Tag            : 0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Route Dist.    : 192.0.2.3:1
MPLS Label     : LABEL 524282

---snip---

Network        : n/a
Nexthop        : 192.0.2.3
From           : 192.0.2.3
Res. Nexthop   : 192.168.23.2
---snip---
Community      : target:64500:1 esi-label:524279/All-Active
---snip---
EVPN type      : AUTO-DISC
ESI            : 01:00:00:00:00:23:00:00:00:01
Tag            : MAX-ET
Route Dist.    : 192.0.2.3:1
MPLS Label     : LABEL 0
---snip---
```

From a service perspective, as soon as CE-11 sends some traffic, the PE learning the CE-11 MAC address will advertise it to the network. The remote PEs (PE-4 and PE-5) will create a new EVPN-MPLS ES destination to ESI-23, with two next-hops: PE-2 and PE-3. The following outputs show the following information:

- PE-4 has learned AD per-EVI/ES routes for ESI-23 from PE-2 and PE-3, as well as the CE-11 MAC address from PE-3 (because MTU-1 picked up its link to PE-3 to send CE-11 frames).

```
[/]
A:admin@PE-4# show router bgp routes evpn auto-disc esi 01:00:00:00:00:23:00:00:00:01
===============================================================================
 BGP Router ID:192.0.2.4        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Auto-Disc Routes
===============================================================================
Flag  Route Dist.         ESI                           NextHop
      Tag                                               Label
-------------------------------------------------------------------------------
u*>i  192.0.2.2:1         01:00:00:00:00:23:00:00:00:01 192.0.2.2
      0                                                 LABEL 524282

u*>i  192.0.2.2:1         01:00:00:00:00:23:00:00:00:01 192.0.2.2
      MAX-ET                                            LABEL 0

u*>i  192.0.2.3:1         01:00:00:00:00:23:00:00:00:01 192.0.2.3
      0                                                 LABEL 524282

u*>i  192.0.2.3:1         01:00:00:00:00:23:00:00:00:01 192.0.2.3
      MAX-ET                                            LABEL 0

-------------------------------------------------------------------------------
Routes : 4
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

PE-4 has learned MAC address 00:00:11:11:11:11 of CE-11 in ESI-23. The BGP EVPN MAC route has PE-3 as next hop:

```
[/]
A:admin@PE-4# show router bgp routes evpn mac rd 192.0.2.3:1
===============================================================================
 BGP Router ID:192.0.2.4        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN MAC Routes
===============================================================================
Flag  Route Dist.         MacAddr           ESI
      Tag                 Mac Mobility      Label1
                          Ip Address
                          NextHop
-------------------------------------------------------------------------------
u*>i  192.0.2.3:1         00:00:11:11:11:11 01:00:00:00:00:23:00:00:00:01
      0                   Seq:0             LABEL 524282
                          n/a
                          192.0.2.3

-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

- In the FDB for VPLS-1, PE-4 has learned the CE-11 MAC address associated with a newly created EVPN-MPLS ES destination:

```
[/]
A:admin@PE-4# show service id 1 fdb mac 00:00:11:11:11:11

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
1          00:00:11:11:11:11 eES:                    Evpn     02/25/21 17:14:43
                             01:00:00:00:00:23:00:00:00:01
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

- Due to the aliasing function, the newly created EVPN-MPLS ES destination to ESI-23 has two next-hops (PE-2 and PE-3), to which PE-4 can load-balance the unicast traffic because **ecmp 2** is configured in the VPLS-1 of PE-4.

```
[/]
A:admin@PE-4# show service id 1 evpn-mpls

===============================================================================
BGP EVPN-MPLS Dest
===============================================================================
TEP Address     Egr Label     Num. MACs   Mcast           Last Change
                Transport:Tnl                             Sup BCast Domain
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
192.0.2.2       524281        0           bum             02/25/2021 16:46:48
                ldp:65538                                 No
192.0.2.3       524281        0           bum             02/25/2021 16:46:48
                ldp:65537                                 No
192.0.2.5       524281        0           bum             02/25/2021 16:46:52
                ldp:65539                                 No
-------------------------------------------------------------------------------
Number of entries : 3
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN-MPLS Ethernet Segment Dest
===============================================================================
Eth SegId                       Num. Macs               Last Change
-------------------------------------------------------------------------------
01:00:00:00:00:23:00:00:00:01   1                       02/25/2021 17:14:43
-------------------------------------------------------------------------------
Number of entries: 1
-------------------------------------------------------------------------------
===============================================================================
---snip---
```

The **show service id 1 evpn-mpls esi esi-1 01:00:00:00:00:23:00:00:00:01** command shows the next-hops that the EVPN-MPLS ES destination is resolved to.

```
[/]
A:admin@PE-4# show service id 1 evpn-mpls esi esi-1 01:00:00:00:00:23:00:00:00:01

===============================================================================
BGP EVPN-MPLS Ethernet Segment Dest
===============================================================================
Eth SegId                       Num. Macs               Last Change
-------------------------------------------------------------------------------
01:00:00:00:00:23:00:00:00:01   1                       02/25/2021 17:14:43
===============================================================================

===============================================================================
BGP EVPN-MPLS Dest TEP Info
===============================================================================
TEP Address             Egr Label               Last Change
                        Transport:Tnl-Id
-------------------------------------------------------------------------------
192.0.2.2               524282                  02/25/2021 17:14:43
                        ldp:65538
192.0.2.3               524282                  02/25/2021 17:14:43
                        ldp:65537
-------------------------------------------------------------------------------
Number of entries : 2
-------------------------------------------------------------------------------
===============================================================================
```

- PE-2 will show the CE-11 MAC address as learned locally in SAP lag-1:1 (because the data plane learning of the CE-11 MAC address happened in PE-2). For PE-3, even though it learned the MAC address from EVPN, it will install it as associated with SAP lag-1:1 because the EVPN route came with ESI-23, which is a local ESI. Because of this, whenever PE-3 receives a frame with MAC DA equal to the CE-11 MAC address, it will be able to forward the frame locally to the SAP lag-1:1. The following output shows the CE-11 MAC address as it is installed in PE-2 and PE-3:

```
[/]
A:admin@PE-2# show service id 1 fdb mac 00:00:11:11:11:11
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
1          00:00:11:11:11:11 sap:lag-1:1             L/90     02/25/21 17:14:43
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

```
[/]
A:admin@PE-3# show service id 1 fdb mac 00:00:11:11:11:11

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
1          00:00:11:11:11:11 sap:lag-1:1             Evpn     02/25/21 17:14:43
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

Single-active multi-homing concepts

Figure 90: EVPN-MPLS single-active multi-homing: mass-withdraw, backup path illustrates two concepts in EVPN single-active multi-homing: mass-withdraw and backup path.

Figure 90: EVPN-MPLS single-active multi-homing: mass-withdraw, backup path

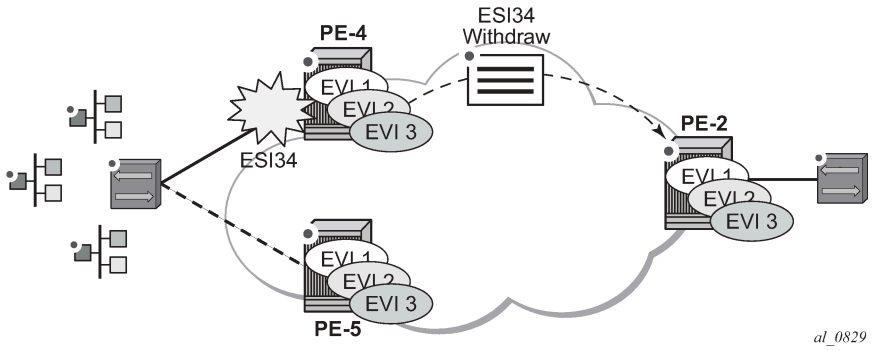

- With mass-withdraw, when ESI-45 goes down, PE-2 does not have to wait for all the MAC routes to be withdrawn to converge all the services. Instead, PE-4 will withdraw the AD per-ES routes (also the AD per-EVI and MAC routes) and that will be used at PE-2 as a notification to stop sending traffic to PE-4 for any MAC address associated with ESI-45.
- With backup path, when PE-2 is notified of the ESI-45 failure due to the withdrawn AD routes, it will not flush any MAC address associated with ESI-45. Instead, it will change the next-hop of the EVPN-MPLS ES destination to the remaining PE in the ESI-45. Backup path only works when there are two PEs in the same ES. If there were more than two PEs in ESI-45, PE-2 would flush all the MAC addresses upon receiving a mass-withdraw notification, because it would not know who the new active PE is.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Single-active multi-homing configuration

The single-active multi-homing configuration example is based on Figure 88: EVPN-MPLS for VPLS services:

MTU-6 is connected to the EVPN network using single-active multi-homing. With the MTU-6 configuration, a VPLS service with active-standby spoke-SDP to PE-4 and PE-5 is configured. In PE-4 and PE-5, the SDP connected to MTU-6 is linked to ESI-45. Both will run the DF election algorithm for EVI 1, and the non-DF PE (PE-4 in this example) will bring down the spoke-SDP and notify MTU-6.

The following output shows the configuration of ESI-45 in PE-4 and PE-5, as well as the SDPs. The configuration of MTU-6 is also shown for completeness. It is important to keep the default — **ignore-standby-signaling false** — on MTU-6 spoke-SDPs because the PW switchover in MTU-6 will be triggered based on the PW status bits sent by PE-4 and PE-5.

SDP 46 with far-end MTU-6 is configured on PE-4:

```
# on PE-4:
configure {
    service {
        sdp 46 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.6
            }
        }
```

Ethernet segment "ESI-45" is configured on PE-4 as follows:

```
# on PE-4:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "ESI-45" {
                        admin-state enable
                        esi 01:00:00:00:00:45:00:00:00:01
                        multi-homing-mode single-active
                        df-election {
                            es-activation-timer 3
                        }
                        association {
                            sdp 46 {
                            }
                        }
                    }
                }
            }
        }
```

On PE-5, SDP 56 is configured as follows:

```
# on PE-5:
configure {
    service {
        sdp 56 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.6
            }
        }
```

Ethernet segment "ESI-45" is configured as follows on PE-5:

```
# on PE-5:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "ESI-45" {
                        admin-state enable
                        esi 01:00:00:00:00:45:00:00:00:01
                        multi-homing-mode single-active
                        df-election {
                            es-activation-timer 3
                        }
                        association {
                            sdp 56 {
                            }
                        }
                    }
                }
            }
        }
```

On MTU-6, the service configuration is as follows:

```
# on MTU-6:
configure {
    service {
        sdp 64 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.4
            }
        }
        sdp 65 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.5
            }
        }
        vpls "VPLS1" {
            admin-state enable
            service-id 1
            customer "1"
            endpoint "CORE" {
            }
            spoke-sdp 64:1 {
                endpoint {
                    name "CORE"
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                stp {
                    admin-state disable
                }
            }
            spoke-sdp 65:1 {
                endpoint {
                    name "CORE"
                }
                stp {
                    admin-state disable
                }
            }
            sap 1/2/1:1 {
            }
        }
```

For a detailed description of the base BGP-EVPN instance and ES configuration, see the All-active multi-homing configuration section. The **es-activation-timer**, **esi**, **service-carving-mode**, **boot-timer**, and **admin-state** commands are used in the same way as for all-active multi-homing. Only the differences compared to all-active multi-homing are described here:

- **multi-homing-mode single-active** must be configured so that the ES acts as single-active. Optionally, **multi-homing-mode single-active-no-esi-label** can be configured, which controls the use of the ESI-label for single-active multi-homing. Although the ESI-label is always used in all-active multi-homing when sending BUM traffic between the PEs in the ES, it is configurable for single-active. However, Nokia recommends to use the default option (using ESI-label) to avoid potential transient issues when there is a DF switchover.
- **association>sdp** <*sdp-id*> is configured so that the ES can be associated with the SDP connected to MTU-6. Although all-active multi-homing only allows LAG associations to the ES, single-active allows LAG, port, and SDP. In this example, SDP is the option, because the access network is MPLS-based.

Similar to the all-active multi-homing case, when configuring the service in PE-4 and PE-5, the service objects are automatically associated with the ESI-45, because they are defined in the SDPs linked to the ESI. The configuration for VPLS 1 on PE-5 is as follows:

```
# on PE-5:
configure {
    service {
        vpls "VPLS1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
            }
            bgp-evpn {
                evi 1
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    ecmp 2
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            spoke-sdp 56:1 {
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

In all-active multi-homing, the non-DF does not bring down the service SAP associated with the ES (it only removes it from the default multicast list). However, in single-active multi-homing, the service spoke-SDP (or SAP, if that was the object associated) is brought operationally down. The following output shows the spoke-SDP state in PE-4 (non-DF), as operationally down with the **StandbyForMHProtocol** flag and the **Local Pw Bits** that are signaled to MTU-6:

```
[/]
A:admin@PE-4# show service system bgp-evpn ethernet-segment name "ESI-45" evi evi-1 1

===============================================================================
EVI DF and Candidate List
===============================================================================
EVI           SvcId         Actv Timer Rem      DF  DF Last Change
-------------------------------------------------------------------------------
1             1             0                   no  02/25/2021 17:19:02
===============================================================================

===============================================================================
DF Candidates                          Time Added
-------------------------------------------------------------------------------
192.0.2.4                              02/25/2021 17:18:56
192.0.2.5                              02/25/2021 17:19:02
-------------------------------------------------------------------------------
Number of entries: 2
===============================================================================
```

Spoke-SDP 46:1 is operationally down on PE-4:

```
[/]
A:admin@PE-4# show service id 1 sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
46:1             Spok     192.0.2.6       Up      Down      524280    524281
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
```

Spoke-SDP 46:1 is operationally down with the StandbyForMHProtocol flag:

```
[/]
A:admin@PE-4# show service id 1 sdp 46:1 detail | match Flag
Flags              : StandbyForMHProtocol
```

The local PW bits (**pwFwdingStandby**) are sent to MTU-6:

```
[/]
A:admin@PE-4# show service id 1 sdp 46:1 detail | match Pw
Local Pw Bits      : pwFwdingStandby
Peer Pw Bits       : None
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Single-active multi-homing operation

The same commands used in the All-active multi-homing operation section can be used for single-active; see that section.

The **show service system bgp-evpn ethernet-segment name ESI-45** command shows an Ethernet-segment **Oper Multi-homing** in addition to the configured **Multi-homing** mode. This occurs because, in spite of configuring the ES as all-active, it may operate as single-active if there is a mismatch between the modes advertised by PE-4 and PE-5 in the AD per-ES routes (per RFC 7432). In this example, the configured and the operational value are the same:

```
[/]
A:admin@PE-4# show service system bgp-evpn ethernet-segment name "ESI-45"

===============================================================================
Service Ethernet Segment
===============================================================================
Name                    : ESI-45
Eth Seg Type            : None
Admin State             : Enabled           Oper State         : Up
ESI                     : 01:00:00:00:00:45:00:00:00:01
Multi-homing            : singleActive      Oper Multi-homing  : singleActive
ES SHG Label            : 524278
Source BMAC LSB         : <none>
Sdp Id                  : 46
ES Activation Timer     : 3 secs
Oper Group              : (Not Specified)
Svc Carving             : auto              Oper Svc Carving   : auto
Cfg Range Type          : primary
===============================================================================
```

As soon as CE-16 sends some traffic, the DF PE (PE-5) will learn the CE-16 MAC address and will advertise it to the network. The remote PEs (PE-2 and PE-3) will create a new EVPN-MPLS ES destination to ESI-45, but this time with only one next-hop, PE-5, because this is single-active multi-homing. The following outputs show the following information:

- PE-2 has learned AD per-EVI/ES routes for ESI-45 from PE-4 and PE-5, as well as the CE-16 MAC address from an ES EVPN-MPLS destination, which is resolved to PE-5 (the DF for ESI-45).

```
[/]
A:admin@PE-2# show router bgp routes evpn auto-disc esi 01:00:00:00:00:45:00:00:00:01
===============================================================================
 BGP Router ID:192.0.2.2        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Auto-Disc Routes
===============================================================================
Flag  Route Dist.         ESI                           NextHop
      Tag                                               Label
-------------------------------------------------------------------------------
u*>i  192.0.2.4:1         01:00:00:00:00:45:00:00:00:01 192.0.2.4
      0                                                 LABEL 524282

u*>i  192.0.2.4:1         01:00:00:00:00:45:00:00:00:01 192.0.2.4
      MAX-ET                                            LABEL 0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
u*>i  192.0.2.5:1         01:00:00:00:00:45:00:00:00:01 192.0.2.5
      0                                                 LABEL 524282

u*>i  192.0.2.5:1         01:00:00:00:00:45:00:00:00:01 192.0.2.5
      MAX-ET                                            LABEL 0

-------------------------------------------------------------------------------
Routes : 4
===============================================================================
```

PE-2 has learned the CE-16 MAC address from an ES EVPN-MPLS destination:

```
[/]
A:admin@PE-2# show service id 1 fdb mac 00:00:16:16:16:16

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
1          00:00:16:16:16:16 eES:                    Evpn     02/25/21 17:20:53
                             01:00:00:00:00:45:00:00:00:01
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

On PE-2, the ES EVPN-MPLS destination is resolved to DF PE-5:

```
[/]
A:admin@PE-2# show service id 1 evpn-mpls esi esi-1 01:00:00:00:00:45:00:00:00:01

===============================================================================
BGP EVPN-MPLS Ethernet Segment Dest
===============================================================================
Eth SegId                       Num. Macs               Last Change
-------------------------------------------------------------------------------
01:00:00:00:00:45:00:00:00:01   1                       02/25/2021 17:20:53
===============================================================================

===============================================================================
BGP EVPN-MPLS Dest TEP Info
===============================================================================
TEP Address             Egr Label               Last Change
                        Transport:Tnl-Id
-------------------------------------------------------------------------------
192.0.2.5               524282                  02/25/2021 17:20:53
                        ldp:65539
-------------------------------------------------------------------------------
Number of entries : 1
-------------------------------------------------------------------------------
===============================================================================
```

- In this case, the local PEs, PE-4 and PE-5, will learn the CE MAC address from an EVPN-MPLS destination and a local spoke-SDP, respectively.

```
[/]
A:admin@PE-4# show service id 1 fdb mac 00:00:16:16:16:16
===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            Transport:Tnl-Id                        Age
-------------------------------------------------------------------------------
1          00:00:16:16:16:16 eES:                    Evpn     02/25/21 17:20:53
                             01:00:00:00:00:45:00:00:00:01
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

The ES EVPN-MPLS destination is resolved to DF PE-5:

```
[/]
A:admin@PE-4# show service id 1 evpn-mpls esi esi-1 01:00:00:00:00:45:00:00:00:01

===============================================================================
BGP EVPN-MPLS Ethernet Segment Dest
===============================================================================
Eth SegId                       Num. Macs               Last Change
-------------------------------------------------------------------------------
01:00:00:00:00:45:00:00:00:01   1                       02/25/2021 17:20:53
===============================================================================

===============================================================================
BGP EVPN-MPLS Dest TEP Info
===============================================================================
TEP Address              Egr Label                Last Change
                         Transport:Tnl-Id
-------------------------------------------------------------------------------
192.0.2.5                524282                   02/25/2021 17:20:53
                         ldp:65539
-------------------------------------------------------------------------------
Number of entries : 1
-------------------------------------------------------------------------------
===============================================================================
```

DF PE-5 learns the CE-16 MAC address from a local spoke SDP:

```
[/]
A:admin@PE-5# show service id 1 fdb mac 00:00:16:16:16:16

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
1          00:00:16:16:16:16 sdp:56:1                L/180    02/25/21 17:20:53
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

Ethernet-segment failures

If either ES fails, a DF re-election will happen and the corresponding AD per-ES/EVI routes will be withdrawn, causing the remote PEs to modify the list of next-hops for the EVPN-MPLS ES destination. The following example illustrates a failure on the SDP between MTU-6 and PE-5 (the DF).

1. A failure occurs in the LSP between MTU-6 and PE-5. This can be any event that brings the SDP down.

```
# log "99" on PE-5:
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
94 2021/02/25 17:25:00.632 CET MINOR: SVCMGR #2303 Base
"Status of SDP 56 changed to admin=up oper=down"
```

2. Immediately, PE-5 gives up the DF role and withdraws the ES route, as well as the AD routes and MAC routes. As soon as PE-4 receives any ES or AD withdraw, it will re-run the DF algorithm and, when the es-activation-timer expires, it will become the DF and activate its spoke-SDP.

```
# log 99 on PE-5:
96 2021/02/25 17:25:00.633 CET MINOR: SVCMGR #2094 Base
"Ethernet Segment:ESI-45, EVI:1, Designated Forwarding state changed to:false"
```

The ES in PE-5 is operational down:

```
[/]
A:admin@PE-5# show service system bgp-evpn ethernet-segment name "ESI-45"

===============================================================================
Service Ethernet Segment
===============================================================================
Name                    : ESI-45
Eth Seg Type            : None
Admin State             : Enabled           Oper State         : Down
ESI                     : 01:00:00:00:00:45:00:00:00:01
Multi-homing            : singleActive      Oper Multi-homing  : singleActive
ES SHG Label            : 524279
Source BMAC LSB         : <none>
Sdp Id                  : 56
ES Activation Timer     : 3 secs
Oper Group              : (Not Specified)
Svc Carving             : auto              Oper Svc Carving   : auto
Cfg Range Type          : primary
===============================================================================
```

PE-5 is no longer the DF and the only DF candidate is PE-4:

```
[/]
A:admin@PE-5# show service system bgp-evpn ethernet-segment name "ESI-45" evi evi-1 1

===============================================================================
EVI DF and Candidate List
===============================================================================
EVI           SvcId         Actv Timer Rem      DF  DF Last Change
-------------------------------------------------------------------------------
1             1             0                   no  02/25/2021 17:25:01
===============================================================================

===============================================================================
DF Candidates                           Time Added
-------------------------------------------------------------------------------
192.0.2.4                               02/25/2021 17:19:03
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
```

PE-4 becomes the DF and the spoke-SDP 46:1 is brought up.

```
# log "99" on PE-4:
102 2021/02/25 17:25:03.598 CET MINOR: SVCMGR #2094 Base
"Ethernet Segment:ESI-45, EVI:1, Designated Forwarding state changed to:true"

103 2021/02/25 17:25:03.598 CET MINOR: SVCMGR #2326 Base
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
"Status of SDP Bind 46:1 in service 1 (customer 1) local PW status bits changed to none"

104 2021/02/25 17:25:03.598 CET MINOR: SVCMGR #2306 Base
"Status of SDP Bind 46:1 in service 1 (customer 1) changed to admin=up oper=up flags="
```

The ES is up in PE-4:

```
[/]
A:admin@PE-4# show service system bgp-evpn ethernet-segment name "ESI-45"

===============================================================================
Service Ethernet Segment
===============================================================================
Name                    : ESI-45
Eth Seg Type            : None
Admin State             : Enabled            Oper State         : Up
ESI                     : 01:00:00:00:00:45:00:00:00:01
Multi-homing            : singleActive       Oper Multi-homing  : singleActive
ES SHG Label            : 524278
Source BMAC LSB         : <none>
Sdp Id                  : 46
ES Activation Timer     : 3 secs
Oper Group              : (Not Specified)
Svc Carving             : auto               Oper Svc Carving   : auto
Cfg Range Type          : primary
===============================================================================
```

PE-4 is the DF and there are no other DF candidates:

```
[/]
A:admin@PE-4# show service system bgp-evpn ethernet-segment name "ESI-45" evi evi-1 1

===============================================================================
EVI DF and Candidate List
===============================================================================
EVI           SvcId         Actv Timer Rem      DF  DF Last Change
-------------------------------------------------------------------------------
1             1             0                   yes 02/25/2021 17:25:04
===============================================================================

===============================================================================
DF Candidates                          Time Added
-------------------------------------------------------------------------------
192.0.2.4                              02/25/2021 17:18:56
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
```

3. The remote PEs, PE-2 and PE-3, receive the BGP-EVPN routes withdrawal and modify the next-hop for the EVPN-MPLS ES destination.

```
# on PE-2:
186 2021/02/25 17:25:00.634 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.5
"Peer 1: 192.0.2.5: UPDATE
Peer 1: 192.0.2.5 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 86
    Flag: 0x90 Type: 15 Len: 82 Multiprotocol Unreachable NLRI:
        Address Family EVPN
        Type: EVPN-AD Len: 25 RD: 192.0.2.5:1 ESI: 01:00:00:00:00:45:00:00:00:01,
              tag: MAX-ET Label: 0
        Type: EVPN-AD Len: 25 RD: 192.0.2.5:1 ESI: 01:00:00:00:00:45:00:00:00:01,
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
              tag: 0 Label: 0
        Type: EVPN-ETH-SEG Len: 23 RD: 192.0.2.5:0
              ESI: 01:00:00:00:00:45:00:00:00:01, IP-Len: 4 Orig-IP-Addr: 192.0.2.5
"
```

The ES EVPN-MPLS destination is resolved to the DF PE-4:

```
[/]
A:admin@PE-2# show service id 1 evpn-mpls esi esi-1 01:00:00:00:00:45:00:00:00:01

===============================================================================
BGP EVPN-MPLS Ethernet Segment Dest
===============================================================================
Eth SegId                       Num. Macs               Last Change
-------------------------------------------------------------------------------
01:00:00:00:00:45:00:00:00:01   1                       02/25/2021 17:25:19
===============================================================================

===============================================================================
BGP EVPN-MPLS Dest TEP Info
===============================================================================
TEP Address             Egr Label               Last Change
                        Transport:Tnl-Id
-------------------------------------------------------------------------------
192.0.2.4               524282                  02/25/2021 17:25:19
                        ldp:65538
-------------------------------------------------------------------------------
Number of entries : 1
-------------------------------------------------------------------------------
===============================================================================
```

The following must be considered:

- The DF election procedure is revertive, that is, when the failed SDP comes back up, PE-5 will take over again as DF and the network will re-converge.
- The DF election is triggered by the following events:
  - Enabling an ES (**admin-state enable**) triggers the DF election for all the services in the ES.
  - A new update/withdrawal of an ES route (containing an ESI configured locally) triggers the DF election for all the services in the ESI.
  - A new update/withdrawal of an AD per-ES route (containing an ESI configured locally) triggers the DF election for all the services associated with the list of RTs received along with the route.
  - A new update of an AD per-ES route with a change in the ESI-label extended community (single-active bit or MPLS label) triggers the DF election for all the services associated with the list of RTs received along with the route.
  - A new update/withdrawal of an AD route per-EVI (containing an ESI configured locally) triggers the DF election for that service.

BGP-EVPN route selection in EVPN networks

The selection of the best route for a MAC address is as follows:

- If a PE receives more than one route for the same MAC address, the best MAC route is chosen:
  - If the route key is equal in two or more routes (that is, the mac, mac-length, ip, ip-length, RD, eth-tag), then regular BGP selection applies:

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

    - If local-pref, AS-path, origin, and MED are equal, the lowest IGP distance to the BGP next-hop is chosen (unless **ignore-nh-metric** is configured). If the BGP next-hop is resolved by an LSP, the cost from the tunnel-table is used.
    - As a last resort tie-breaker, the route with the lowest originator ID, or received from the peer with the lowest BGP Identifier, is chosen (unless **ignore-router-id** is configured and the routes being compared are EBGP routes).
  - If the mac-length, mac, ip-length, ip, eth-tag are equal, and the RD is different, the EVPN selection process is applied in the following order:
    - Conditional static MAC addresses (local protected MAC addresses)
    - EVPN static MAC addresses (remote protected MAC addresses)
    - Data plane learned MAC addresses (regular learning on SAPs/SDP-bindings)
    - EVPN MAC addresses with higher sequence number
    - Lowest IP address (next-hop IP of the EVPN NLRI)
    - Lowest Ethernet tag (will be normally zero)
    - Lowest RD
- After a MAC route is selected, the system checks for an associated ES.
  - If it has an ES, the system uses the MAC address as the EVPN-MPLS ES destination. The ES destination is constructed based on the AD per-EVI routes received for that ES (regardless of MAC address priorities with the ES).
  - The system selects the first ECMP number of AD per-EVI routes arranged by the IP address of PEs (lower IPs are selected first).
  - If the same PE has advertised multiple RDs, the system selects the route with the lowest RD for that PE.

In the example of Figure 88: EVPN-MPLS for VPLS services, PE-4 resolves the next-hops for ESI-23 as described in the second choice above, that is, because ECMP=2, the two available next-hops are chosen. If ECMP is changed to 1, PE-4 will pick up the lower IP (in the BGP next-hop). This is illustrated in the following output:

```
[/]
A:admin@PE-4# show service id 1 evpn-mpls esi esi-1 01:00:00:00:00:23:00:00:00:01

===============================================================================
BGP EVPN-MPLS Ethernet Segment Dest
===============================================================================
Eth SegId                       Num. Macs               Last Change
-------------------------------------------------------------------------------
01:00:00:00:00:23:00:00:00:01   1                       02/25/2021 17:14:43
===============================================================================

===============================================================================
BGP EVPN-MPLS Dest TEP Info
===============================================================================
TEP Address             Egr Label               Last Change
                        Transport:Tnl-Id
-------------------------------------------------------------------------------
192.0.2.2               524282                  02/25/2021 17:14:43
                        ldp:65538
192.0.2.3               524282                  02/25/2021 17:14:43
                        ldp:65537
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
Number of entries : 2
-------------------------------------------------------------------------------
===============================================================================
```

When ECMP equals 1, only the BGP next hop with the lower IP is chosen:

```
# on PE-4:
configure {
    service {
        vpls "VPLS1" {
            bgp-evpn {
                mpls 1 {
                    ecmp 1
```

```
[/]
A:admin@PE-4# show service id 1 evpn-mpls esi esi-1 01:00:00:00:00:23:00:00:00:01

===============================================================================
BGP EVPN-MPLS Ethernet Segment Dest
===============================================================================
Eth SegId                       Num. Macs               Last Change
-------------------------------------------------------------------------------
01:00:00:00:00:23:00:00:00:01   1                       02/25/2021 17:37:29
===============================================================================

===============================================================================
BGP EVPN-MPLS Dest TEP Info
===============================================================================
TEP Address             Egr Label               Last Change
                        Transport:Tnl-Id
-------------------------------------------------------------------------------
192.0.2.2               524282                  02/25/2021 17:14:43
                        ldp:65538
-------------------------------------------------------------------------------
Number of entries : 1
-------------------------------------------------------------------------------
===============================================================================
```

Comparing EVPN multi-homing and BGP multi-homing

EVPN-MPLS services support EVPN-MH (EVPN multi-homing) and also BGP-MH as in chapter BGP Multi-Homing for VPLS Networks. While EVPN-MH is the standard way of providing access resiliency in RFC 7432, BGP-MH is also a standard mechanism supported in VPLS or EVPN networks. The following table provides some comparison between both technologies.

Table 5: Comparing EVPN multi-homing and BGP multi-homing

VPN Requirements	EVPN-MH	BGP-MH	Comments
All-active MH (flow-based load-balancing)	Yes	No	EVPN-MH provides better bandwidth utilization
Single-active MH (service-based load-balancing)	Yes	Yes	

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

VPN Requirements	EVPN-MH	BGP-MH	Comments
DF PE election - automatic service balancing	Yes Service-carving	No Requires vsi policies and LP manipulation	EVPN-MH provides better automation
DF PE election – manual configuration per service	Yes	No	EVPN-MH allows for manual DF config for EVIs and ISIDs (2 PEs)
Split-horizon indication in the data plane	Yes ESI-label	No	Prevents transient loops when dual-active DFs show up
DF indication in the control plane	No	Yes	BGP MH guarantees one DF at a time. EVPN relies on timers to ensure one DF at a time
Allows multiple SAPs or SDP-bindings per service on the same site	No	Yes Through the use of SHGs	
Boot timer and site(es)-activation-timers	Yes	Yes	BGP-MH supports more granular configuration (service level)
Support for oper-groups	No	Yes	
Non-DF notification to the CE (MPLS and CFM)	Yes	Yes	Avoids blackholing

In addition to the preceding comparison, the following configuration excerpt compares EVPN-MH with BGP-MH on a bgp-evpn VPLS service and shows that, while EVPN-MH does not have any configuration at service level, BGP-MH is configured within the VPLS context, which gives a more granular control over the redundancy provided. See the BGP Multi-Homing for VPLS Networks chapter for more information about BGP-MH.

```
[ex:/configure service system]
A:admin@PE-4# info
    bgp {
        evpn {
            ethernet-segment "ESI-45" {
                admin-state enable
                esi 0x01000000004500000001
                multi-homing-mode single-active
                df-election {
                    es-activation-timer 3
                }
                association {
                    sdp 46 {
                    }
                }
            }
        }
    }
```

```
[ex:/configure service vpls "VPLS1"]
A:admin@PE-4# info
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    admin-state enable
    service-id 1
    customer "1"
    bgp 1 {
    }
    bgp-evpn {
        evi 1
        mpls 1 {
            admin-state enable
            ingress-replication-bum-label true
            ecmp 2
            auto-bind-tunnel {
                resolution any
            }
        }
    }
    spoke-sdp 46:1 {
    }
```

For BGP multi-homing, site "site-1" is configured, as follows. The RD needs to be configured in the **bgp** context.

```
[ex:/configure service vpls "VPLS1"]
A:admin@PE-4# info
    admin-state enable
    service-id 1
    customer "1"
    bgp 1 {
        route-distinguisher "192.0.2.4:1"
    }
    bgp-evpn {
        evi 1
        mpls 1 {
            admin-state enable
            ingress-replication-bum-label true
            ecmp 2
            auto-bind-tunnel {
                resolution any
            }
        }
    }
    spoke-sdp 46:1 {
    }
    bgp-mh-site "site-1" {
        admin-state enable
        id 1
        activation-timer 3
        spoke-sdp 46:1
    }
```

Proxy-ARP/ND configuration for EVPN-MPLS networks

Although not strictly a BGP-EVPN configuration, **vpls>proxy-arp** and **vpls>proxy-nd** functions are typically enabled along with EVPN-MPLS in order to reduce the amount of flooding in the network. The proxy-ARP/ND agent in the VPLS service will snoop ARP-requests and/or Neighbor Solicitation messages and will reply to those messages locally (if the information is known) without having to flood the requests to the network.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The configuration options for proxy-ARP are the following:

```
[ex:/configure service vpls "VPLS1"]
A:admin@PE-2# proxy-arp ?

 proxy-arp

 admin-state           - Administrative state of the proxy
 age-time              - Aging timer for proxy entries, where entries are flushed
                         upon timer expiry
 apply-groups          - Apply a configuration group at this level
 apply-groups-exclude  - Exclude a configuration group at this level
 duplicate-detect      + Enter the duplicate-detect context
 dynamic-arp           + Enter the dynamic-arp context
 dynamic-populate      - Populate proxy ARP entries from snooped GARP/ARP/ND
                         messages on SAPs/SDP-bindings
 evpn                  + Enter the evpn context
 send-refresh          - Time at which to send a refresh message
 static-arp            + Enter the static-arp context
 table-size            - Maximum number of learned and static entries allowed in
                         the proxy table of this service
```

The configuration options for proxy-ND are the following:

```
[ex:/configure service vpls "VPLS1"]
A:admin@PE-2# proxy-nd ?

 proxy-nd

 admin-state           - Administrative state of the proxy
 age-time              - Aging timer for proxy entries, where entries are flushed
                         upon timer expiry
 apply-groups          - Apply a configuration group at this level
 apply-groups-exclude  - Exclude a configuration group at this level
 duplicate-detect      + Enter the duplicate-detect context
 dynamic-neighbor      + Enter the dynamic-neighbor context
 dynamic-populate      - Populate proxy ARP entries from snooped GARP/ARP/ND
                         messages on SAPs/SDP-bindings
 evpn                  + Enter the evpn context
 send-refresh          - Time at which to send a refresh message
 static-neighbor       + Enter the static-neighbor context
 table-size            - Maximum number of learned and static entries allowed in
                         the proxy table of this service
```

When proxy-ARP/ND is enabled, the following configuration guidelines must be followed:

- **dynamic-populate** should be used only in networks with a consistent configuration of this command in all PEs.
- When using **dynamic-populate**, the **age-time** value should be configured to a value equal to three times the **send-refresh** value. This will help reduce the EVPN withdrawals and re-advertisements in the network.
- With large **age-time** values, it would be sufficient to configure the **send-refresh** value to half of the **proxy-ARP**/**ND age-time** or **FDB age-time**.
- In scaled environments (with thousands of services), it is not recommended to set the send-refresh value to less than 300 s. In such scenarios, Nokia recommends using a minimum proxy-ARP/ND **age-time** and FDB **age** of 900 s.
- The use of the following commands reduces or suppresses the ARP/ND flooding in an EVPN network, because EVPN MAC routes replace the function of the regular data plane ARP/ND messages:

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

  - **proxy-arp>evpn>flood> gratuitous-arp false**
  - **proxy-arp>evpn>flood> unknown-arp-req false**
  - **proxy-nd>evpn>flood> unknown-neighbor-solicitation false**
  - **proxy-nd>evpn>flood> unknown-neighbor-advertise-router false**
  - **proxy-nd>evpn>flood> unknown-neighbor-advertise-host false**
- Nokia recommends using the preceding commands only in EVPN networks where the CEs are routers directly connected to an SR OS node acting as the PE. Networks using aggregation switches between the host/routers and the PEs should flood GARP/ND messages in EVPN to make sure the remote caches are updated and BGP does not miss the advertisement of these entries.
- When **duplicate-detect anti-spoof-mac** is used with proxy-ARP/ND, ingress filters (in the access SAPs/SDP-bindings) should be configured to drop all traffic with destination anti-spoof-mac. The same MAC address should be configured in all PEs where duplicate-detect is active.
- When proxy-ND is used, the configuration of the following commands should be consistent in all the PEs in the network:
  - **proxy-nd>evpn>flood> unknown-neighbor-advertise-router**
  - **proxy-nd>evpn>flood> unknown-neighbor-advertise-host**
  - **proxy-nd>evpn>advertise-neighbor-type**
- Because EVPN does not propagate the **router** flag in IPv6--> MAC address advertisements, in a mixed network with hosts and routers where **evpn>advertise-neighbor-type router** is configured, unsolicited host NA messages should be flooded so that the entire network gets to learn all of the host and router ND entries. In the same way, **evpn>advertise-neighbor-type host** should be configured so that unsolicited router NA messages are flooded.

Finally, along with proxy-ARP/ND, **vpls>fdb>discard-unknown true** may be used in some EVPN-MPLS deployments where all the CEs are routers and they announce themselves to the network by sending GARPs or NAs (Neighbor Solicitation messages). According to RFC 7432, whether or not to flood packets to unknown destination MAC addresses should be an administrative choice, depending on how learning happens between CEs and PEs. **Discard-unknown** provides that administrative choice in case all the MAC addresses in an EVI can be learned even before any traffic is exchanged.

Proxy-ARP/ND along with **discard-unknown** helps reduce the BUM traffic in an EVPN network significantly; however, their use must be analyzed and considered, depending on the type of CEs in the EVI.

An example of proxy-ARP configuration is as follows. This configuration should be added to all PEs. When a new ARP message is received on any of the PEs, they will learn the IP-MAC address pair and will advertise it to the network.

```
# on PE-2, PE-3, PE-4, PE-5:
configure {
    service {
        vpls "VPLS1" {
            proxy-arp {
                admin-state enable
                dynamic-populate true
                age-time 900
                send-refresh 300
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Enabling proxy-ARP increases the number of MAC/IP routes being sent by the PEs. This is due to the following reasons:

- An additional MAC/IP route will be advertised per new learned IP-MAC address pair, regardless of having advertised the same MAC address already.
- A MAC per VPLS service will be advertised with a system MAC address. That MAC address will be used as MAC SA for proxy-ARP confirm messages when an IP moves to a different PE.

The following output shows the MAC/IP routes on PE-2 when proxy-ARP is enabled in the network.

```
[/]
A:admin@PE-2# show router bgp routes evpn mac
===============================================================================
 BGP Router ID:192.0.2.2        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN MAC Routes
===============================================================================
Flag  Route Dist.         MacAddr           ESI
      Tag                 Mac Mobility      Label1
                          Ip Address
                          NextHop
-------------------------------------------------------------------------------
u*>i  192.0.2.3:1         02:17:ff:00:03:3a ESI-0
      0                   Static            LABEL 524282
                          n/a
                          192.0.2.3

u*>i  192.0.2.4:1         02:1b:ff:00:03:3a ESI-0
      0                   Static            LABEL 524282
                          n/a
                          192.0.2.4

u*>i  192.0.2.5:1         00:00:16:16:16:16 01:00:00:00:00:45:00:00:00:01
      0                   Seq:0             LABEL 524282
                          n/a
                          192.0.2.5

u*>i  192.0.2.5:1         02:1f:ff:00:03:3a ESI-0
      0                   Static            LABEL 524282
                          n/a
                          192.0.2.5

-------------------------------------------------------------------------------
Routes : 4
===============================================================================
```

Troubleshooting and debug commands

When troubleshooting an EVPN-MPLS network, the following show commands and debug commands are recommended, as already discussed throughout this chapter:

- **show redundancy bgp-evpn-multi-homing**
- **show router bgp routes evpn (and filters)**

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- **show service evpn-mpls [<TEP ip-address>]**
- **show service id bgp-evpn**
- **show service id evpn-mpls (and modifiers)**
- **show service id fdb (and modifiers)**
- **show service system bgp-evpn**
- **show service system bgp-evpn ethernet-segment (and modifiers)**
- **debug router bgp update** (in classic CLI)
- **log-id "99"**

In addition to the preceding commands, the following tools dump commands may also help:

- **tools dump service evpn usage** — This command shows the amount of EVPN-MPLS (and EVPN-VXLAN) destinations consumed in the system.
- **tools dump service system bgp-evpn ethernet-segment** <*name*> **evi <[1..65535]> df** — This command computes the DF election for a specific ESI and EVI. Note: The **show service system bgp-evpn ethernet-segment** commands shows whether the local PE is DF or non-DF for a specific EVI, but it does not show who the DF is if it is not the local PE. In case of more than 2 PEs in the ES, this command may be especially useful.

Some examples are provided below for PE-2. PE-2 is showing seven EVPN-MPLS destinations due to the following:

- Each remote PE consumes one EVPN-MPLS destination for unicast (if they advertise MAC/IP routes to PE-2 and the ingress-replication-bum-label is configured in all the PEs). PE-2 has three remote unicast EVPN-MPLS destinations.
- Each remote PE consumes one EVPN-MPLS destination for multicast (if they advertise inclusive multicast routes to PE-2). PE-2 has three remote multicast EVPN-MPLS destinations.
- Each remote ES consumes one EVPN-MPLS destination (it is only one per ES, regardless of the multi-homing mode and the number of PEs in the ES). PE-2 has one remote ES (ESI-45).

```
[/]
A:admin@PE-2# tools dump service evpn usage

vxlan-evpn-mpls usage statistics at 02/25/2021 17:40:37:

MPLS-TEP                                        :             3
VXLAN-TEP                                       :             0
Total-TEP                                       :      3/ 16383

Mpls Dests (TEP, Egress Label + ES + ES-BMAC)   :             7
Mpls Etree Leaf Dests                           :             0
Vxlan Dests (TEP, Egress VNI + ES)              :             0
Total-Dest                                      :      7/196607

Sdp Bind +  Evpn Dests                          :      8/245759
ES L2/L3 PBR                                    :      0/ 32767
Evpn Etree Remote BUM Leaf Labels               :             0
```

To compute the DF election for EVI 1:

```
[/]
A:admin@PE-2# tools dump service system bgp-evpn ethernet-segment "ESI-23" evi 1 df
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
[02/25/2021 17:40:47] Computed DF: 192.0.2.3 (Remote) (Boot Timer Expired: Yes)
```

Conclusion

SR OS has a full RFC 7432 EVPN-MPLS implementation including single-active and all-active multi-homing. This example has shown how to configure and operate EVPN-MPLS for a simple non multi-homing configuration as well as a multi-homing configuration. Other topics, such as the integration of VPLS objects with EVPN-MPLS and proxy-ARP/ND, have also been discussed.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

EVPN for MPLS Tunnels in Epipe Services (EVPN-VPWS)

This chapter provides information about EVPN for MPLS tunnels in Epipe services (EVPN-VPWS).

Topics in this chapter include:

Applicability

This chapter was initially written for SR OS Release 14.0.R4, but the MD-CLI in the current edition is based on SR OS Release 22.10.R1. Ethernet Virtual Private Network - Virtual Private Wire Service (EVPN-VPWS) is supported in SR OS Release 14.0.R1 and later. EVPN-VPWS in multi-homing scenarios is supported in SR OS Release 14.0.R4 and later.

Chapter EVPN for MPLS Tunnels is prerequisite reading.

Overview

Service providers prefer an optimized, standardized, and unified control plane for VPNs. EVPN-VPWS is supported in MPLS networks that also run EVPN-MPLS in VPLS services. From a control plane perspective, EVPN-VPWS is a simplified point-to-point version of RFC 7432 – *BGP MPLS-Based Ethernet VPN*, because there is no need to advertise MAC routes in VPWS. EVPN-VPWS is described in RFC 8214 – *Virtual Private Wire Service Support in Ethernet VPN*.

EVPN-VPWS supports all-active multi-homing (per-flow load-balancing multi-homing) as well as single-active multi-homing (per-service load-balancing multi-homing), using the same Ethernet segments (ESs) used for EVPN-MPLS VPLS services. EVPN-VPWS uses route-type 1 and route-type 4; it does not use route-types 2, 3, or 5, because MAC/IP routes, inclusive multicast, or IP-prefix routes are not required.

The figure Figure 91: Route types and NLRIs for EVPN-VPWS shows the encoding of the required extensions for the route-types 1 and 4 for EVPN-VPWS.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 91: Route types and NLRIs for EVPN-VPWS

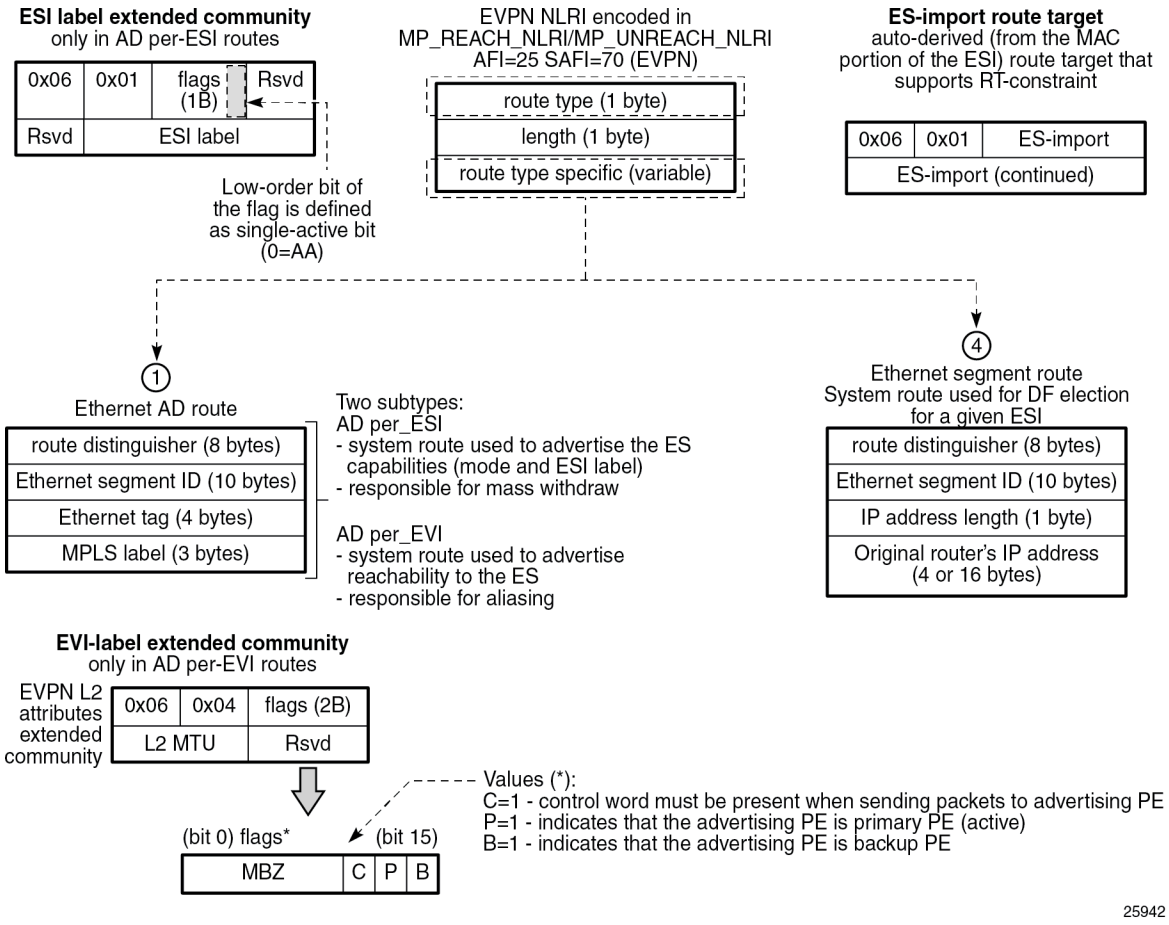

Two sub-types are defined for route-type 1. Route-type 4 has no sub-types. The route types used for EVPN-VPWS have the following purposes:

- Route-type 1 - Auto-discovery per EVPN instance (AD per-EVI). This route type is used in all EVPN-VPWS scenarios, with or without multi-homing. For EVPN-VPWS, the Ethernet tag field is encoded with the local Attachment Circuit (AC) of the advertising PE. This value is configured using the **service>epipe>bgp-evpn>local-attachment-circuit>eth-tag** <*value*> command. The route distinguisher (RD), MPLS label, and the Ethernet segment ID (ESI) are encoded as for EVPN-MPLS. The MPLS label field is used as service label. In case of multi-homing, AD per-EVI routes containing the same ESI are used to provide aliasing and a backup path to the PEs part of the ES. The L2 MTU is encoded with the service MTU configured in the Epipe. The following flags are used for EVPN-VPWS:
  - Flag C is set if a control word is configured in the service.
  - Flag P is set if the advertising PE is primary PE.
    - If no multi-homing is used, there is no primary PE (P=0).
    - In all-active multi-homing, all PEs in the ES are primary (P=1).
    - In single-active multi-homing, only one PE per-EVI in the ES is primary (P=1).

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

– Flag B is set if the advertising PE is backup PE.
  - The B-flag is only set in case of single-active multi-homing and only for one PE, even if more than two PEs are present in the same single-active ES. The backup PE is the winner of the second Designated Forwarder (DF) election (excluding the DF). The remaining non-DF PEs send B=0.

If there is no multi-homing, the ESI, flag P, and flag B will be zero.

- Route-type 1 - AD per Ethernet segment (AD per-ES). Same encoding as for EVPN-MPLS. AD per-ES is only used in multi-homing scenarios where it is advertised per ES from the PE. It carries the ESI label (used for split-horizon, but only for VPLS services and not for Epipe services) and can affect procedures such as the DF election, as well as the aliasing on remote PEs.
- Route-type 4 - ES route. Same encoding as for EVPN-MPLS. Route-type 4 is only used in multi-homing scenarios. This route advertises a local configured ES. The exchange of this route can discover remote PEs that are part of the same ES and the DF election algorithm among them.

Configuration

The figure Figure 92: EVPN-VPWS example topology shows the example topology that will be used throughout this chapter.

Figure 92: EVPN-VPWS example topology

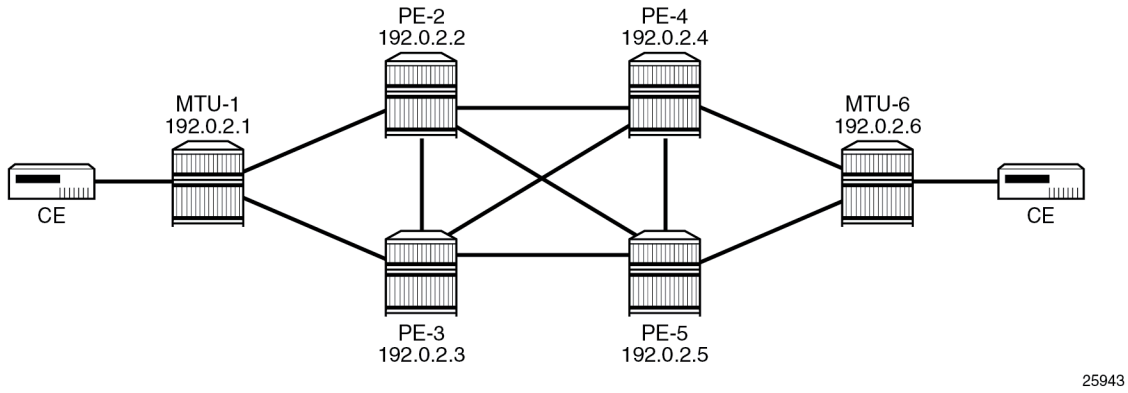

The example topology consists of six SR OS nodes with the following initial configuration:

- Network (or hybrid) ports interconnect the core PEs with configured router interfaces.
- MTU-1 is a pure Ethernet aggregator. The ports toward the core PEs are access ports. Likewise, the ports on PE-2 and PE-3 toward MTU-1 are access ports.
- Core PEs and MTU-6 run IS-IS on all router interfaces. Point-to-point adjacencies are established for the exchange of system IP addresses.
- Link LDP is configured between all PEs, and toward/from MTU-6.
- EVPN uses BGP for exchanging reachability at service level. Therefore, BGP peering sessions must be established among the core PEs for the EVPN family. Although typically a separate router is used, in this chapter, PE-2 is used as route reflector (RR) with the following BGP configuration:

```
# on RR PE-2:
configure {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            split-horizon true
            rapid-update {
                evpn true
            }
            group "internal" {
                peer-as 64500
                family {
                    evpn true
                }
                cluster {
                    cluster-id 192.0.2.2
                }
            }
            neighbor "192.0.2.3" {
                group "internal"
            }
            neighbor "192.0.2.4" {
                group "internal"
            }
            neighbor "192.0.2.5" {
                group "internal"
            }
        }
```

The BGP configuration on the other PEs is as follows:

```
# on PE-3, PE-4, PE-5:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            split-horizon true
            rapid-update {
                evpn true
            }
            group "internal" {
                peer-as 64500
                family {
                    evpn true
                }
            }
            neighbor "192.0.2.2" {
                group "internal"
            }
        }
```

The following EVPN-VPWS scenarios are described in the following sections:

- EVPN for MPLS tunnels in Epipe services without multi-homing
- EVPN for MPLS tunnels in Epipe services with all-active multi-homing
- EVPN for MPLS tunnels in Epipe services with single-active multi-homing

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

EVPN for MPLS tunnels in Epipe services without multi-homing

BGP-EVPN can be enabled in Epipe services with either SAPs or spoke-SDPs at the access, as shown in the figure Figure 93: Example topology for EVPN-VPWS without multi-homing.

Figure 93: Example topology for EVPN-VPWS without multi-homing

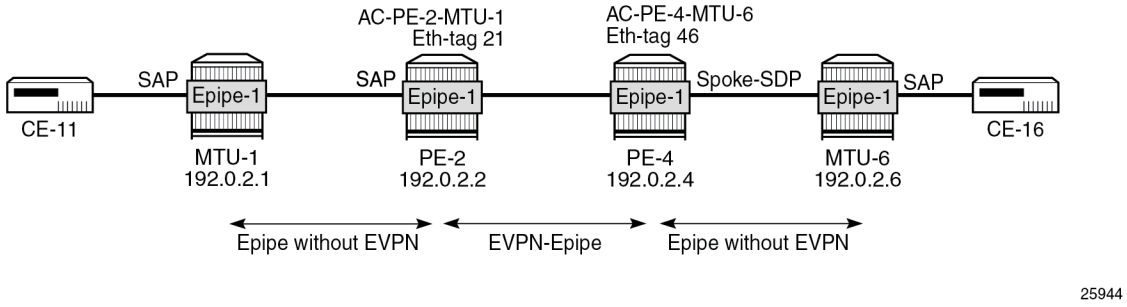

On PE-2, Epipe 1 is configured as follows:

```
# on PE-2:
configure {
    service {
        epipe "Epipe-1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
            }
            sap 1/1/c11/1:1 {
            }
            bgp-evpn {
                evi 1
                local-attachment-circuit "AC-PE-2-MTU-1" {
                    eth-tag 21
                }
                remote-attachment-circuit "AC-PE-4-MTU-6" {
                    eth-tag 46
                }
                mpls 1 {
                    admin-state enable
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
        }
```

On PE-4, the service configuration is as follows:

```
# on PE-4:
configure {
    service {
        epipe "Epipe-1" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
            }
            spoke-sdp 460:1 {
            }
            bgp-evpn {
                evi 1
                local-attachment-circuit "AC-PE-4-MTU-6" {
                    eth-tag 46
                }
                remote-attachment-circuit "AC-PE-2-MTU-1" {
                    eth-tag 21
                }
                mpls 1 {
                    admin-state enable
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
        }
        sdp 460 {
            admin-state enable
            far-end {
                ip-address 192.0.2.6
            }
        }
```

Where the following commands are relevant for the EVPN-VPWS configuration:

- **bgp 1** enables the context for the BGP configuration relevant to the service. The **bgp** context configures the common BGP parameters for all BGP families in the service, such as route distinguisher and route target. Even if the general BGP parameters for the service are auto-derived, the **bgp** context must be enabled.

```
[ex:/configure service epipe "Epipe-1"]
A:admin@PE-2# bgp 1 ?

 bgp

 adv-service-mtu       - Advertised service MTU value
 apply-groups          - Apply a configuration group at this level
 apply-groups-exclude  - Exclude a configuration group at this level
 pw-template-binding   + Enter the pw-template-binding list instance
 route-distinguisher   - High-order 6 bytes that are used as string to compose VSI-ID for
                         use in NLRI
 route-target          + Enter the route-target context
 vsi-export            - VSI export policies
 vsi-import            - VSI import policies
```

- The following parameters can be configured in the **bgp-evpn** context:

```
[ex:/configure service epipe "Epipe-1"]
A:admin@PE-2# bgp-evpn ?

 bgp-evpn

 apply-groups          - Apply a configuration group at this level
 apply-groups-exclude  - Exclude a configuration group at this level
 evi                   - EVPN ID
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
local-attachment-     + Enter the local-attachment-circuit list instance
 circuit
mpls                  + Enter the mpls list instance
remote-attachment-    + Enter the remote-attachment-circuit list instance
 circuit
segment-routing-v6    + Enter the segment-routing-v6 list instance
vxlan                 + Enter the vxlan list instance
```

– The **evi** is a two-byte or three-byte identifier used for auto-deriving the service RD (only for two-byte EVI), service RT, and for the DF election in multi-homing. The auto-derivation of RD and RT for a two-byte EVI is as follows:

  • RD <system IP address>:<evi>

  • RT <autonomous system number>:<evi>

  The EVI values must be unique in the system, regardless of the type of service they are assigned to (Epipe or VPLS).

  Note: Three-byte EVI values are supported in SR OS Release 21.10.R1 and later. For auto-derived RT as per RFC 8365, the **evi-three-byte-auto-rt** command must be configured, as described in the Three-byte EVI in EVPN Services chapter.

– The **local-attachment-circuit** and **remote-attachment-circuit** identify the two attachment circuits connected by the EVPN-VPWS service. The configured Ethernet tag for the local AC is advertised in the Ethernet tag field of the AD per-EVI route for the Epipe, along with the corresponding RD, RT, and MPLS label. Both local and remote Ethernet tags are mandatory to bring up the Epipe service. If the received Ethernet tag for the Epipe service matches the configured remote AC Ethernet tag, it will create an EVPN-MPLS destination to the next hop.

– The following configuration options are available for Epipes in the **bgp-evpn>mpls <bgp-instance>** context:

```
[ex:/configure service epipe "Epipe-1" bgp-evpn]
A:admin@PE-2# mpls 1 ?

 mpls

 admin-state           - Administrative state of BGP EVPN MPLS
 apply-groups          - Apply a configuration group at this level
 apply-groups-exclude  - Exclude a configuration group at this level
 auto-bind-tunnel      + Enter the auto-bind-tunnel context
 control-word          - Enable the CW bit in the label message
 default-route-tag     - Default route tag
 dynamic-egress-       - Enables dynamic egress label limit
  label-limit
 ecmp                  - Maximum ECMP routes information
 entropy-label         - Enable use of entropy-labels
 evi-three-byte-auto-  - Auto-derive the BGP EVPN route target
  rt
 force-vc-forwarding   - VC forwarding action
 oper-group            - Operational group identifier
 route-next-hop        + Enter the route-next-hop context
 send-tunnel-encap     + Enter the send-tunnel-encap context
```

This is a subset of the options for VPLS services; see chapter EVPN for MPLS Tunnels.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

When the local AC (SAP 1/1/c11/1:1) is up, PE-2 sends a BGP EVPN AD per-EVI route that contains Ethernet tag 21 for the local AC:

```
# on PE-2:
3 2022/11/30 11:33:31.729 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.4
"Peer 1: 192.0.2.4: UPDATE
Peer 1: 192.0.2.4 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 81
    Flag: 0x90 Type: 14 Len: 36 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
        Type: EVPN-AD Len: 25 RD: 192.0.2.2:1 ESI: ESI-0, tag: 21 Label: 8388512 (Raw Label:
 0x7fffa0) PathId:
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:1
        l2-attribute:MTU: 1514 C: 0 P: 0 B: 0
        bgp-tunnel-encap:MPLS
"
```

The auto-derived RD is 192.0.2.2:1 and the RT is 64500:1.

When the remote AC on PE-4 (spoke-SDP 460:1) is up, PE-2 receives the following BGP update from PE-4:

```
# on PE-2:
5 2022/11/30 11:33:50.377 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.4
"Peer 1: 192.0.2.4: UPDATE
Peer 1: 192.0.2.4 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 81
    Flag: 0x90 Type: 14 Len: 36 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.4
        Type: EVPN-AD Len: 25 RD: 192.0.2.4:1 ESI: ESI-0, tag: 46 Label: 8388512 (Raw Label:
 0x7fffa0) PathId:
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:1
        l2-attribute:MTU: 1514 C: 0 P: 0 B: 0
        bgp-tunnel-encap:MPLS
"
```

When the received RT matches and the received Ethernet tag matches the configured remote AC, the EVPN-MPLS destination (comprised of a termination endpoint (TEP) and egress label) is created on PE-2 and PE-4:

```
[/]
A:admin@PE-2# show service id 1 evpn-mpls

===============================================================================
BGP EVPN-MPLS Dest
===============================================================================
TEP Address                          Egr Label               Last Change
                                      Transport:Tnl-id
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
192.0.2.4                           524282                11/30/2022 11:33:50
                                    ldp:65538
-------------------------------------------------------------------------------
Number of entries : 1
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN-MPLS Ethernet Segment Dest
===============================================================================
Eth SegId                             Last Change
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

The MPLS label in the debug message is not the same as in the service, because the router will strip the extra four lowest bits to get the 20-bit MPLS label. The egress label for the EVPN-MPLS destination on PE-4 is 524282. The 24-bit label value in the BGP update debug is 16 (2^4) times as high: 524282*16 = 8388512. This is because the debug message is shown before the router can parse the label field and see if it corresponds to an MPLS label (20 bits) or a VXLAN VNI (24 bits).

The BGP AD per-EVI routes for Ethernet tag 46 can be shown with the following command:

```
[/]
A:admin@PE-2# show router bgp routes evpn auto-disc tag 46
===============================================================================
 BGP Router ID:192.0.2.2        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Auto-Disc Routes
===============================================================================
Flag  Route Dist.         ESI                           NextHop
      Tag                                               Label
-------------------------------------------------------------------------------
u*>i  192.0.2.4:1         ESI-0                         192.0.2.4
      46                                                LABEL 524282

-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

The following command shows the BGP EVPN information for Epipe 1:

```
[/]
A:admin@PE-2# show service id 1 bgp-evpn

===============================================================================
BGP EVPN Table
===============================================================================
EVI               : 1                 Creation Origin    : manual

-------------------------------------------------------------------------------
Local AC Name              Eth Tag  Endpoint                   Ingress Label
-------------------------------------------------------------------------------
AC-PE-2-MTU-1              21                                  0
-------------------------------------------------------------------------------
Number of local ACs : 1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
Remote AC Name                    Eth Tag  Endpoint
-------------------------------------------------------------------------------
AC-PE-4-MTU-6                     46
-------------------------------------------------------------------------------
Number of Remote ACs : 1
===============================================================================

===============================================================================
BGP EVPN MPLS Information
===============================================================================
Admin Status       : Enabled            Bgp Instance       : 1
Force Vlan Fwding  : Disabled
Force Qinq Fwding  : none
Route NextHop Type : system-ipv4
Control Word       : Disabled
Max Ecmp Routes    : 1
Entropy Label      : Disabled
Default Route Tag  : none
Oper Group         :
Evi 3-byte Auto-RT : Disabled
Dyn Egr Lbl Limit  : Disabled
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN MPLS Auto Bind Tunnel Information
===============================================================================
Allow-Flex-Algo-Fallback  : false
Resolution                : any               Strict Tnl Tag   : false
Max Ecmp Routes           : 1
Bgp Instance              : 1
Filter Tunnel Types       : (Not Specified)
Weighted Ecmp             : false
-------------------------------------------------------------------------------
===============================================================================
```

Note: Each PE sends its service MTU into the L2 MTU field in the L2-attribute in the AD per-EVI route for the Epipe service. The received L2 MTU will be checked. In case of a mismatch between the received MTU and the configured service MTU, the router will not set up the EVPN destination and, therefore, the service will not come up.

EVPN for MPLS tunnels in Epipe services with multi-homing

SR OS supports EVPN multi-homing as per RFC 8214.

The EVPN multi-homing implementation is based on the concept of the Ethernet segment (ES). An ES is a logical structure that can be defined in one or more PEs and identifies the CE (or access network) multi-homed to the EVPN PEs. An ES is associated with a port, LAG, or SDP object, and is shared by all the services defined on those objects. It can also be shared between Epipe and VPLS services.

Each ES has a unique Ethernet Segment Identifier (ESI) that is 10 bytes and is manually configured.

Note: Auto-derived EVPN ESI type 1 as per RFC 7432 is supported in SR OS Release 21.5.R1 and later, as described in the EVPN ESI Type 1 chapter.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The ESI is advertised in the control plane to all the PEs in an EVPN network; therefore, it is very important to ensure that the 10-byte ESI value is unique throughout the entire network. Single-homed CEs are assumed to be connected to an ES with ESI = 0 (single-homed ESs are not explicitly configured).

The ES is part of the base BGP-EVPN configuration and is not applied to any EVPN-MPLS service, by default. An ES can be shared by multiple services; the association of a specific SAP or spoke-SDP to an ES is automatically made when the SAP is defined in the same LAG or port configured in the ES, or when the spoke-SDP is defined in the same SDP configured in the ES.

Regardless of the multi-homing mode, the local Ethernet tag values must match on all the PEs that are part of the same ES. The PEs in the ES will use the AD per-EVI routes from the peer PEs to validate the PEs as DF election candidates for an EVI. The DF election is only relevant for single-active multi-homing ESs. For Epipes defined in an all-active multi-homing ES, there is no DF election required, because all PEs are forwarding traffic and all traffic is treated as unicast.

Aliasing is supported when sending traffic to an ES destination. Assuming ECMP is enabled on the ingress PE (and shared queuing or ingress policing), per-flow load-balancing will be performed among all the PEs that advertised P=1. PEs advertising P=0 are not considered as next hops for an ES destination.

The following sections show the configuration of:

- an all-active multi-homing ES with a LAG associated with it
- a single-active multi-homing ES linked to an SDP

The figure Figure 94: Example topology EVPN-VPWS with multi-homing shows an all-active ES and a single-active ES. The all-active multi-homing ES "AA-ESI-23" on PE-2 and PE-3 has a LAG associated to it; the single-active multi-homing ES "SA-ESI-45" on PE-4 and PE-5 has an SDP associated to it.

Figure 94: Example topology EVPN-VPWS with multi-homing

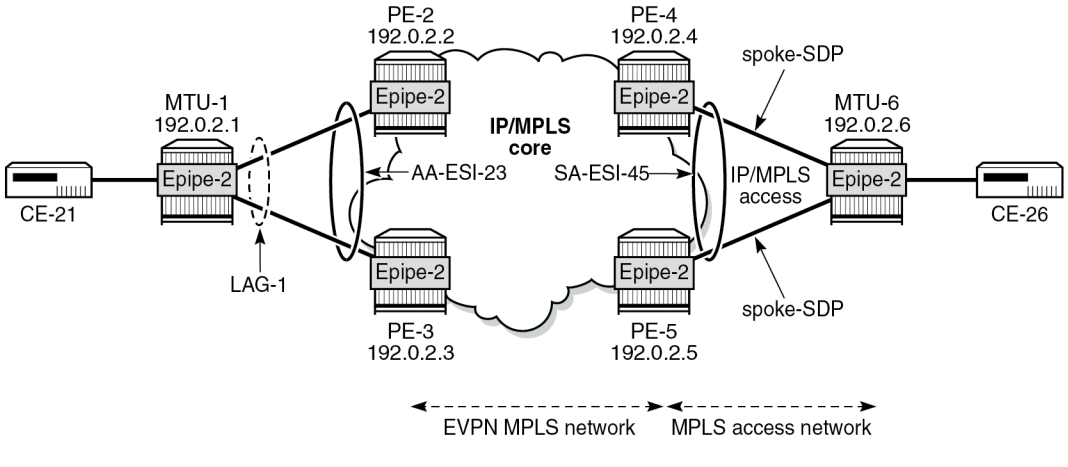

EVPN for MPLS tunnels in Epipe services with all-active multi-homing

All-active multi-homing allows for per-flow load-balancing. Unlike EVPN-MPLS in VPLS services, EVPN-VPWS has no DF election in all-active multi-homing. All PEs in the ES are active and the remote PE will

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

do per-flow load-balancing. ES "AA-ESI-23" is configured on PE-2 and PE-3 in all-active multi-homing with LAG 1 associated to it. This LAG is used as a SAP in Epipe 2 on both PE-2 and PE-3. The configuration of the ES and Epipe 2 is identical on PE-2 and PE-3, including the local AC and remote AC names and Ethernet tags:

```
# on PE-2, PE-3:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "AA-ESI-23" {
                        admin-state enable
                        esi 01:00:00:00:00:23:00:00:00:01
                        multi-homing-mode all-active
                        df-election {
                            es-activation-timer 3
                        }
                        association {
                            lag "lag-1" {
                            }
                        }
                    }
                }
            }
        }
        epipe "Epipe-2" {
            admin-state enable
            service-id 2
            customer "1"
            bgp 1 {
            }
            sap lag-1:2 {
            }
            bgp-evpn {
                evi 2
                local-attachment-circuit "AC-AA-ESI-23-MTU-1" {
                    eth-tag 231
                }
                remote-attachment-circuit "AC-SA-ESI-45-MTU-6" {
                    eth-tag 456
                }
                mpls 1 {
                    admin-state enable
                    ecmp 2
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
        }
```

See chapter EVPN for MPLS Tunnels for a detailed explanation of the configuration parameters of the ES.

In EVPN-VPWS multi-homing scenarios, three route types are exchanged: AD per-EVI, AD per-ES, and ES routes. The following ES route (route-type 4) for ESI 01:00:00:00:00:23:00:00:00:01 sent by PE-2 is imported at PE-3:

```
# on PE-3:
3 2022/11/30 11:44:28.466 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Received BGP UPDATE:
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    Withdrawn Length = 0
    Total Path Attr Length = 71
    Flag: 0x90 Type: 14 Len: 34 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
        Type: EVPN-ETH-SEG Len: 23 RD: 192.0.2.2:0 ESI: 01:00:00:00:00:23:00:00:00:01, IP-Len:
 4 Orig-IP-Addr: 192.0.2.2
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        df-election::DF-Type:Auto/DP:0/DF-Preference:0/AC:1
        target:00:00:00:00:23:00
"
```

The target 00:00:00:00:23:00 in the extended community is derived from the ESI (bytes 2 to 7) and is only imported by the PEs that are part of the same ES; that is, PE-2 and PE-3 in this example.

At the same time, the following AD per-ES route (route-type 1) with maximum Ethernet tag (MAX-ET, all Fs) and label 0 is sent by RR PE-2 and imported by the rest of the PEs. The following two BGP updates with MAX-ET are received by PE-4:

```
# on PE-4:
6 2022/11/30 11:44:28.466 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 81
    Flag: 0x90 Type: 14 Len: 36 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
        Type: EVPN-AD Len: 25 RD: 192.0.2.2:2 ESI: 01:00:00:00:00:23:00:00:00:01, tag: MAX-ET
 Label: 0 (Raw Label: 0x0) PathId:
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:2
        esi-label:524280/All-Active
        bgp-tunnel-encap:MPLS
"
```

```
8 2022/11/30 11:44:30.124 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 95
    Flag: 0x90 Type: 14 Len: 36 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.3
        Type: EVPN-AD Len: 25 RD: 192.0.2.3:2 ESI: 01:00:00:00:00:23:00:00:00:01, tag: MAX-ET
 Label: 0 (Raw Label: 0x0) PathId:
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0x80 Type: 9 Len: 4 Originator ID: 192.0.2.3
    Flag: 0x80 Type: 10 Len: 4 Cluster ID:
        192.0.2.2
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:2
        esi-label:524281/All-Active
        bgp-tunnel-encap:MPLS
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
"
```

The ESI label is in the extended community, as well as the indication that the multi-homing is all-active. Epipe services do not require ESI labels because BUM traffic is not recognized as such in EVPN-VPWS services. However, because the ES can be shared by Epipe and VPLS services, the AD per-ES route still includes a non-zero ESI label.

The following AD per-EVI routes (route-type 1) with Ethernet tag 231 sent by RR PE-2 are received and imported on PE-4:

```
# on PE-4:
5 2022/11/30 11:44:28.466 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 81
    Flag: 0x90 Type: 14 Len: 36 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
        Type: EVPN-AD Len: 25 RD: 192.0.2.2:2 ESI: 01:00:00:00:00:23:00:00:00:01, tag: 231
 Label: 8388496 (Raw Label: 0x7fff90) PathId:
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:2
        l2-attribute:MTU: 1514 C: 0 P: 1 B: 0
        bgp-tunnel-encap:MPLS
"
```

```
7 2022/11/30 11:44:30.124 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 95
    Flag: 0x90 Type: 14 Len: 36 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.3
        Type: EVPN-AD Len: 25 RD: 192.0.2.3:2 ESI: 01:00:00:00:00:23:00:00:00:01, tag: 231
 Label: 8388512 (Raw Label: 0x7fffa0) PathId:
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0x80 Type: 9 Len: 4 Originator ID: 192.0.2.3
    Flag: 0x80 Type: 10 Len: 4 Cluster ID:
        192.0.2.2
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:2
        l2-attribute:MTU: 1514 C: 0 P: 1 B: 0
        bgp-tunnel-encap:MPLS
"
```

This route contains the flags for control word (C), primary (P), and backup (B). In all-active multi-homing, all nodes are primary (P=1).

PE-4 has learned AD per-EVI/ES routes for AA-ESI-23 from PE-2 and PE-3, as shown in the following output:

```
[/]
A:admin@PE-4# show router bgp routes evpn auto-disc esi 01:00:00:00:00:23:00:00:00:01
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
 BGP Router ID:192.0.2.4        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Auto-Disc Routes
===============================================================================
Flag  Route Dist.         ESI                           NextHop
      Tag                                               Label
-------------------------------------------------------------------------------
u*>i  192.0.2.2:2         01:00:00:00:00:23:00:00:00:01 192.0.2.2
      231                                               LABEL 524281

u*>i  192.0.2.2:2         01:00:00:00:00:23:00:00:00:01 192.0.2.2
      MAX-ET                                            LABEL 0

u*>i  192.0.2.3:2         01:00:00:00:00:23:00:00:00:01 192.0.2.3
      231                                               LABEL 524282

u*>i  192.0.2.3:2         01:00:00:00:00:23:00:00:00:01 192.0.2.3
      MAX-ET                                            LABEL 0

-------------------------------------------------------------------------------
Routes : 4
===============================================================================
```

For Epipe 2 on PE-4, the EVPN MPLS destination is not pointing at a specific TEP, but AA-ESI-23, as shown in the following output:

```
[/]
A:admin@PE-4# show service id 2 evpn-mpls

===============================================================================
BGP EVPN-MPLS Dest
===============================================================================
TEP Address                        Egr Label              Last Change
                                    Transport:Tnl-id
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================

===============================================================================
BGP EVPN-MPLS Ethernet Segment Dest
===============================================================================
Eth SegId                             Last Change
-------------------------------------------------------------------------------
01:00:00:00:00:23:00:00:00:01         11/30/2022 11:44:50
-------------------------------------------------------------------------------
Number of entries: 1
-------------------------------------------------------------------------------
===============================================================================
```

When ECMP > 1 on the ingress PE, multiple TEPs can correspond to a specific ESI (aliasing). In this case, ECMP=2 and PE-4 and PE-5 have two TEP addresses and egress labels for ESI 01:00:00:00:00:23:00:00:00:01, as shown for PE-4:

```
[/]
A:admin@PE-4# show service id 2 evpn-mpls esi esi-1 01:00:00:00:00:23:00:00:00:01
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
BGP EVPN-MPLS Ethernet Segment Dest
===============================================================================
Eth SegId                              Last Change
-------------------------------------------------------------------------------
01:00:00:00:00:23:00:00:00:01          11/30/2022 11:44:50
===============================================================================

===============================================================================
BGP EVPN-MPLS Dest TEP Info
===============================================================================
TEP Address             Egr Label               Last Change
                        Transport:Tnl-Id
-------------------------------------------------------------------------------
192.0.2.2               524281                  11/30/2022 11:44:50
                        ldp:65538
192.0.2.3               524282                  11/30/2022 11:44:50
                        ldp:65537
-------------------------------------------------------------------------------
Number of entries : 2
-------------------------------------------------------------------------------
===============================================================================
```

Note: Even if ECMP is configured, the ingress router will not load-balance the traffic unless shared queuing or ingress policing is configured. This is not specific to EVPN, but generic to the way Epipes forward traffic.

In all-active multi-homing for EVPN-VPWS, there is no DF election and all PEs in the ES are active. For AA-ESI-23, both PE-2 and PE-3 are active/primary/DF, but there are no DF candidates, because there is no DF election:

```
[/]
A:admin@PE-2# show service system bgp-evpn ethernet-segment name "AA-ESI-23" evi evi-1 2

===============================================================================
EVI DF and Candidate List
===============================================================================
EVI           SvcId         Actv Timer Rem      DF  DF Last Change
-------------------------------------------------------------------------------
2             2             0                   yes 11/30/2022 11:44:28
===============================================================================

===============================================================================
DF Candidates                          Time Added           Oper Pref  Do Not
                                                              Value    Preempt
-------------------------------------------------------------------------------
No entries found
===============================================================================
```

Similarly, on PE-3:

```
[/]
A:admin@PE-3# show service system bgp-evpn ethernet-segment name "AA-ESI-23" evi evi-1 2

===============================================================================
EVI DF and Candidate List
===============================================================================
EVI           SvcId         Actv Timer Rem      DF  DF Last Change
-------------------------------------------------------------------------------
2             2             0                   yes 11/30/2022 11:44:30
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
DF Candidates                           Time Added           Oper Pref  Do Not
                                                               Value    Preempt
-------------------------------------------------------------------------------
No entries found
===============================================================================
```

To confirm that all-active multi-homing is working correctly, the following command shows all information related to a specific ESI; in this case, AA-ESI-23 on PE-2:

```
[/]
A:admin@PE-2# show service system bgp-evpn ethernet-segment name "AA-ESI-23" all

===============================================================================
Service Ethernet Segment
===============================================================================
Name                    : AA-ESI-23
Eth Seg Type            : None
Admin State             : Enabled            Oper State         : Up
ESI                     : 01:00:00:00:00:23:00:00:00:01
Oper ESI                : 01:00:00:00:00:23:00:00:00:01
Auto-ESI Type           : None
AC DF Capability        : Include
Multi-homing            : allActive          Oper Multi-homing  : allActive
ES SHG Label            : 524280
Source BMAC LSB         : None
Lag                     : lag-1
ES Activation Timer     : 3 secs
Oper Group              : (Not Specified)
Svc Carving             : auto               Oper Svc Carving   : auto
Cfg Range Type          : primary
Vprn NextHop EVI Ranges : <none>
===============================================================================

===============================================================================
EVI Information
===============================================================================
EVI                 SvcId               Actv Timer Rem      DF
-------------------------------------------------------------------------------
2                   2                   0                   yes
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
---snip---
```

EVPN for MPLS tunnels in Epipe services with single-active multi-homing

Single-active multi-homing allows for per-service load-balancing. Single-active multi-homing is configured on PE-4 and PE-5 with ES "SA-ESI-45". Both PEs have an SDP to MTU-6, which is associated with the ES and to the Epipe service. The configuration of the local and remote AC names and Ethernet tags is identical on PE-4 and PE-5.

On PE-4, the service configuration is as follows:

```
# on PE-4:
configure {
    service {
        system {
            bgp {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                evpn {
                    ethernet-segment "SA-ESI-45" {
                        admin-state enable
                        esi 01:00:00:00:00:45:00:00:00:01
                        multi-homing-mode single-active
                        df-election {
                            es-activation-timer 3
                        }
                        association {
                            sdp 46 {
                            }
                        }
                    }
                }
            }
        }
        epipe "Epipe-2" {
            admin-state enable
            service-id 2
            customer "1"
            bgp 1 {
            }
            spoke-sdp 46:2 {
            }
            bgp-evpn {
                evi 2
                local-attachment-circuit "AC-SA-ESI-45-MTU-6" {
                    eth-tag 456
                }
                remote-attachment-circuit "AC-AA-ESI-23-MTU-1" {
                    eth-tag 231
                }
                mpls 1 {
                    admin-state enable
                    ecmp 2
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
        }
        sdp 46 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.6
            }
        }
```

On PE-5, the configuration is similar, but with a different SDP:

```
# on PE-5:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "SA-ESI-45" {
                        admin-state enable
                        esi 01:00:00:00:00:45:00:00:00:01
                        multi-homing-mode single-active
                        df-election {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                            es-activation-timer 3
                        }
                        association {
                            sdp 56 {
                            }
                        }
                    }
                }
            }
        }
        epipe "Epipe 2" {
            admin-state enable
            service-id 2
            customer "1"
            bgp 1 {
            }
            spoke-sdp 56:2 {
            }
            bgp-evpn {
                evi 2
                local-attachment-circuit "AC-SA-ESI-45-MTU-6" {
                    eth-tag 456
                }
                remote-attachment-circuit "AC-AA-ESI-23-MTU-1" {
                    eth-tag 231
                }
                mpls 1 {
                    admin-state enable
                    ecmp 2
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
        }
        sdp 56 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.6
            }
        }
```

Three route types will be exchanged between the core PEs: AD per-EVI, AD per-ES, and ES routes.

PE-4 and PE-5 advertise ES routes that are only imported by them. As an example, the following is the ES route with originator PE-4 sent by RR PE-2 to PE-5. It contains a target 00:00:00:00:45:00 in the extended community that is derived from the ESI:

```
# on PE-2:
56 2022/11/30 11:45:03.406 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.5
"Peer 1: 192.0.2.5: UPDATE
Peer 1: 192.0.2.5 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 85
    Flag: 0x90 Type: 14 Len: 34 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.4
        Type: EVPN-ETH-SEG Len: 23 RD: 192.0.2.4:0 ESI: 01:00:00:00:00:45:00:00:00:01, IP-Len:
 4 Orig-IP-Addr: 192.0.2.4
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0x80 Type: 9 Len: 4 Originator ID: 192.0.2.4
    Flag: 0x80 Type: 10 Len: 4 Cluster ID:
        192.0.2.2
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        df-election::DF-Type:Auto/DP:0/DF-Preference:0/AC:1
        target:00:00:00:00:45:00
"
```

The AD per-ES route has a maximum Ethernet tag (MAX-ET) and an ESI label in the extended community. The multi-homing mode is single-active. As in the case of all-active multi-homing, the ESI label is not used in Epipe services. The following BGP update with originator PE-4 is sent by RR PE-2 to its client PE-5:

```
# on PE-2:
36 2022/11/30 11:44:47.394 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.5
"Peer 1: 192.0.2.5: UPDATE
Peer 1: 192.0.2.5 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 95
    Flag: 0x90 Type: 14 Len: 36 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.4
        Type: EVPN-AD Len: 25 RD: 192.0.2.4:2 ESI: 01:00:00:00:00:45:00:00:00:01, tag: MAX-ET
 Label: 0 (Raw Label: 0x0) PathId:
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0x80 Type: 9 Len: 4 Originator ID: 192.0.2.4
    Flag: 0x80 Type: 10 Len: 4 Cluster ID:
        192.0.2.2
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:2
        esi-label:524279/Single-Active
        bgp-tunnel-encap:MPLS
"
```

The AD per-EVI route contains flags for primary and backup, which will be different for routes received from PE-4 and PE-5. In this case, PE-4 is primary in the single-active multi-homing ES (P=1):

```
# on PE-2:
64 2022/11/30 11:45:06.415 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.5
"Peer 1: 192.0.2.5: UPDATE
Peer 1: 192.0.2.5 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 95
    Flag: 0x90 Type: 14 Len: 36 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.4
        Type: EVPN-AD Len: 25 RD: 192.0.2.4:2 ESI: 01:00:00:00:00:45:00:00:00:01, tag: 456
 Label: 8388480 (Raw Label: 0x7fff80) PathId:
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0x80 Type: 9 Len: 4 Originator ID: 192.0.2.4
    Flag: 0x80 Type: 10 Len: 4 Cluster ID:
        192.0.2.2
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:2
        l2-attribute:MTU: 1514 C: 0 P: 1 B: 0
        bgp-tunnel-encap:MPLS
"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

PE-5 is backup in the single-active multi-homing ES (B=1):

```
# on PE-2:
72 2022/11/30 11:45:10.872 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.5
"Peer 1: 192.0.2.5: UPDATE
Peer 1: 192.0.2.5 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 81
    Flag: 0x90 Type: 14 Len: 36 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.5
        Type: EVPN-AD Len: 25 RD: 192.0.2.5:2 ESI: 01:00:00:00:00:45:00:00:00:01, tag: 456
 Label: 8388512 (Raw Label: 0x7fffa0) PathId:
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:2
        l2-attribute:MTU: 1514 C: 0 P: 0 B: 1
        bgp-tunnel-encap:MPLS
"
```

The BGP EVPN AD routes can be shown with the following command:

```
[/]
A:admin@PE-2# show router bgp routes evpn auto-disc esi 01:00:00:00:00:45:00:00:00:01
===============================================================================
 BGP Router ID:192.0.2.2        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Auto-Disc Routes
===============================================================================
Flag  Route Dist.         ESI                           NextHop
      Tag                                               Label
-------------------------------------------------------------------------------
u*>i  192.0.2.4:2         01:00:00:00:00:45:00:00:00:01 192.0.2.4
      456                                               LABEL 524280

u*>i  192.0.2.4:2         01:00:00:00:00:45:00:00:00:01 192.0.2.4
      MAX-ET                                            LABEL 0

u*>i  192.0.2.5:2         01:00:00:00:00:45:00:00:00:01 192.0.2.5
      456                                               LABEL 524282

u*>i  192.0.2.5:2         01:00:00:00:00:45:00:00:00:01 192.0.2.5
      MAX-ET                                            LABEL 0

-------------------------------------------------------------------------------
Routes : 4
===============================================================================
```

For each PE in the single-active ES, there are two AD routes: the routes with MAX-ET are AD per-ES routes and the routes with a configured Ethernet tag are AD per-EVI routes.

The EVPN MPLS destination for Epipe 2 on PE-2 is SA-ESI-45, as shown in the following output:

```
[/]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
A:admin@PE-2# show service id 2 evpn-mpls

===============================================================================
BGP EVPN-MPLS Dest
===============================================================================
TEP Address                          Egr Label              Last Change
                                      Transport:Tnl-id
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================

===============================================================================
BGP EVPN-MPLS Ethernet Segment Dest
===============================================================================
Eth SegId                              Last Change
-------------------------------------------------------------------------------
01:00:00:00:00:45:00:00:00:01          11/30/2022 11:45:06
-------------------------------------------------------------------------------
Number of entries: 1
-------------------------------------------------------------------------------
===============================================================================
```

The ESI is resolved to the TEP address of the primary (DF) PE-4, as follows:

```
[/]
A:admin@PE-2# show service id 2 evpn-mpls esi esi-1 01:00:00:00:00:45:00:00:00:01

===============================================================================
BGP EVPN-MPLS Ethernet Segment Dest
===============================================================================
Eth SegId                              Last Change
-------------------------------------------------------------------------------
01:00:00:00:00:45:00:00:00:01          11/30/2022 11:45:06
===============================================================================

===============================================================================
BGP EVPN-MPLS Dest TEP Info
===============================================================================
TEP Address              Egr Label                Last Change
                         Transport:Tnl-Id
-------------------------------------------------------------------------------
192.0.2.4                524280                   11/30/2022 11:45:06
                         ldp:65538
-------------------------------------------------------------------------------
Number of entries : 1
-------------------------------------------------------------------------------
===============================================================================
```

The DF election is key for the forwarding and backup functions in single-active multi-homing ESs. The PE elected as DF will be the primary for the ES in the Epipe and will unblock the SAP/spoke-SDP for upstream and downstream traffic. The rest of the PEs in the ES will bring their ES SAPs or spoke-SDPs operationally down.

PE-5 is a non-DF, as follows:

```
[/]
A:admin@PE-5# show service system bgp-evpn ethernet-segment name "SA-ESI-45" evi evi-1 2

===============================================================================
EVI DF and Candidate List
===============================================================================
EVI           SvcId         Actv Timer Rem      DF  DF Last Change
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
2             2             0                    no  11/30/2022 11:44:55
===============================================================================

===============================================================================
DF Candidates                              Time Added           Oper Pref  Do Not
                                                                  Value    Preempt
-------------------------------------------------------------------------------
192.0.2.4                                  11/30/2022 11:45:03  0          Disabl*
192.0.2.5                                  11/30/2022 11:45:08  0          Disabl*
-------------------------------------------------------------------------------
Number of entries: 2
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

In single-active multi-homing, the service spoke-SDP (or SAP) is brought operationally down on the non-DF, as shown in the following output:

```
[/]
A:admin@PE-5# show service id 2 sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
56:2             Spok     192.0.2.6       Up      Down      524280    524280
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
```

The spoke-SDP 56:2 is operationally down with a StandbyForMHProtocol flag:

```
[/]
A:admin@PE-5# show service id 2 sdp 56:2 detail | match Flag
Flags              : StandbyForMHProtocol
```

Two consecutive DF elections take place: the first DF election includes all PEs in the ES for that Epipe and determines which PE is the primary PE (flags P=1, B=0). The second DF election excludes this DF and determines which PE is the backup (P=0, B=1). All other PEs signal flags P=0 and B=0.

When the primary PE fails, AD per-ES/EVI withdrawal messages are sent to the remote PE, which will update its next hop to the backup. The backup PE takes over immediately without waiting for the **es-activation-timer** to bring up its SAP/spoke-SDP.

Ethernet segment failures

When the SDP toward the primary (DF) fails, the backup PE needs to take over. An SDP failure is emulated and log 99 on PE-4 shows that SDP 46 is operational down and PE-4 is no longer the DF:

```
140 2022/11/30 12:09:36.765 CET MINOR: SVCMGR #2094 Base
"Ethernet Segment:SA-ESI-45, EVI:2, Designated Forwarding state changed to:false"

139 2022/11/30 12:09:36.764 CET MINOR: SVCMGR #2326 Base
"Status of SDP Bind 46:2 in service 2 (customer 1) local PW status bits changed to psnIngress
Fault psnEgressFault "
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
138 2022/11/30 12:09:36.764 CET MINOR: SVCMGR #2303 Base
"Status of SDP 46 changed to admin=up oper=down"
```

Remote PEs receive route withdrawal updates (unreachable NLRI) from former DF PE-4, for example on RR PE-2:

```
# on PE-2:
76 2022/11/30 12:09:36.765 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.4
"Peer 1: 192.0.2.4: UPDATE
Peer 1: 192.0.2.4 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 34
    Flag: 0x90 Type: 15 Len: 30 Multiprotocol Unreachable NLRI:
        Address Family EVPN
        Type: EVPN-AD Len: 25 RD: 192.0.2.4:2 ESI: 01:00:00:00:00:45:00:00:00:01, tag: MAX-ET
 Label: 0 (Raw Label: 0x0) PathId:
"

75 2022/11/30 12:09:36.765 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.4
"Peer 1: 192.0.2.4: UPDATE
Peer 1: 192.0.2.4 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 59
    Flag: 0x90 Type: 15 Len: 55 Multiprotocol Unreachable NLRI:
        Address Family EVPN
        Type: EVPN-AD Len: 25 RD: 192.0.2.4:2 ESI: 01:00:00:00:00:45:00:00:00:01, tag: 456
 Label: 0 (Raw Label: 0x0) PathId:
        Type: EVPN-ETH-SEG Len: 23 RD: 192.0.2.4:0 ESI: 01:00:00:00:00:45:00:00:00:01, IP-Len:
 4 Orig-IP-Addr: 192.0.2.4
"
```

The backup PE-5 is promoted to primary (P=1, B=0) and sends BGP updates accordingly. The following AD per-EVI is received on PE-2:

```
# on PE-2:
79 2022/11/30 12:09:36.768 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.5
"Peer 1: 192.0.2.5: UPDATE
Peer 1: 192.0.2.5 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 81
    Flag: 0x90 Type: 14 Len: 36 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.5
        Type: EVPN-AD Len: 25 RD: 192.0.2.5:2 ESI: 01:00:00:00:00:45:00:00:00:01, tag: 456
 Label: 8388512 (Raw Label: 0x7fffa0) PathId:
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:2
        l2-attribute:MTU: 1514 C: 0 P: 1 B: 0
        bgp-tunnel-encap:MPLS
"
```

PE-5 brings up its spoke-SDP without waiting for the **es-activation-timer** and takes over immediately. It is now the only DF candidate, and therefore the DF, as follows:

```
[/]
A:admin@PE-5# show service system bgp-evpn ethernet-segment name "SA-ESI-45" evi evi-1 2

===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
EVI DF and Candidate List
===============================================================================
EVI           SvcId         Actv Timer Rem      DF  DF Last Change
-------------------------------------------------------------------------------
2             2             0                   yes 11/30/2022 11:44:55
===============================================================================

===============================================================================
DF Candidates                              Time Added           Oper Pref  Do Not
                                                                  Value    Preempt
-------------------------------------------------------------------------------
192.0.2.5                                  11/30/2022 11:45:08  0          Disabl*
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

BGP updates are exchanged and the remote PEs will resolve the ESI to the TEP address 192.0.2.5. For example, on PE-2:

```
[/]
A:admin@PE-2# show service id 2 evpn-mpls esi esi-1 01:00:00:00:00:45:00:00:00:01

===============================================================================
BGP EVPN-MPLS Ethernet Segment Dest
===============================================================================
Eth SegId                                  Last Change
-------------------------------------------------------------------------------
01:00:00:00:00:45:00:00:00:01              11/30/2022 12:09:37
===============================================================================

===============================================================================
BGP EVPN-MPLS Dest TEP Info
===============================================================================
TEP Address              Egr Label                Last Change
                         Transport:Tnl-Id
-------------------------------------------------------------------------------
192.0.2.5                524282                   11/30/2022 12:09:37
                         ldp:65539
-------------------------------------------------------------------------------
Number of entries : 1
-------------------------------------------------------------------------------
===============================================================================
```

This process is revertive; as soon as the SDP 46 is operationally up again, a new DF election is triggered with two DF candidates and PE-4 will be elected as DF.

Troubleshooting and debugging

The following show and debug commands can be used in EVPN-VPWS:

- **show redundancy bgp-evpn-multi-homing**
- **show router bgp routes evpn** (and filters)
- **show service evpn-mpls** *[<TEP ip-address>]*
- **show service id bgp-evpn**
- **show service id evpn-mpls** (and modifiers)
- **show service system bgp-evpn**

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- **show service system bgp-evpn ethernet-segment** (and modifiers)
- **debug router bgp update**
- **show log log-id 99**

Most of these commands have been shown in the preceding sections; some commands are shown in this section.

Information about the configured boot timers (before DF election) and ES activation timer (after the system has been elected DF) can be shown as follows:

```
[/]
A:admin@PE-2# show redundancy bgp-evpn-multi-homing

===============================================================================
Redundancy BGP EVPN Multi-homing Information
===============================================================================
Boot-Timer              : 10 secs
Boot-Timer Remaining    : 0 secs
ES Activation Timer     : 3 secs
===============================================================================
```

See chapter EVPN for MPLS Tunnels for a description of these timers.

The following command shows that the BGP route-type 4 (ES route) messages are only imported by the PEs in the same ES; for example, on PE-3:

```
[/]
A:admin@PE-3# show router bgp routes evpn eth-seg
===============================================================================
 BGP Router ID:192.0.2.3        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Eth-Seg Routes
===============================================================================
Flag  Route Dist.         ESI                           NextHop
      OrigAddr
-------------------------------------------------------------------------------
u*>i  192.0.2.2:0         01:00:00:00:00:23:00:00:00:01 192.0.2.2
      192.0.2.2

-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

On PE-4:

```
[/]
A:admin@PE-4# show router bgp routes evpn eth-seg
===============================================================================
 BGP Router ID:192.0.2.4        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
BGP EVPN Eth-Seg Routes
===============================================================================
Flag  Route Dist.         ESI                          NextHop
      OrigAddr
-------------------------------------------------------------------------------
u*>i  192.0.2.5:0         01:00:00:00:00:45:00:00:00:01 192.0.2.5
      192.0.2.5

-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

The following command shows all the EVPN MPLS destinations toward TEP 192.0.2.4. Epipe 1 has an EVPN MPLS destination toward TEP 192.0.2.4 directly and Epipe 2 has an EVPN MPLS destination to SA-ESI-45, which can be resolved to TEP 192.0.2.4. This is shown in the following output:

```
[/]
A:admin@PE-2# show service evpn-mpls 192.0.2.4

===============================================================================
BGP EVPN-MPLS Dest
===============================================================================
Service Id                              Egr Label                     Instance
-------------------------------------------------------------------------------
1                                       524282                        1
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN-MPLS Ethernet Segment Dest
===============================================================================
Service Id          Eth Seg Id                              Egr Label
-------------------------------------------------------------------------------
2                   01:00:00:00:00:45:00:00:00:01           524280
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN-MPLS ES BMac Dest
===============================================================================
Service Id               ES BMac                 Egr Label
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

The following command lists all configured ESs on the system:

```
[/]
A:admin@PE-2# show service system bgp-evpn ethernet-segment

===============================================================================
Service Ethernet Segment
===============================================================================
Name                            ESI                           Admin     Oper
-------------------------------------------------------------------------------
AA-ESI-23                       01:00:00:00:00:23:00:00:00:01 Enabled   Up
-------------------------------------------------------------------------------
Entries found: 1
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

In addition to the preceding commands, the following tools dump commands may be useful:

- **tools dump service evpn usage** – This command shows the number of EVPN-MPLS (and EVPN-VXLAN) destinations in the system.
- **tools dump service system bgp-evpn ethernet-segment <name> evi <..> df** – This command computes the DF election for a specific ESI and EVI. For all-active, there is no DF election and all PEs forward traffic. For single-active, one PE will be active for a service while another PE will be backup. This command shows the DF (primary), even if it is not the local PE.

The usage of EVPN resources can be shown as follows:

```
[/]
A:admin@PE-2# tools dump service evpn usage

vxlan-srv6-evpn-mpls usage statistics at 11/30/2022 12:15:35:

MPLS-TEP                                        :             1
VXLAN-TEP                                       :             0
SRV6-TEP                                        :             0
Total-TEP                                       :      1/ 16383

Mpls Dests (TEP, Egress Label + ES + ES-BMAC)   :             2
Mpls Etree Leaf Dests                           :             0
Vxlan Dests (TEP, Egress VNI + ES)              :             0
Srv6 Dests (TEP, SID  + ES)                     :             0
Total-Dest                                      :      2/196607

Sdp Bind +  Evpn Dests                          :      2/245759
ES L2/L3 PBR                                    :      0/ 32767
Evpn Etree Remote BUM Leaf Labels               :             0
```

On PE-2, there is one MPLS-TEP (192.0.2.4 in Epipe 1 and Epipe 2) and there are two MPLS destinations: 192.0.2.4 and ESI 01:00:00:00:00:45:00:00:00:01. PE-5 is not an MPLS-TEP for PE-2, because it is not a primary and, therefore, not forwarding any traffic.

In all-active multi-homing, the DF election is not applicable:

```
[/]
A:admin@PE-2# tools dump service system bgp-evpn ethernet-segment "AA-ESI-23" evi 2 df

[11/30/2022 12:15:50] All Active VPWS or IP-ALIASING - DF N/A
```

In single-active multi-homing, the following command shows which PE is the DF:

```
[/]
A:admin@PE-5# tools dump service system bgp-evpn ethernet-segment "SA-ESI-45" evi 2 df

[11/30/2022 12:16:04] Computed DF: 192.0.2.4 (Remote) (Boot Timer Expired: Yes)
[11/30/2022 12:16:04] Computed Backup: 192.0.2.5 (This Node)
```

The command is launched on PE-5, which is a backup. The computed DF is PE-4 and the boot timer has expired, meaning there is no DF re-election pending.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Conclusion

EVPN-VPWS is a simplified point-to-point version of RFC 7432 - *BGP MPLS-Based Ethernet VPN*. When used for Epipe and VPLS services, EVPN provides a unified control plane mechanism that simplifies the network deployment and operation. Single-active and all-active multi-homing can be used in Epipes; EVPN-VPWS is a differentiator of EVPN compared to traditional TLDP or BGP Epipe redundancy mechanisms. The Ethernet Segments used for multi-homing can be shared between EVPN VPLS and EVPN Epipes.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

EVPN for MPLS Tunnels in Routed VPLS

This chapter provides information about EVPN for MPLS tunnels in routed VPLS.

Topics in this chapter include:

Applicability

This chapter was initially written for SR OS Release 15.0.R4, but the MD-CLI in the current edition is based on SR OS Release 21.10.R3. EVPN-MPLS and IP-prefix advertisement in routed VPLS (R-VPLS) without Multi-homing (MH) is supported in SR OS Release 14.0.R1, and later. EVPN-MPLS and IP-prefix advertisement in R-VPLS with all-active and single-active MH is supported in SR OS Release 14.0.R4, and later. Virtual Router Redundancy Protocol (VRRP) in passive mode is also supported in SR OS Release 14.0.R4, and later.

Chapter EVPN for VXLAN Tunnels (Layer 3) is prerequisite reading.

Overview

The EVPN-MPLS in R-VPLS feature matches the EVPN-VXLAN in R-VPLS feature, which is described in chapter EVPN for VXLAN Tunnels (Layer 3) The following capabilities are supported in an R-VPLS service where **bgp-evpn mpls** is enabled:

- R-VPLS with Virtual Router Redundancy Protocol (VRRP) support on the VPRN interfaces
- R-VPLS support including IP route advertisement (IP prefix routes — BGP-EVPN route type 5) with regular interfaces
- R-VPLS support including IP route advertisement with **evpn-tunnel** interfaces
- R-VPLS with IPv6 support on the VPRN IP interface

All-active and single-active MH Ethernet segments (ESs) are supported in R-VPLS. When Ethernet Segments (ESs) are used along with R-VPLS services in two or more PEs, passive VRRP provides an "anycast default gateway" that optimizes inter-subnet forwarding for hosts in the R-VPLS. Passive VRRP is described in the following section.

Passive VRRP

VRRP can be configured in passive mode, which suppresses the transmission and reception of keepalive messages. Passive VRRP can be configured in the base router, in an IES, or in a VPRN, using the following commands:

```
[/]
A:admin@PE-2# tree flat detail | match vrrp | match passive
configure groups group <string> router <string> interface <string> ipv4 vrrp <string
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
 | number> passive <boolean>
configure groups group <string> router <string> interface <string> ipv6 vrrp <string | number>
 passive <boolean>
configure groups group <string> service ies <string> interface <string> ipv4 vrrp <string |
 number> passive <boolean>
configure groups group <string> service ies <string> interface <string> ipv6 vrrp <string |
 number> passive <boolean>
configure groups group <string> service vprn <string> interface <string> ipv4 vrrp <string |
 number> passive <boolean>
configure groups group <string> service vprn <string> interface <string> ipv6 vrrp <string |
 number> passive <boolean>
configure router <string> interface <string> ipv4 vrrp <number> passive <boolean>
configure router <string> interface <string> ipv6 vrrp <number> passive <boolean>
configure service ies <string> interface <string> ipv4 vrrp <number> passive <boolean>
configure service ies <string> interface <string> ipv6 vrrp <number> passive <boolean>
configure service vprn <string> interface <string> ipv4 vrrp <number> passive <boolean>
configure service vprn <string> interface <string> ipv6 vrrp <number> passive <boolean>
```

All PEs configured with passive VRRP become VRRP master and take ownership of the virtual IP and MAC addresses. Figure 95: Passive VRRP - vMAC/vIP advertised by GARP shows the use of passive VRRP where the VRID and default gateway (GW) are identical for all nodes, and therefore, the vMAC/vIP are identical. Each PE sends Gratuitous Address Resolution Protocol (GARP) messages with the same vMAC/vIP.

Figure 95: Passive VRRP - vMAC/vIP advertised by GARP

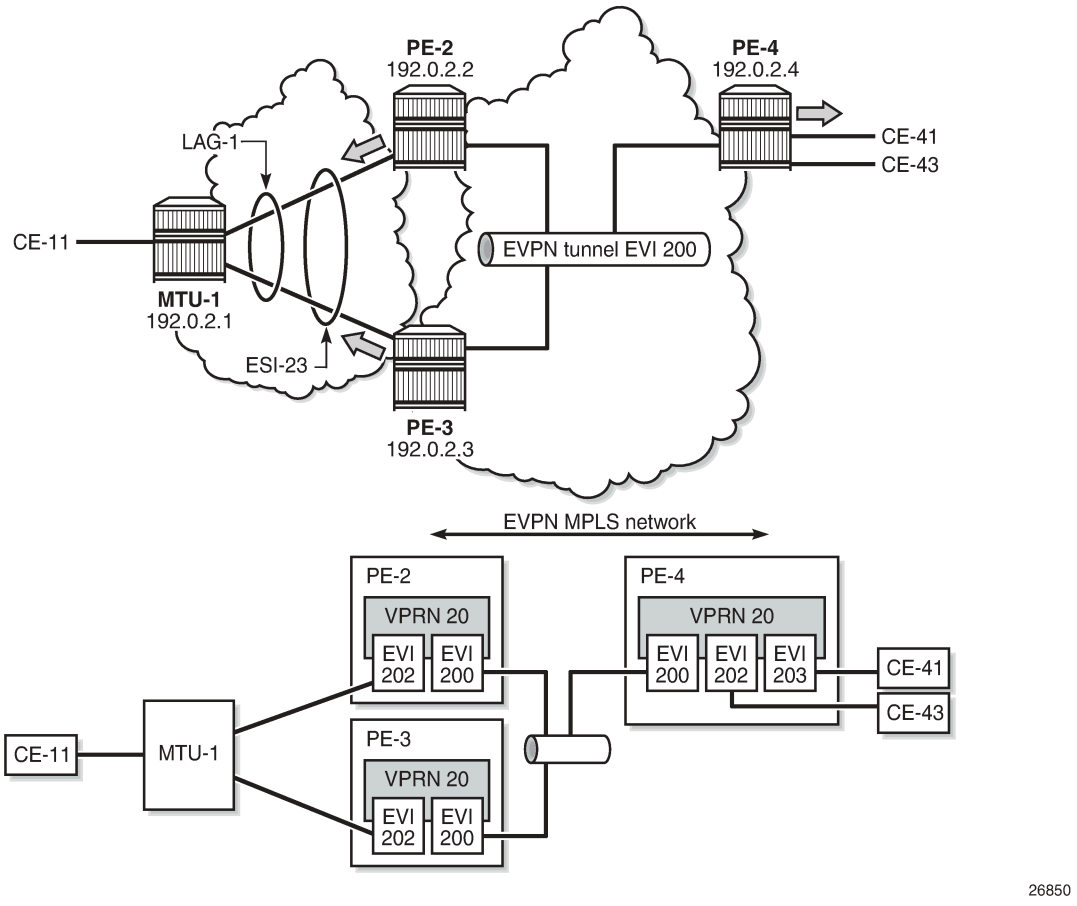

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Ethernet VPN instance (EVI) 202 is configured on all PEs as an R-VPLS with passive VRRP. Each individual R-VPLS interface has a unique MAC/IP, but they all have the same vMAC/vIP because they share the same VRID and backup IP address. The vMAC address is auto-derived out of 00:00:5e:00:00:<VRID>, as per RFC 3768.

The behavior is as follows:

- PEs advertise their real MAC/IP and their vMAC/vIP in EVPN for EVI 202.
- All hosts in EVI 202 have a unique configured default GW.
- When a CE sends upstream traffic to a remote subnet, the packets are routed by the closest PE because the vMAC address is local on each PE.
- In case of ES failure, or in case of single-active MH if the traffic arrives at the non-Designated Forwarder (NDF) PE, the traffic will not be discarded at the peer ES PE. Virtual MAC addresses bypass the R-VPLS interface protection, so traffic can be forwarded between the PEs without being dropped. Note that if passive VRRP was not used in this case and the same regular interface anycast MAC/IP was used instead, the peer PE would discard the traffic due to the MAC Source Address (SA).

Passive VRRP provides an efficient anycast default gateway solution, with the following advantages compared to regular VRRP:

- No need for multiple VRRP instances to achieve default GW load-balancing. Only one VRRP instance is in the R-VPLS, so only one default GW is needed for all hosts.
- Fast convergence because all the nodes in the VRID are master.
- Better scalability because there is no need for keepalive messages or BFD to detect failures.

Passive VRRP provides the following advantages compared to using the same anycast MAC/IP in all the Integrated Routing Bridging (IRB) interfaces:

- VRRP vMAC SA bypasses the protection in the receiving R-VPLS service; therefore, frames with MAC SA matching the local vMAC address are not discarded, and VRRP vMAC SAs can be used in combination with EVPN multi-homing.
- PEs will not show traps claiming duplicate IP addresses.
- vMAC addresses are auto-derived from the VRID, so no need to configure the same MAC address in all the IRB interfaces.
- PEs can still use their real (unique) IRB IP addresses when sending ICMP packets for troubleshooting purposes.

Configuration

In this section, the following use cases are described:

- EVPN-MPLS R-VPLS without multi-homing
- EVPN-MPLS R-VPLS with all-active multi-homing ES
- EVPN-MPLS R-VPLS with single-active multi-homing ES

EVPN-MPLS R-VPLS without multi-homing

The first scenario describes R-VPLS support including IP route advertisement (BGP-EVPN route type 5) with EVPN tunnel interfaces, without multi-homing. VPLS 101 does not have any connected host, but the

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

linked VPRN has SAP 1/2/1:10. Figure 96: R-VPLS with EVPN tunnel, without multi-homing shows the example topology used for R-VPLS with EVPN tunnel but without multi-homing. IP prefixes are advertised.

Figure 96: R-VPLS with EVPN tunnel, without multi-homing

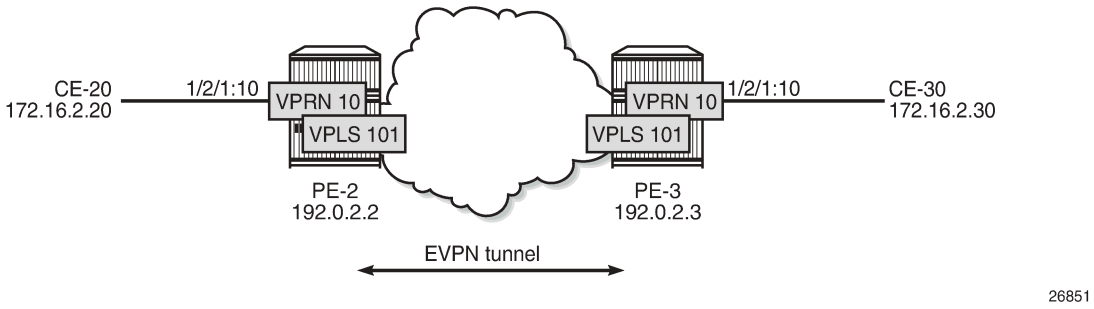

The initial configuration includes the following:

- Cards, MDAs, ports
- Router interface between PE-2 and PE-3
- IS-IS (or OSPF)
- LDP enabled on the router interface between PE-2 and PE-3

BGP is configured for the EVPN address family on PE-2 and PE-3. The BGP configuration on PE-2 is as follows. The BGP configuration on PE-3 is similar.

```
# on PE-2:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            family {
                ipv4 false
                evpn true
            }
            rapid-update {
                evpn true
            }
            group "internal" {
                peer-as 64500
            }
            neighbor "192.0.2.3" {
                group "internal"
            }
        }
```

The CEs are connected to SAP 1/2/1:10 in VPRN 10. R-VPLS 101 is bound to VPRN 10 and VPRN 10 has a dedicated interface "int-evi-101" for the EVPN tunnel. In general, if only one route-target (RT) is used for import and export in the EVPN-VPLS, it is good to add the EVI and have the route distinguisher (RD) and

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

RT auto-derived from the EVI. It is simpler and avoids configuration mistakes. The service configuration on PE-2 is as follows:

```
# on PE-2:
configure {
    service {
        vpls "evi-101" {
            admin-state enable
            service-id 101
            customer "1"
            routed-vpls {
            }
            bgp 1 {              # RD and RT are not manually configured in BGP context
            }
            bgp-evpn {
                evi 101       # RD and RT will be auto-derived from the EVI
                routes {
                    ip-prefix {
                        advertise true
                    }
                }
                mpls 1 {
                    admin-state enable
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
        }
        vprn "VPRN 10" {
            admin-state enable
            service-id 10
            customer "1"
            interface "int-PE-2-CE-20" {
                ipv4 {
                    primary {
                        address 172.16.2.1
                        prefix-length 24
                    }
                }
                sap 1/2/1:10 {
                }
            }
            interface "int-evi-101" {
                vpls "evi-101" {
                    evpn-tunnel {
                    }
                }
            }
        }
```

- The **routed-vpls** command is required so that R-VPLS "evi-101" can be bound to VPRN 10.
- The service name "evi-101" must match the name in the VPRN 10 VPLS interface.
- The VPRN 10 VPLS interface is configured with the keyword **evpn-tunnel**. This configuration has the advantage of not having to allocate IP addresses to the R-VPLS interfaces, however, it cannot be used when the R-VPLS has local SAPs.

The configuration is similar on PE-3. It is important that the RD is different on PE-2 and PE-3, but it is automatically the case when the RD is auto-derived from the configured EVI, as in the example. The RD on PE-2 is 192.0.2.2:101; on PE-3, the RD is 192.0.2.3:101.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

PE-3 receives the following BGP-EVPN IP prefix route for prefix 172.16.2.0/24 from PE-2:

```
2 2022/02/24 11:00:28.145 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 90
    Flag: 0x90 Type: 14 Len: 45 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
        Type: EVPN-IP-PREFIX Len: 34 RD: 192.0.2.2:101, tag: 0,
           ip_prefix: 172.16.2.0/24 gw_ip 0.0.0.0 Label: 8388496 (Raw Label: 0x7fff90)
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:101
        mac-nh:02:13:ff:ff:ff:a2
        bgp-tunnel-encap:MPLS
"
```

GW IP 0.0.0.0 is an indication that an EVPN tunnel is in use. With EVPN tunnels, no IRB IP address needs to be configured in the VPRN. EVPN tunnels make provisioning easier to automate and save IP addresses from the tenant IP space.

The BGP tunnel encapsulation is MPLS, but the MPLS label in the debug message is not the same as in the service, because the router will strip the extra four lowest bits to get the 20-bit MPLS label. In the debug message, the label is 8388496. This is because the debug message is shown before the router can parse the label field and see if it corresponds to an MPLS label (20 bits) or a VXLAN VNI (24 bits). The MPLS label is calculated by dividing the label value by 24 (16), as follows: 8388496/16 = 524281.

The MAC next-hop extended community 02:13:ff:ff:ff:a2 is the MAC address of the interface "int-evi-101" in VPRN 10 on PE-2, as follows:

```
[/]
A:admin@PE-2# show router 10 interface "int-evi-101" detail | match "MAC Address"
MAC Address      : 02:13:ff:ff:ff:a2    Mac Accounting    : Disabled
```

The routing table for VPRN 10 on PE-3 contains the route for prefix 172.16.2.0/24 as the EVPN-IFF (IFF stands for Interface-ful) route with next-hop "int-evi-101" and interface name "ET-02:13:ff:ff:ff:a2" (ET stands for EVPN Tunnel), as follows:

```
[/]
A:admin@PE-3# show router 10 route-table

===============================================================================
Route Table (Service: 10)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
172.16.2.0/24                                 Remote  EVPN-IFF  01h43m58s  169
       int-evi-101 (ET-02:13:ff:ff:ff:a2)                          0
172.16.3.0/24                                 Local   Local     01h43m59s  0
       int-PE-3-CE-30                                             0
-------------------------------------------------------------------------------
No. of Routes: 2
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
       S = Sticky ECMP requested
===============================================================================
```

The forwarding database (FDB) for VPLS 101 on PE-3 shows an entry for MAC address 02:13:ff:ff:ff:a2 that is learned via EVPN. The MAC address is static (S) and protected (P). The MPLS label is 524281.

```
[/]
A:admin@PE-3# show service id 101 fdb detail

===============================================================================
Forwarding Database, Service 101
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
101        02:13:ff:ff:ff:a2 mpls-1:                 EvpnS:P  02/24/22 11:00:35
                             192.0.2.2:524281
           ldp:65538
101        02:17:ff:ff:ff:a2 cpm                     Intf     02/24/22 11:00:34
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

When the CEs have IPv6 addresses, the VPRN configuration is similar on the PEs, but the **ipv6** context must be enabled in the EVPN tunnel interface, so that the router can advertise and process BGP-EVPN routes type 5 with IPv6 prefixes. The configuration of the VPLS is identical for IPv4 and IPv6.

```
# on PE-2:
configure {
    service {
        vpls "evi-106" {
            admin-state enable
            service-id 106
            customer "1"
            routed-vpls {
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 106
                routes {
                    ip-prefix {
                        advertise true
                    }
                }
                mpls 1 {
                    admin-state enable
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
        }
        vprn "VPRN 16" {
            admin-state enable
            service-id 16
            customer "1"
            interface "int-PE-2-CE-26" {
                sap 1/2/1:16 {
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                ipv6 {
                    address 2001:db8:16::2:1 {
                        prefix-length 120
                    }
                }
            }
            interface "int-evi-106" {
                vpls "evi-106" {
                    evpn-tunnel {
                    }
                }
                ipv6 {
                }
            }
        }
```

When advertising IPv6 prefixes, the GW IP field in the route type 5 is always populated with the IPv6 address of the R-VPLS interface. In this example, because no specific IPv6 global address is configured, the GW IP will be populated with the auto-created link local address. The following BGP update is received by PE-3 for IPv6 prefix 2001:db8:16::2:0/120:

```
# on PE-3:
9 2022/02/24 11:00:35.338 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 106
    Flag: 0x90 Type: 14 Len: 69 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
        Type: EVPN-IP-PREFIX Len: 58 RD: 192.0.2.2:106, tag: 0,
              ip_prefix: 2001:db8:16::2:0/120 gw_ip fe80::14:1ff:fe02:1
              Label: 8388480 (Raw Label: 0x7fff80)
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:106
        bgp-tunnel-encap:MPLS
"
```

The IPv6 route-table on PE-3 is as follows:

```
[/]
A:admin@PE-3# show router 16 route-table ipv6

===============================================================================
IPv6 Route Table (Service: 16)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
2001:db8:16::2:0/120                          Remote  EVPN-IFF  01h50m01s  169
       fe80::14:1ff:fe02:1-"int-evi-106"                             0
2001:db8:16::3:0/120                          Local   Local     01h50m01s  0
       int-PE-3-CE-36                                                0
-------------------------------------------------------------------------------
No. of Routes: 2
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
```

EVPN-MPLS R-VPLS with all-active MH

Figure 97: EVPN-MPLS R-VPLS with all-active MH ES shows the example topology with all-active multi-homing ES "AA-ESI-23".

Figure 97: EVPN-MPLS R-VPLS with all-active MH ES

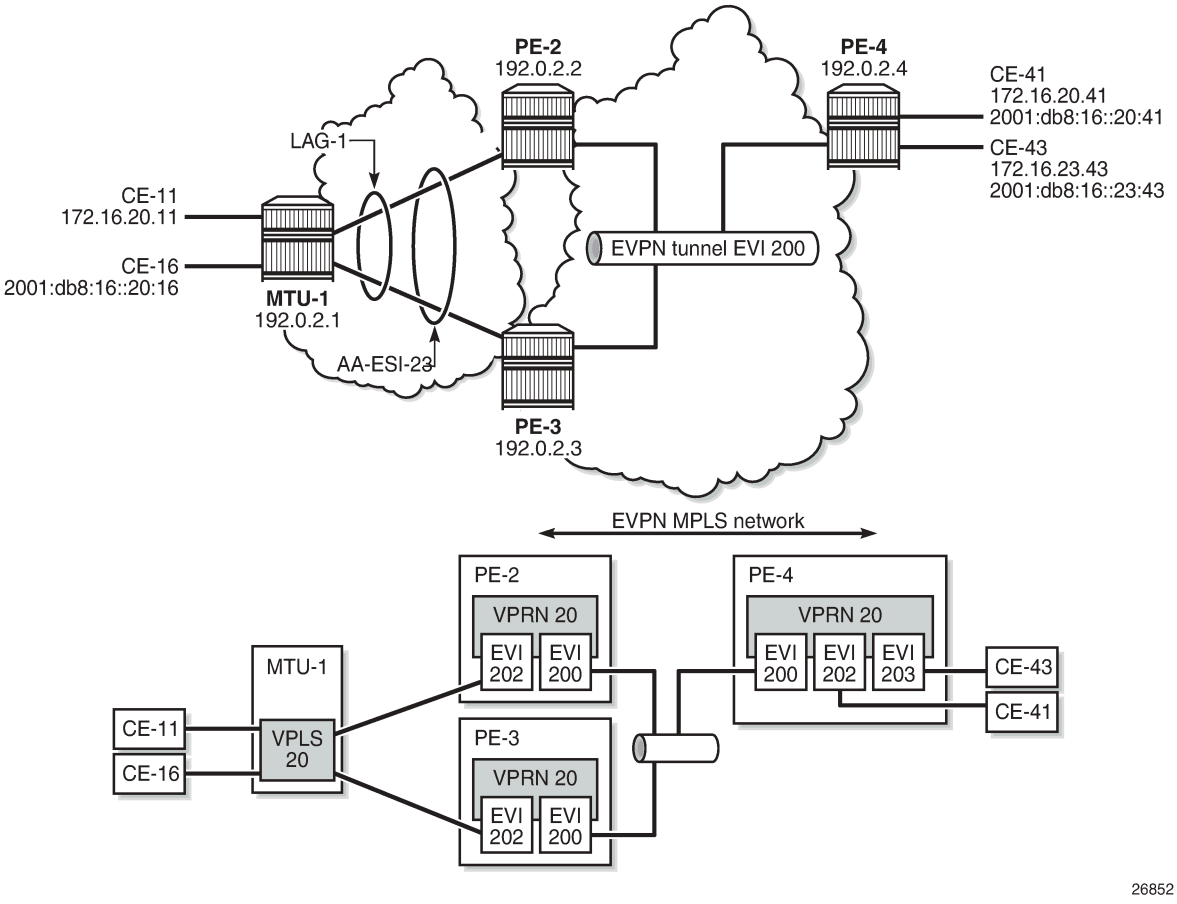

BGP is configured between PE-2, PE-3, and PE-4 for address family EVPN. The configuration on PE-2 is as follows:

```
# on PE-2:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            family {
                ipv4 false
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                evpn true
            }
            rapid-update {
                evpn true
            }
            group "internal" {
                peer-as 64500
            }
            neighbor "192.0.2.3" {
                group "internal"
            }
            neighbor "192.0.2.4" {
                group "internal"
            }
```

All-active multi-homing Ethernet segment "AA-ESI-23" is configured on PE-2 and PE-3, as follows:

```
# on PE-2, PE-3:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "AA-ESI-23" {
                        admin-state enable
                        esi 01:00:00:00:00:23:00:00:00:01
                        multi-homing-mode all-active
                        df-election {
                            es-activation-timer 3
                        }
                        association {
                            lag "lag-1" {
                            }
                        }
                    }
                }
            }
        }
```

The following services are configured on the PEs:

- VPRN 20 has interfaces bound to VPLS 200 and VPLS 202. On PE-4, VPRN 20 also has an interface bound to VPLS 203.
- VPLS 200 is configured as an EVPN tunnel that connects the PEs.
- VPLS 202 and VPLS 203 have attachment circuits to CEs.

The services are configured on PE-2 as follows. The configuration on PE-3 and PE-4 is similar.

```
# on PE-2:
configure {
    service {
        vpls "evi-200" {
            admin-state enable
            service-id 200
            customer "1"
            routed-vpls {
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 200
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                routes {
                    ip-prefix {
                        advertise true
                    }
                }
                mpls 1 {
                    admin-state enable
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
        }
        vpls "evi-202" {
            admin-state enable
            service-id 202
            customer "1"
            routed-vpls {
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 202
                mpls 1 {
                    admin-state enable
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            sap lag-1:20 {
            }
        }
        vprn "VPRN 20" {
            admin-state enable
            service-id 20
            customer "1"
            interface "int-evi-200" {
                vpls "evi-200" {
                    evpn-tunnel {
                    }
                }
                ipv6 {
                }
            }
            interface "int-evi-202" {
                mac 00:ca:fe:00:02:02
                ipv4 {
                    primary {
                        address 172.16.20.2
                        prefix-length 24
                    }
                    vrrp 1 {
                        backup [172.16.20.254]
                        passive true
                        ping-reply true
                        traceroute-reply true
                    }
                }
                vpls "evi-202" {
                }
                ipv6 {
                    link-local-address {
                        address fe80::16:20:2
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    duplicate-address-detection false
                }
                address 2001:db8:16::20:2 {
                    prefix-length 120
                }
                vrrp 1 {
                    backup [fe80::16:20:fe]
                    passive true
                    ping-reply true
                    traceroute-reply true
                }
            }
        }
        ipv6 {
            router-advertisement {
                interface "int-evi-202" {
                    admin-state enable
                    use-virtual-mac true
                }
            }
        }
    }
```

The IPv6 VRRP backup address is in the same subnet as the link local address of the interface "int-evi-202". The option **duplicate-address-detection false** is configured on the link local address to disable Duplicate Address Detection (DAD) and set the IPv6 address as preferred. Also for IPv6, router advertisement must be enabled and configured to use the virtual MAC address.

Passive VRRP

EVI 202 is configured as an R-VPLS with passive VRRP. A passive-VRRP VRID instance suppresses the transmission and reception of keepalive messages. All PEs configured with passive VRRP become VRRP master and take ownership of the virtual IP and MAC address.

Each individual R-VPLS interface has a different MAC/IP on each PE. The MAC/IPs for "int-evi-202" on PE-2 are MAC 00:ca:fe:00:02:02 and IP 172.16.20.2/24 for IPv4 and the same MAC address with IPv6 2001:db8:16::20:2 and fe80::16:20:2. However, the R-VPLS interfaces on all PEs share the same VRID 1 and backup IP address 172.16.20.254, so the same vMAC/vIP 00:00:5e:00:01:01/172.16.20.254 and vMAC/vIP 00:00:5e:00:02:01/ fe80::16:20:fe are advertised by all PEs. PE-2 advertises the following EVPN MAC routes:

```
83 2022/02/24 15:09:15.841 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.4
"Peer 1: 192.0.2.4: UPDATE
Peer 1: 192.0.2.4 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 285
    Flag: 0x90 Type: 14 Len: 240 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
        Type: EVPN-MAC Len: 49 RD: 192.0.2.2:202 ESI: ESI-0, tag: 0, mac len: 48
            mac: 00:00:5e:00:02:01, IP len: 16, IP: fe80::16:20:fe, label1: 8388416
        Type: EVPN-MAC Len: 37 RD: 192.0.2.2:202 ESI: ESI-0, tag: 0, mac len: 48
            mac: 00:00:5e:00:01:01, IP len: 4, IP: 172.16.20.254, label1: 8388416
        Type: EVPN-MAC Len: 49 RD: 192.0.2.2:202 ESI: ESI-0, tag: 0, mac len: 48
            mac: 00:ca:fe:00:02:02, IP len: 16, IP: fe80::16:20:2, label1: 8388416
        Type: EVPN-MAC Len: 49 RD: 192.0.2.2:202 ESI: ESI-0, tag: 0, mac len: 48
            mac: 00:ca:fe:00:02:02, IP len: 16, IP: 2001:db8:16::20:2, label1: 8388416
        Type: EVPN-MAC Len: 37 RD: 192.0.2.2:202 ESI: ESI-0, tag: 0, mac len: 48
            mac: 00:ca:fe:00:02:02, IP len: 4, IP: 172.16.20.2, label1: 8388416
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:202
        bgp-tunnel-encap:MPLS
        mac-mobility:Seq:0/Static
"
```

The three PEs advertise the same (anycast) vMAC/vIP in EVI 202 as protected, but each PE keeps its own MAC entry in the FDB. The following FDB shows that the source identifier for vMAC 00:00:5e:00:01:01 and vMAC 00:00:5e:00:02:01 is the CPM. These two vMAC entries with source identifier CPM are seen on all PEs.

```
[/]
A:admin@PE-2# show service id 202 fdb detail

===============================================================================
Forwarding Database, Service 202
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
202        00:00:01:00:00:11 sap:lag-1:20            L/0      02/24/22 15:09:21
202        00:00:01:00:00:16 sap:lag-1:20            L/0      02/24/22 15:09:22
202        00:00:04:00:00:41 mpls-1:                 Evpn     02/24/22 15:09:14
                             192.0.2.4:524281
           ldp:65539
202        00:00:5e:00:01:01 cpm                     Intf     02/24/22 15:08:50
202        00:00:5e:00:02:01 cpm                     Intf     02/24/22 15:08:50
202        00:ca:fe:00:02:02 cpm                     Intf     02/24/22 15:08:50
202        00:ca:fe:00:02:03 mpls-1:                 EvpnS:P  02/24/22 15:09:03
                             192.0.2.3:524276
           ldp:65538
202        00:ca:fe:00:02:04 mpls-1:                 EvpnS:P  02/24/22 15:09:14
                             192.0.2.4:524281
           ldp:65539
-------------------------------------------------------------------------------
No. of MAC Entries: 8
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

The interface MAC 00:ca:fe:00:02:02 is local, so it also has the CPM as source identifier. MAC 00:ca:fe:00:02:03 is the PE-3's R-VPLS interface MAC and it is learned via EVPN-MPLS (mpls-1) as static (S) and protected (P). MAC address 00:ca:fe:00:02:04 on PE-4 is also static and protected.

PE-4 sends the following IP prefix route (BGP-EVPN route type 5) for prefix 172.16.23.0/24 to the other PEs:

```
37 2022/02/24 15:09:13.665 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 90
    Flag: 0x90 Type: 14 Len: 45 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.4
        Type: EVPN-IP-PREFIX Len: 34 RD: 192.0.2.4:200, tag: 0,
                             ip_prefix: 172.16.23.0/24 gw_ip 0.0.0.0
                             Label: 8388512 (Raw Label: 0x7fffa0)
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:200
        mac-nh:02:1b:ff:00:00:05
        bgp-tunnel-encap:MPLS
"
```

The IP prefixes are advertised with next-hop equal to the EVPN-tunnel GW MAC "int-evi-200", as follows:

```
[/]
A:admin@PE-4# show router 20 interface "int-evi-200" detail | match "MAC Address"
MAC Address      : 02:1b:ff:00:00:05    Mac Accounting    : Disabled
```

The routing table for VPRN 20 on PE-2 contains IP-prefix 172.16.23.0/24 with next-hop 02:1b:ff:00:00:05, as follows:

```
[/]
A:admin@PE-2# show router 20 route-table

===============================================================================
Route Table (Service: 20)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
172.16.20.0/24                                Local   Local     00h17m12s  0
       int-evi-202                                                 0
172.16.23.0/24                                Remote  EVPN-IFF  00h16m48s  169
       int-evi-200 (ET-02:1b:ff:00:00:05)                         0
-------------------------------------------------------------------------------
No. of Routes: 2
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

The following IPv6 routing table for VPRN 20 on PE-2 contains prefix 2001:db8:16::23:0/120, which has also been advertised by PE-4. The next-hop is again "int-evi-200", only this time the link local IPv6 address is displayed (GW IP) instead of the MAC address. The next-hop is the GW IP value in the route type 5, as long as it is non-zero. When the GW IP address is zero, the route type 5 is expected to contain a mac-nh extended community. The MAC encoded in the extended community is used as next-hop in that case.

```
[/]
A:admin@PE-2# show router 20 route-table ipv6

===============================================================================
IPv6 Route Table (Service: 20)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
2001:db8:16::20:0/120                         Local   Local     00h17m10s  0
       int-evi-202                                                 0
2001:db8:16::23:0/120                         Remote  EVPN-IFF  00h16m46s  169
       fe80::a5:9124:c1ed:83ce-"int-evi-200"                       0
-------------------------------------------------------------------------------
No. of Routes: 2
Flags: n = Number of times nexthop is repeated
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

The EVPN tunnel service VPLS 200 has all the MAC addresses of the EVPN interfaces within VPRN 20 as static (S) and protected (P), as follows:

```
[/]
A:admin@PE-2# show service id "evi-200" fdb detail

===============================================================================
Forwarding Database, Service 200
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
200        02:13:ff:00:00:05 cpm                     Intf     02/24/22 15:08:50
200        02:17:ff:00:00:05 mpls-1:                 EvpnS:P  02/24/22 15:09:03
                             192.0.2.3:524277
           ldp:65538
200        02:1b:ff:00:00:05 mpls-1:                 EvpnS:P  02/24/22 15:09:14
                             192.0.2.4:524282
           ldp:65539
-------------------------------------------------------------------------------
No. of MAC Entries: 3
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

The VRRP instance in each PE is master, as follows:

```
[/]
A:admin@PE-2# show router 20 vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                   VR Id Own Adm  State       Base Pri   Msg Int
                                 IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-evi-202                      1     No  Up   Master       100       1
                                 IPv4      Up   n/a         100        No
  Backup Addr: 172.16.20.254
int-evi-202                      1     No  Up   Master       100       1
                                 IPv6      Up   n/a         100        Yes
  Backup Addr: fe80::16:20:fe
-------------------------------------------------------------------------------
Instances : 2
===============================================================================
```

```
[/]
A:admin@PE-3# show router 20 vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                   VR Id Own Adm  State       Base Pri   Msg Int
                                 IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-evi-202                      1     No  Up   Master       100       1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                                IPv4      Up   n/a         100        No
  Backup Addr: 172.16.20.254
int-evi-202                     1     No  Up   Master       100       1
                                IPv6      Up   n/a         100        Yes
  Backup Addr: fe80::16:20:fe
-------------------------------------------------------------------------------
Instances : 2
===============================================================================
```

```
[/]
A:admin@PE-4# show router 20 vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                  VR Id Own Adm  State       Base Pri   Msg Int
                                IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-evi-202                     1     No  Up   Master       100       1
                                IPv4      Up   n/a         100        No
  Backup Addr: 172.16.20.254
int-evi-203                     2     No  Up   Master       100       1
                                IPv4      Up   n/a         100        No
  Backup Addr: 172.16.23.254
int-evi-202                     1     No  Up   Master       100       1
                                IPv6      Up   n/a         100        Yes
  Backup Addr: fe80::16:20:fe
int-evi-203                     2     No  Up   Master       100       1
                                IPv6      Up   n/a         100        Yes
  Backup Addr: fe80::16:23:fe
-------------------------------------------------------------------------------
Instances : 4
===============================================================================
```

Operation

On PE-4, VPRN 20 has one interface bound to VPLS 202 and another interface bound to VPLS 203. CE-41 is attached to VPLS 202, whereas CE-43 is attached to VPLS 203. When ping messages are sent from CE-41 to CE-43, or vice versa, the messages go via VPRN 20, which has routes to both CEs, as follows:

```
[/]
A:admin@PE-4# show router 20 route-table

===============================================================================
Route Table (Service: 20)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
172.16.20.0/24                                Local   Local     00h19m37s  0
       int-evi-202                                                 0
172.16.23.0/24                                Local   Local     00h19m37s  0
       int-evi-203                                                 0
-------------------------------------------------------------------------------
No. of Routes: 2
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
```

```
[/]
A:admin@PE-4# show router 20 route-table ipv6

===============================================================================
IPv6 Route Table (Service: 20)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
2001:db8:16::20:0/120                         Local   Local     00h19m36s  0
       int-evi-202                                                 0
2001:db8:16::23:0/120                         Local   Local     00h19m36s  0
       int-evi-203                                                 0
-------------------------------------------------------------------------------
No. of Routes: 2
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

When traffic is sent between CE-11 and CE-41, which are both associated with VPLS 202, the forwarding is done by the VPLS and not via the VPRN. The FDB for VPLS 202 on PE-3 is as follows:

```
[/]
A:admin@PE-3# show service id 202 fdb detail

===============================================================================
Forwarding Database, Service 202
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
202        00:00:01:00:00:11 sap:lag-1:20            L/0      02/24/22 15:28:41
202        00:00:01:00:00:16 sap:lag-1:20            L/0      02/24/22 15:28:45
202        00:00:04:00:00:41 mpls-1:                 Evpn     02/24/22 15:28:40
                             192.0.2.4:524281
           ldp:65539
202        00:00:5e:00:01:01 cpm                     Intf     02/24/22 15:09:03
202        00:00:5e:00:02:01 cpm                     Intf     02/24/22 15:09:03
202        00:ca:fe:00:02:02 mpls-1:                 EvpnS:P  02/24/22 15:09:04
                             192.0.2.2:524276
           ldp:65538
202        00:ca:fe:00:02:03 cpm                     Intf     02/24/22 15:09:03
202        00:ca:fe:00:02:04 mpls-1:                 EvpnS:P  02/24/22 15:09:14
                             192.0.2.4:524281
           ldp:65539
-------------------------------------------------------------------------------
No. of MAC Entries: 8
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

MAC 00:00:01:00:00:11 corresponds to CE-11 and is learned on SAP lag-1:20 on PE-3 and advertised via an EVPN MAC route to the BGP peers. MAC 00:00:04:00:00:41 corresponds to CE-41 and was advertised via an EVPN MAC route from PE-4, where the MAC was learned on SAP 1/2/1:41 of VPLS 202, as shown in the following FDB:

```
[/]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
A:admin@PE-4# show service id 202 fdb detail

===============================================================================
Forwarding Database, Service 202
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
202        00:00:01:00:00:11 eES:                    Evpn     02/24/22 15:28:41
                             01:00:00:00:00:23:00:00:00:01
202        00:00:01:00:00:16 eES:                    Evpn     02/24/22 15:28:45
                             01:00:00:00:00:23:00:00:00:01
202        00:00:04:00:00:41 sap:1/2/1:41            L/90     02/24/22 15:28:40
202        00:00:5e:00:01:01 cpm                     Intf     02/24/22 15:09:14
202        00:00:5e:00:02:01 cpm                     Intf     02/24/22 15:09:14
202        00:ca:fe:00:02:02 mpls-1:                 EvpnS:P  02/24/22 15:09:16
                             192.0.2.2:524276
           ldp:65538
202        00:ca:fe:00:02:03 mpls-1:                 EvpnS:P  02/24/22 15:09:16
                             192.0.2.3:524276
           ldp:65539
202        00:ca:fe:00:02:04 cpm                     Intf     02/24/22 15:09:14
-------------------------------------------------------------------------------
No. of MAC Entries: 8
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

CE-43's MAC address is not present in VPLS 202's FDB. VPLS 203's FDB shows the CE-43's MAC address, but not CE-41's. Traffic between these two VPLS services goes via the VPRN and cannot use Layer 2 forwarding.

```
[/]
A:admin@PE-4# show service id 203 fdb detail

===============================================================================
Forwarding Database, Service 203
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
203        00:00:04:00:00:43 sap:1/2/1:43            L/90     02/24/22 15:28:40
203        00:00:5e:00:01:02 cpm                     Intf     02/24/22 15:09:14
203        00:00:5e:00:02:02 cpm                     Intf     02/24/22 15:09:14
203        00:ca:fe:00:23:04 cpm                     Intf     02/24/22 15:09:14
-------------------------------------------------------------------------------
No. of MAC Entries: 4
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

EVPN-MPLS R-VPLS with single-active MH

Figure 98: EVPN-MPLS R-VPLS with single-active multi-homing shows the example topology with single-active multi-homing ES "SA-ESI-23". The difference is that the ES is single-active and SDPs are used instead of a LAG.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 98: EVPN-MPLS R-VPLS with single-active multi-homing

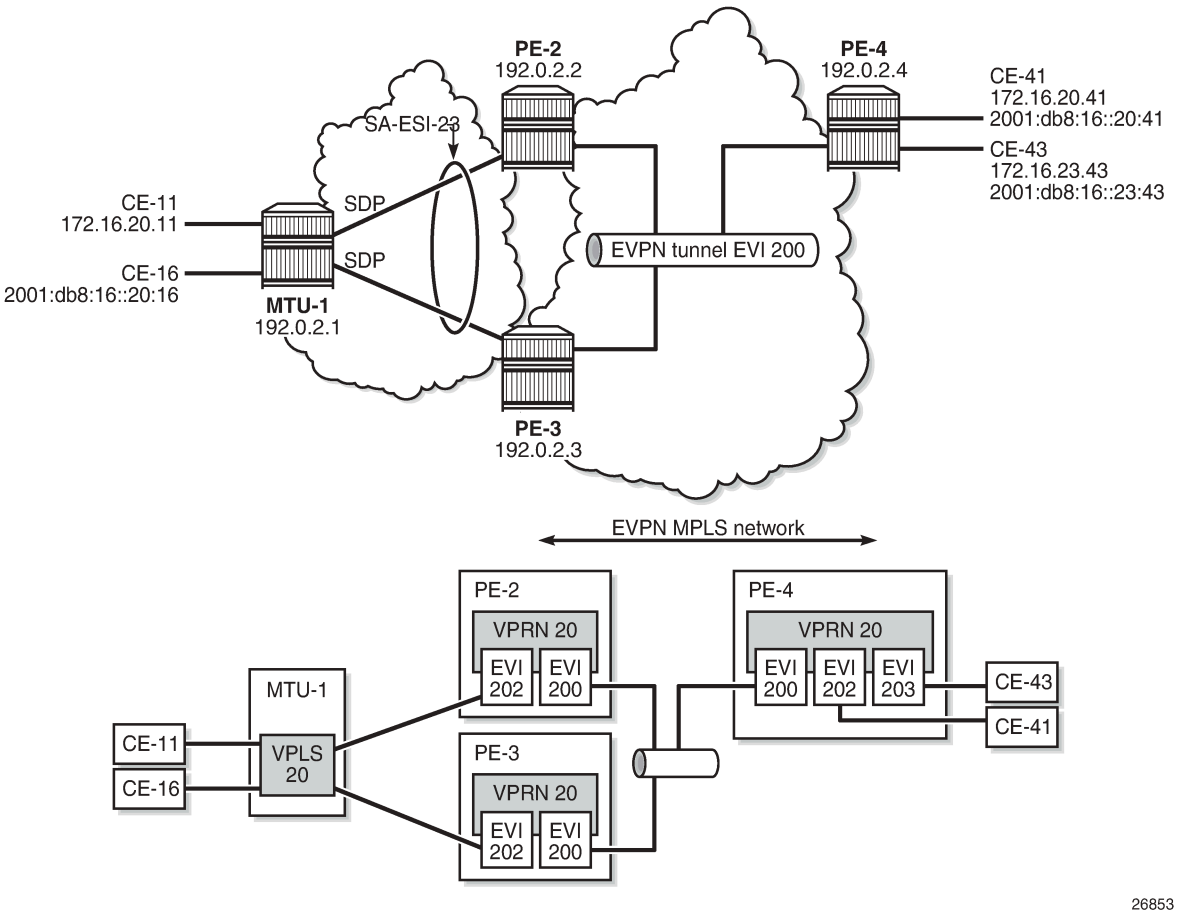

The configuration is modified as follows:

- LAG 1 is removed from MTU-1, PE-2, and PE-3.
- Network interfaces are configured between MTU-1 and PE-2/PE-3 with IS-IS and LDP enabled.
- SDPs are configured.
- Ethernet segment "SA-ESI-23" is configured as single-active multi-homing. The SDP is associated with this ES.
- VPLS 202 on PE-2 and PE-3 no longer has a SAP, but a spoke-SDP instead.
- No changes are required on VPRN 20 or VPLS 200.

The service configuration on PE-2 is as follows. The configuration on PE-3 is similar. No changes are required on PE-4.

```
# on PE-2:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "SA-ESI-23" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        admin-state enable
                        esi 01:00:00:00:00:23:00:00:00:02
                        multi-homing-mode single-active
                        df-election {
                            es-activation-timer 3
                        }
                        association {
                            sdp 21 {
                            }
                        }
                    }
                }
            }
        }
---snip---
        sdp 21 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.1
            }
        }
---snip---
        vprn "VPRN 20" {
            admin-state enable
            service-id 20
            customer "1"
            interface "int-evi-200" {
                vpls "evi-200" {
                    evpn-tunnel {
                    }
                }
                ipv6 {
                }
            }
            interface "int-evi-202" {
                mac 00:ca:fe:00:02:02
                ipv4 {
                    primary {
                        address 172.16.20.2
                        prefix-length 24
                    }
                    vrrp 1 {
                        backup [172.16.20.254]
                        passive true
                        ping-reply true
                        traceroute-reply true
                    }
                }
                vpls "evi-202" {
                }
                ipv6 {
                    link-local-address {
                        address fe80::16:20:2
                        duplicate-address-detection false
                    }
                    address 2001:db8:16::20:2 {
                        prefix-length 120
                    }
                    vrrp 1 {
                        backup [fe80::16:20:fe]
                        passive true
                        ping-reply true
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        traceroute-reply true
                    }
                }
            }
            ipv6 {
                router-advertisement {
                    interface "int-evi-202" {
                        admin-state enable
                        use-virtual-mac true
                    }
                }
            }
        }
        vpls "evi-200" {
            admin-state enable
            service-id 200
            customer "1"
            routed-vpls {
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 200
                routes {
                    ip-prefix {
                        advertise true
                    }
                }
                mpls 1 {
                    admin-state enable
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
        }
        vpls "evi-202" {
            admin-state enable
            service-id 202
            customer "1"
            routed-vpls {
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 202
                mpls 1 {
                    admin-state enable
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            spoke-sdp 21:20 {
            }
        }
```

PE-2 is the Designated Forwarder (DF) in the single-active ES, as shown in the following output:

```
[/]
A:admin@PE-2# show service id 202 ethernet-segment
No sap entries
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
SDP Ethernet-Segment Information
===============================================================================
SDP                  Eth-Seg                          Status
-------------------------------------------------------------------------------
21:20                SA-ESI-23                        DF
===============================================================================
No vxlan instance entries
```

```
[/]
A:admin@PE-3# show service id 202 ethernet-segment
No sap entries

===============================================================================
SDP Ethernet-Segment Information
===============================================================================
SDP                  Eth-Seg                          Status
-------------------------------------------------------------------------------
31:20                SA-ESI-23                        NDF
===============================================================================
No vxlan instance entries
```

When traffic has been sent between CE-11 and CE-41, the FDB on PE-2 is as follows. MAC address 00:00:01:00:00:11 corresponds to CE-11 and has been learned on spoke-SDP 21:20; MAC address 00:00:04:00:00:41 corresponds to CE-41 and has been advertised by PE-4 in an EVPN-MAC route.

```
[/]
A:admin@PE-2# show service id 202 fdb detail

===============================================================================
Forwarding Database, Service 202
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
202        00:00:01:00:00:11 sdp:21:20               L/30     02/24/22 15:36:52
202        00:00:01:00:00:16 sdp:21:20               L/30     02/24/22 15:37:00
202        00:00:04:00:00:41 mpls-1:                 Evpn     02/24/22 15:36:56
                             192.0.2.4:524281
           ldp:65539
202        00:00:5e:00:01:01 cpm                     Intf     02/24/22 15:08:50
202        00:00:5e:00:02:01 cpm                     Intf     02/24/22 15:08:50
202        00:ca:fe:00:02:02 cpm                     Intf     02/24/22 15:08:50
202        00:ca:fe:00:02:03 mpls-1:                 EvpnS:P  02/24/22 15:09:03
                             192.0.2.3:524276
           ldp:65538
202        00:ca:fe:00:02:04 mpls-1:                 EvpnS:P  02/24/22 15:09:14
                             192.0.2.4:524281
           ldp:65539
-------------------------------------------------------------------------------
No. of MAC Entries: 8
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

When the SDP between MTU-1 and DF PE-2 goes down, traffic from CE-41 to CE-11 is forwarded by PE-4 to DF PE-2. PE-2 cannot forward the packets to CE-11 directly, and will forward the packets to its ES peer PE-3. PE-3 will forward to CE-11 even if the MAC SA matches its own vMAC. Virtual MACs bypass the R-VPLS interface protection, so traffic can be forwarded between the PEs without being dropped.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Conclusion

EVPN can be used as the unified control plane VPN technology, not only for providing Layer 2 connectivity, but also Layer 3 (inter-subnet forwarding). EVPN for MPLS tunnels, along with multi-homing and passive VRRP, provides efficient layer-2/layer-3 connectivity to distributed hosts and routers.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

EVPN for PBB over MPLS (PBB-EVPN)

This chapter provides information about EVPN for PBB over MPLS (PBB-EVPN).

Topics in this chapter include:

Applicability

This chapter was initially written for SR OS Release 13.0.R6. The MD-CLI in the current edition is based on SR OS Release 21.2.R1.

Note:
A prerequisite is to read the EVPN for MPLS Tunnels chapter.

Overview

EVPN for Provider Backbone Bridging (PBB) over MPLS (hereafter called PBB-EVPN) is specified in RFC 7623, *Provider Backbone Bridging Combined with Ethernet VPN (PBB-EVPN)*. It provides a simplified version of EVPN-MPLS for cases where the network requires very high scalability and does not need all the advanced features supported by EVPN-MPLS (but still requires single-active and all-active multi-homing capabilities). Table 6: EVPN and PBB-EVPN SR OS feature comparison provides a comparison between the capabilities of EVPN and PBB-EVPN in SR OS, and may help to choose between them when designing a VPN service.

Table 6: EVPN and PBB-EVPN SR OS feature comparison

VPN requirements	EVPN	PBB-EVPN	Comments
All-active Multi-Homing (MH) (flow-based load-balancing)	Yes	Yes	Allows better bandwidth utilization
Single-active MH (service-based load-balancing)	Yes	Yes	
Ethernet Local Area Network (E-LAN) and point-to-point E-Line services	Yes	Yes	
Inter-subnet-forwarding	Yes	No	Allows combined Layer 2 / Layer 3 services. EVPN

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

VPN requirements	EVPN	PBB-EVPN	Comments
Proxy-Address Resolution Protocol / Neighbor Discovery (Proxy-ARP/ND) and IP-duplication protection	Yes	No	Allows Broadcast, Unknown unicast and Multicast (BUM) traffic reduction and better security
Customer MAC (CMAC) protection	Yes	No	Allows protecting key static CMACs
Data Center integration	Yes	No	Integration with VXLAN and Nuage Virtualized Services Directory (VSD)
Control plane overhead	Medium	Low	PBB-EVPN only advertises Backbone MACs (BMACs) and no route type 1s
Confinement of CMAC learning	No	Yes	CMACs are only learned on PEs with flows using those CMACs
CMAC summarization	No	Yes	Aggregation of CMACs into BMACs

PBB-EVPN is a combination of 802.1ah PBB and RFC 7432, *BGP MPLS-Based Ethernet VPN* (EVPN-MPLS), and reuses the PBB-Virtual Private LAN Service (VPLS) service model, where Border Gateway Protocol BGP-EVPN is enabled in the backbone VPLS (B-VPLS) domain. EVPN is used as the control plane in the B-VPLS domain to control the distribution of BMACs and set up per-backbone service instance identifier (ISID) flooding trees for service instance VPLS (I-VPLS) services. The learning of the CMACs, either on local SAPs/SDP-bindings or associated with remote BMACs, is still performed in the data plane. Only the learning of BMACs in the B-VPLS is performed through BGP.

The SR OS PBB-EVPN implementation supports I-VPLS and PBB-Epipe services, including single-active and all-active multi-homing.

Because PBB-EVPN is based on the same control plane model as EVPN for MPLS, it is recommended to read the EVPN for MPLS Tunnels chapter before configuring PBB-EVPN. PBB-EVPN uses a subset of the BGP-EVPN routes described in EVPN for MPLS Tunnels as shown in Figure 99: EVPN route types.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 99: EVPN route types

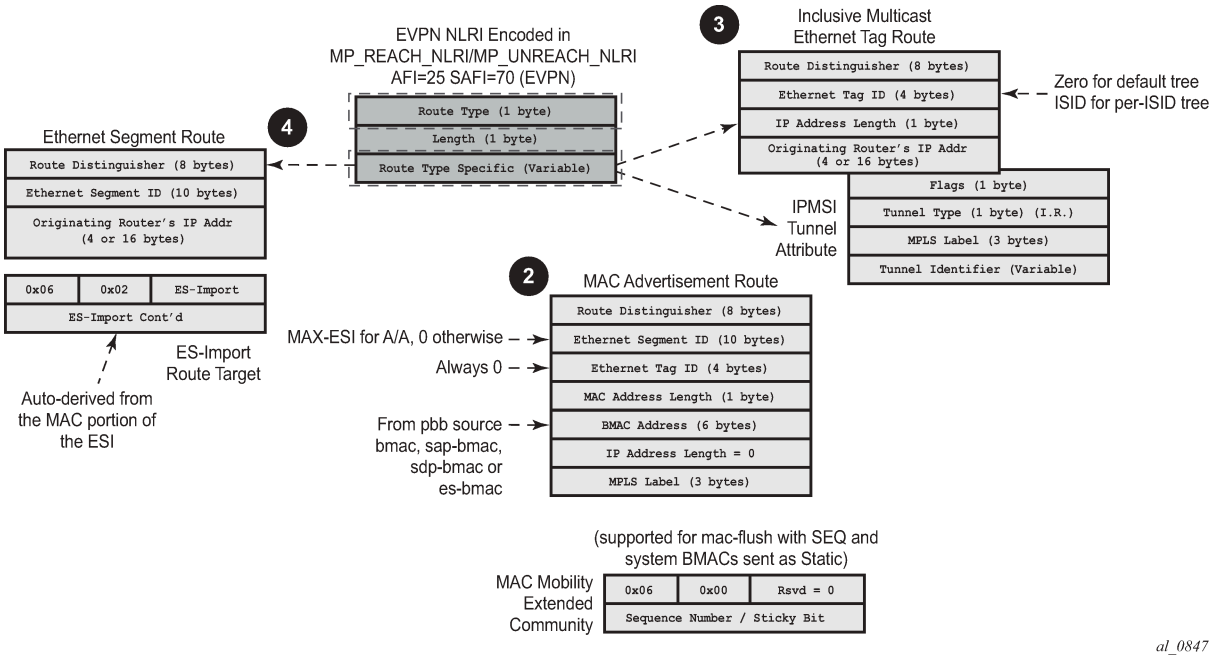

When no EVPN multi-homing is used in the network, only the base routes are used. Route types 2 and 3 are considered the base and mandatory routes:

- Route type 2 — (B) MAC route — In PBB-EVPN, this route type is used for the advertisement of BMAC addresses that will be installed in the remote Forwarding Data Bases (FDBs). There are no IP addresses advertised in PBB-EVPN. The MAC mobility extended community is used for advertising system BMACs as **protected** (with the sticky bit set) and it is also used for CMAC flush in some single-homing scenarios that will be described later.
- Route type 3 — Inclusive Multicast route — This route type is used for the advertisement of the I-VPLS ISIDs (no Epipes) and the desired multicast tree for each of them. The ISIDs are encoded in the Ethernet-tag field of the Network Layer Reachability Information (NLRI). When the B-VPLS is created and enabled, an Inclusive Multicast route with ISID = 0 is advertised. This is for the creation of the default multicast tree.

When EVPN multi-homing is used in an ISID, route type 4 (Ethernet Segment (ES) route) is used. In PBB-EVPN, there is no route type 1 advertised when multi-homing is used on the ISID services (I-VPLS and Epipes). Only route type 4 is used, and in the same way as it is for EVPN-MPLS. See the EVPN for MPLS Tunnels example for more information about ES routes, how they are formed, and how their RT/RD values are populated.

Configuration

This example describes the basic PBB-EVPN configuration first (without multi-homing) and how the flood containment is handled in PBB-EVPN. Flood containment refers to the efficient distribution of the BUM traffic generated for an ISID.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Networks are not always greenfield, so a smooth migration of PBB-EVPN from PBB-VPLS is required to minimize the effect on existing services. This example also describes this migration, starting from a common PBB-VPLS configuration.

Finally, this example describes the configuration of PBB-EVPN multi-homing.

The same setup described in the VPN for MPLS tunnels example is used:

- Four PEs in the core (PE-2, PE-3, PE-4, and PE-5).
- The PEs are interconnected in the same way as explained in EVPN for MPLS Tunnels with the same IP addressing, IS-IS, transport LDP, and BGP peering configuration. There is not any difference with the basic infrastructure. See the EVPN for MPLS Tunnels chapter if more information is required.
- When configuring multi-homing, MTU-1 and MTU-6 are connected to the core.

PBB-EVPN configuration without multi-homing

Figure 100: PBB-EVPN network without multi-homing shows the example topology used in this chapter.

Figure 100: PBB-EVPN network without multi-homing

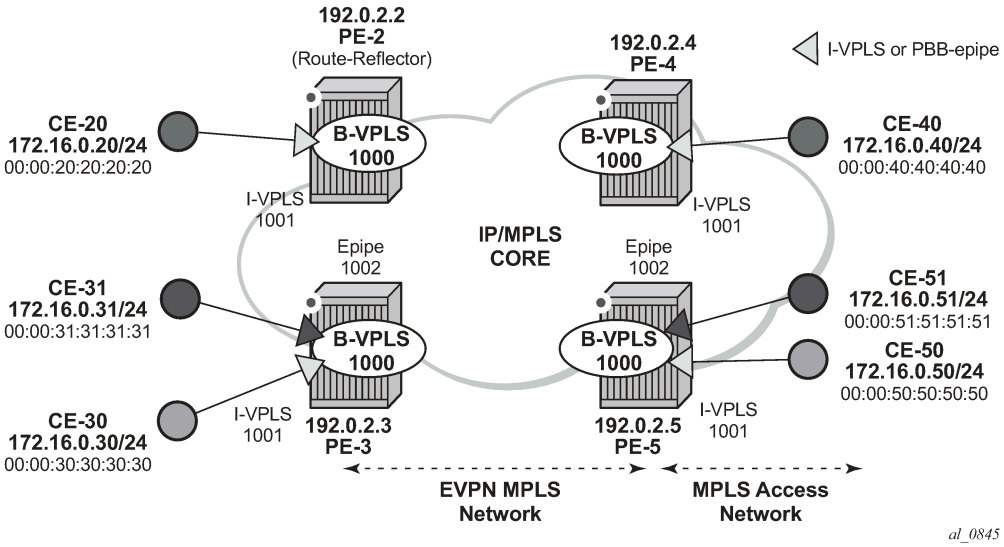

When configuring PBB-EVPN:

- There is no difference at the access side (I-VPLS and Epipe configuration) compared to other PBB technologies supported in SR OS, such as Shortest Path Bridging for MAC (SPBM) or PBB-VPLS.
- The B-VPLS becomes an EVPN-MPLS service, where bgp-evpn mpls is added.

The following output shows an example of a basic configuration in PE-3. B-VPLS 1000 is bgp-evpn enabled and I-VPLS 1001 and Epipe 1002 are linked to B-VPLS 1000.

```
# on PE-3:
configure {
    service {
        vpls "B-VPLS 1000" {
            admin-state enable
            service-id 1000
            customer "1"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            service-mtu 2000
            pbb-type b-vpls
            pbb {
                source-bmac {
                    address 00:00:00:00:00:03
                }
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 1000
                mpls 1 {
                    admin-state enable
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
        }
        vpls "I-VPLS 1001" {
            admin-state enable
            service-id 1001
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS 1000" {
                    isid 1001
                }
            }
            sap 1/2/1:1001 {
            }
        }
        epipe "Epipe 1002" {
            admin-state enable
            service-id 1002
            customer "1"
            pbb {
                tunnel {
                    backbone-vpls-service-name "B-VPLS 1000"
                    isid 1002
                    backbone-dest-mac 00:00:00:00:00:05
                }
            }
            sap 1/2/1:1002 {
            }
        }
```

In the preceding output, there is no new configuration needed for I-VPLS/Epipe services. As for the B-VPLS, the output shows the minimum configuration required. If needed, the following parameters can be modified in the **bgp-evpn** context:

```
[ex:/configure service vpls "B-VPLS 1000"]
A:admin@PE-2# bgp-evpn ?

 bgp-evpn

 accept-ivpls-evpn-    - Accept non-zero ethernet-tag MAC routes and process for CMAC
  flush                  flushing
 apply-groups          - Apply a configuration group at this level
 apply-groups-exclude  - Exclude a configuration group at this level
 evi                   - EVPN ID
 incl-mcast-orig-ip    - Originating IP address
 isid-route-target     + Enter the isid-route-target context
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
 mac-duplication       + Enter the mac-duplication context
 mpls                  + Enter the mpls list instance
 routes                + Enter the routes context
 vxlan                 + Enter the vxlan list instance
```

The following parameters can be modified in the **bgp-evpn mpls 1** context:

```
[ex:/configure service vpls "B-VPLS 1000" bgp-evpn]
A:admin@PE-2# mpls 1 ?

 mpls

 admin-state           - Administrative state of BGP EVPN MPLS
 apply-groups          - Apply a configuration group at this level
 apply-groups-exclude  - Exclude a configuration group at this level
 auto-bind-tunnel      + Enter the auto-bind-tunnel context
 control-word          - Enable/disable setting the CW bit in the label message.
 default-route-tag     - Default route tag
 ecmp                  - Maximum ECMP routes information
 entropy-label         - Enable/disable use of entropy-label.
 fdb                   + Enter the fdb context
 force-vc-forwarding   - VC forwarding action
 ingress-replication-  - Use the same label as the one advertised for unicast traffic
  bum-label
 oper-group            - Operational-Group identifier.
 route-next-hop        + Enter the route-next-hop context
 send-tunnel-encap     + Enter the send-tunnel-encap context
 split-horizon-group   - Split horizon group
```

A detailed description of these commands is included in the EVPN for MPLS Tunnels chapter. In addition to the preceding commands, the following **service>(b-)vpls>pbb** commands are relevant for PBB-EVPN in the B-VPLS service:

- **force-qtag-forwarding** allows the transparent transport of the customer 802.1p bits across the B-VPLS services.
- **source-bmac>address** can modify the source BMAC for all the PBB packets containing traffic from non-multi-homed I-VPLS and Epipe services.
- **source-bmac>use-es-bmac-lsb true** instructs the system to use an ES-specific BMAC for traffic coming from an ES on an I-VPLS or Epipe.
- **source-bmac>use-mclag-bmac-lsb true** instructs the system to use a SAP-specific BMAC for traffic coming from an MC-LAG I-VPLS/Epipe SAP.

Flood containment for I-VPLS services

In general, PBB technologies in SR OS support a way to contain flooding for a specified I-VPLS ISID, so that BUM traffic for that ISID only reaches the PEs where the ISID is locally defined. Each PE creates a Multicast Forwarding Information Base (MFIB) per I-VPLS ISID on the B-VPLS instance. That MFIB supports SAP/SDP-binding endpoints that can be populated by:

- Multiple MAC Registration Protocol (MMRP) in regular PBB-VPLS
- IS-IS in SPBM

In PBB-EVPN, B-VPLS EVPN destinations can be added to the MFIBs using EVPN Inclusive Multicast Ethernet tag routes when they include the ISID in the Ethernet-tag. By default, when a B-VPLS is successfully enabled (**admin-state enable**), the PE advertises:

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- An Inclusive Multicast route for ISID = 0 — This allows the remote PEs to add the advertising PE to the default-multicast-list for the B-VPLS.
- An Inclusive Multicast route for each local ISID defined in the system (a local ISID includes configured I-VPLS and static-ISIDs) — This allows the remote PEs to create MFIB entries in the B-VPLS for the received ISIDs.

Because EVPN destinations, B-SAPs, and B-spoke-SDPs can coexist in the same B-VPLS, be aware of the different flooding lists created and how they are used in a B-VPLS. Figure 101: PBB-EVPN — flooding lists illustrates this concept with an example for B-VPLS 1000 in PE-1. The assumptions are:

- I-VPLS 1001 is created in PE-1, PE-2, and PE-4 only.
- PE-1, PE-2, PE-3, PE-4, and PE-5 support BGP-EVPN in B-VPLS 1000.
- PE-6 and PE-7 only support spoke-SDPs.
- PE-1 is connected to all six PEs.

Figure 101: PBB-EVPN — flooding lists

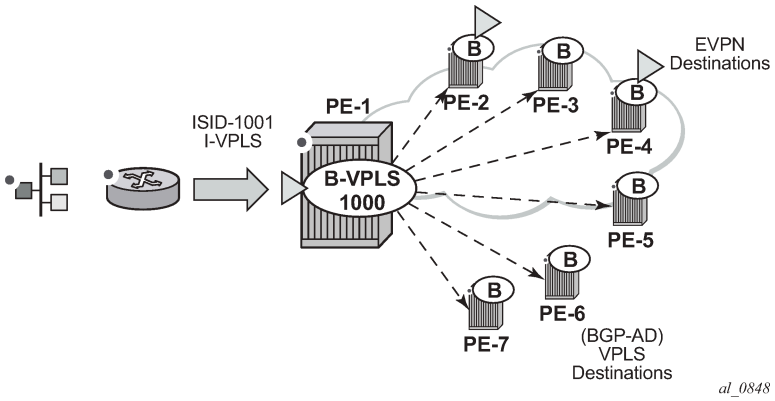

In this situation, PE-1 creates two flooding lists in B-VPLS 1000:

- Default-multicast-list — composed of:
  - All the EVPN PEs that advertised ISID = 0 (PE-2, PE-3, PE-4, PE-5).
  - All the B-spoke-SDPs (or B-SAPs) (PE-6, PE-7).
  - All the EVPN PEs that advertised ISID 1001 and no ISID 0 (if an isid-policy is created in PE-1 stating **use-def-mcast** for ISID 1001). Note: third-party PEs may not advertise ISID = 0, but only non-zero ISIDs.
- MFIB for ISID 1001 is composed of:
  - All the EVPN PEs that advertised ISID 1001 (PE-2 and PE-4) unless there is an ISID-policy in PE-1 stating **use-def-mcast** for ISID 1001.
  - Static-ISIDs defined in manual B-spoke-SDPs and B-SAPs (static-ISIDs cannot be created on BGP-AD auto-discovered B-spoke-SDPs).

Based on the above, when BUM traffic is sent to I-VPLS 1001 on PE-1:

- The traffic is encapsulated in PBB with the group BMAC for ISID 1001 and sent (by default) to the MFIB created for ISID 1001 (PE-2 and PE-4).
- If an ISID-policy is added with **use-def-mcast** for ISID 1001, the BUM traffic is encapsulated in PBB with the group BMAC for ISID 1001 and sent to the default-multicast-list, that is, all six remote PEs.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Referring to Figure 100: PBB-EVPN network without multi-homing, the following output illustrates the use of the ISID-policy in PBB-EVPN. PE-2 does not have any ISID-policy configured; when it receives BUM traffic from the local I-VPLS 1001, it uses the MFIB for ISID 1001:

```
# on PE-2:
configure {
    service {
        vpls "B-VPLS 1000" {
            admin-state enable
            service-id 1000
            customer "1"
            service-mtu 2000
            pbb-type b-vpls
            pbb {
                source-bmac {
                    address 00:00:00:00:00:02
                }
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 1000
                mpls 1 {
                    admin-state enable
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
        }
```

```
[/]A:admin@PE-2# show service id 1000 mfib

===============================================================================
Multicast FIB, Service 1000
===============================================================================
Source Address  Group Address         Port Id                     Svc Id   Fwd
                                                                         Blk
-------------------------------------------------------------------------------
*               01:1e:83:00:03:e9     b-mpls:192.0.2.3:524282     Local    Fwd
                                      b-mpls:192.0.2.4:524282     Local    Fwd
                                      b-mpls:192.0.2.5:524282     Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
```

An ISID-policy can be added to modify this behavior and allow PE-2 to use the default multicast list. If I-VPLS 1001 exists in all the remote PEs (as in this example), using the default multicast list is as efficient as using the MFIB and saves expensive MFIB resources. In the following output, as soon as the ISID-policy is added, the MFIB entries for ISID 1001 are removed and PE-2 starts using the default multicast list.

```
# on PE-2:
configure {
    service {
        vpls "B-VPLS 1000" {
            isid-policy {
                entry 10 {
                    use-def-mcast true
                    range {
                        start 1001
                        end 2000
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    }
                }
```

The MFIB on PE-2 does not contain any entries for ISID 1001 anymore, as follows:

```
[/]
A:admin@PE-2# show service id 1000 mfib

================================================================================
Multicast FIB, Service 1000
================================================================================
Source Address  Group Address         Port Id                     Svc Id   Fwd
                                                                         Blk
--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
Number of entries: 0
================================================================================
```

PBB-VPLS to PBB-EVPN migration

The principles required for migrating a PBB-VPLS network to PBB-EVPN are explained in the VPLS to EVPN-MPLS integration section of the EVPN for MPLS Tunnels chapter. Those principles are also applicable to EVPN destinations and spoke-SDPs in the B-VPLS and can be summarized in three points:

- Systems with an EVPN destination and SDP-binding to the same far-end IP bring down the SDP-binding. This avoids loops when both constructs exist in the same network.
- SDP-bindings and EVPN destinations can be placed in the same Split-Horizon Group (SHG). When traffic from an SDP-binding/EVPN destination belonging to that SHG is received on a PE, it is never forwarded to another SDP-binding or EVPN destination on the same SHG.
- MAC addresses learned on an SDP-binding or SAP, that belong to an SHG where EVPN destinations are also created, are not advertised in BGP-EVPN.

Based on those principles, this section describes how to migrate a PBB-VPLS network to PBB-EVPN. The network in Figure 100: PBB-EVPN network without multi-homing represents a regular PBB-VPLS network that needs to be migrated to PBB-EVPN.

In that network, the four PEs are running BGP-AD and TLDP for the discovery and setup of the pseudowires in the B-VPLS instance. The advantage of this configuration is that the migration can be done node by node and with minimum impact on customer service.

Initial configuration

Initially, the network is configured for PBB-VPLS with BGP-AD in B-VPLS 1000. The EVPN family is to be added. At the access, I-VPLS 1001 is connected to the CEs. As an example, the configuration in PE-3 is shown. An equivalent configuration exists in the other three PEs.

Note:
The EVPN family is added to the BGP configuration because PBB-EVPN uses this address family. Assuming there are redundant Route Reflectors (RRs), the addition of EVPN can be done without service impact. In this example, the assumption is that the PEs are already configured with the EVPN family.

```
# on PE-3:
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            split-horizon true
            rapid-update {
                evpn true
            }
            group "internal" {
                peer-as 64500
                family {
                    l2-vpn true
                    evpn true
                }
            }
            neighbor "192.0.2.2" {
                group "internal"
            }
        }
```

```
# on PE-3:
configure {
    service {
        pw-template "PW1" {
            pw-template-id 1
            split-horizon-group {
                name "CORE"
            }
        }
        vpls "B-VPLS 1000" {
            admin-state enable
            service-id 1000
            customer "1"
            service-mtu 2000
            pbb-type b-vpls
            pbb {
                source-bmac {
                    address 00:00:00:00:00:03
                }
            }
            bgp 1 {
                pw-template-binding "PW1" {
                }
            }
            bgp-ad {
                admin-state enable
                vpls-id "64500:1000"
            }
        }
        vpls "I-VPLS 1001" {
            admin-state enable
            service-id 1001
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS 1000" {
                    isid 1001
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            sap 1/2/1:1001 {
            }
        }
```

```
[/]
A:admin@PE-3# show service id 1000 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 1000               Vpn Id            : 0
Service Type      : b-VPLS

---snip---

Oper Backbone Src : 00:00:00:00:00:03
Use SAP B-MAC     : Disabled
i-Vpls Count      : 1
Epipe Count       : 1
Use ESI B-MAC     : Disabled

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                              Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sdp:32765:4294967293 SB(192.0.2.5)      BgpAd        0       8978    Up   Up
sdp:32766:4294967294 SB(192.0.2.4)      BgpAd        0       8978    Up   Up
sdp:32767:4294967295 SB(192.0.2.2)      BgpAd        0       8978    Up   Up
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

Multiple MAC Registration Protocol (MMRP) is not used in the B-VPLS instance. If it were enabled, MMRP would have to be disabled in the network before this migration. If there are ISIDs using B-VPLS SDP-bindings to reach some remote locations and B-VPLS EVPN destinations to reach others, the default multicast list must be used in the current release (the MFIB cannot be used if there is a mix of both types). Therefore, during the migration process, the ISIDs must be added to the default multicast list.

1. Add service-level SHG (if not already there).

   From the first node being migrated to PBB-EVPN to all nodes migrated, PBB-VPLS and PBB-EVPN have to coexist within the same meshed network. That is, EVPN-MPLS destinations and SDP-bindings need to be defined in the same split-horizon group. Therefore, if there is no split-horizon group defined in the B-VPLS, the first step is to add it. In this example, the split-horizon group is defined at the **config>service>pw-template>level**; therefore, it has to be added at the B-VPLS level.

   Note:
   When the **service>split-horizon-group** is removed, an eval-pw-template must be performed.

   Note:
   After adding the **split-horizon-group** at the service level, an eval-pw-template must be performed again so that the SDP-bindings take the new SHG configuration.

   Note:
   During the time between the **split-horizon-group** being removed and added back again, the SDP-bindings can forward BUM traffic to each other, so this operation must be done carefully to avoid loops.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Assuming that the first node to be migrated is PE-3, the following output shows the procedure for adding the **split-horizon-group** at the service level.

```
# on PE-3:
configure exclusive
    service {
        pw-template "PW1" {
            pw-template-id 1
            delete split-horizon-group
        }
        commit
```

```
[/]
A:admin@PE-3# tools perform service id 1000 eval-pw-template 1 allow-service-impact
eval-pw-template succeeded for Svc 1000 32765:4294967293 Policy 1
eval-pw-template succeeded for Svc 1000 32766:4294967294 Policy 1
eval-pw-template succeeded for Svc 1000 32767:4294967295 Policy 1
```

```
# on PE-3:
configure exclusive
    service {
        vpls "B-VPLS 1000" {
            bgp 1 {
                pw-template-binding "PW1" {
                    split-horizon-group "CORE"
                }
            }
            split-horizon-group "CORE" {
            }
            commit
```

```
[/]
A:admin@PE-3# tools perform service id 1000 eval-pw-template 1 allow-service-impact
eval-pw-template succeeded for Svc 1000 32765:4294967293 Policy 1
eval-pw-template succeeded for Svc 1000 32766:4294967294 Policy 1
eval-pw-template succeeded for Svc 1000 32767:4294967295 Policy 1
```

```
[ex:configure service vpls "B-VPLS 1000"]
A:admin@PE-3# info
    admin-state enable
    service-id 1000
    customer "1"
    service-mtu 2000
    pbb-type b-vpls
    pbb {
        source-bmac {
            address 00:00:00:00:00:03
        }
    }
    bgp 1 {
        pw-template-binding "PW1" {
            split-horizon-group "CORE"
        }
    }
    bgp-ad {
        admin-state enable
        vpls-id "64500:1000"
    }
    split-horizon-group "CORE" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    }
```

2. Add BGP-EVPN and ISID-policy configuration to the B-VPLS.

   After the B-VPLS is configured with the split horizon group, the BGP-EVPN configuration (with admin-state disabled) and ISID-policy can be added, as follows.

```
# on PE-3:
configure {
    service {
        vpls "B-VPLS 1000" {
            bgp-evpn {
                evi 1000
                mpls 1 {
                    admin-state disable     # must be disabled
                    split-horizon-group "CORE"
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            isid-policy {
                entry 10 {
                    use-def-mcast true
                    range {
                        start 1001
                        end 3000
                    }
                }
            }
```

3. Enable BGP-EVPN MPLS on the PE.

   When the configuration is ready, the **bgp-evpn mpls 1** context can be enabled, as follows:

```
# on PE-3:
configure {
    service {
        vpls "B-VPLS 1000" {
            bgp-evpn {
                mpls 1 {
                    admin-state enable
                }
            }
```

   Enabling the **bgp-evpn mpls 1** context triggers a route-refresh message for the EVPN family from PE-3, but no changes happen because PE-3 does not create any EVPN destinations until it imports EVPN routes from the other PEs. The three spoke-SDPs to the remote PEs are still up.

4. Repeat steps 1 to 3 for the second PE (PE-5).

   The same steps 1 to 3 are repeated for PE-5. When the **bgp-evpn mpls 1** context is enabled, PE-5 sends a route-refresh and gets the BGP-EVPN routes from PE-3. As a result of that, PE-3 brings down the spoke-SDP to PE-5 and creates an EVPN destination to PE-5. The same process happens in PE-5. The following CLI output shows the received routes in PE-3 and spoke-SDP going down.

```
45 2021/03/05 09:57:37.206 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Received BGP UPDATE:
    Withdrawn Length = 0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    Total Path Attr Length = 110
    Flag: 0x90 Type: 14 Len: 47 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.5
        Type: EVPN-INCL-MCAST Len: 17 RD: 64500:1000, tag: 1001, orig_addr len: 32,
                              orig_addr: 192.0.2.5
        Type: EVPN-INCL-MCAST Len: 17 RD: 64500:1000, tag: 0, orig_addr len: 32,
                              orig_addr: 192.0.2.5
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0x80 Type: 9 Len: 4 Originator ID: 192.0.2.5
    Flag: 0x80 Type: 10 Len: 4 Cluster ID:
        1.1.1.1
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:1000
        bgp-tunnel-encap:MPLS
    Flag: 0xc0 Type: 22 Len: 9 PMSI:
        Tunnel-type Ingress Replication (6)
        Flags: (0x0)[Type: None BM: 0 U: 0 Leaf: not required]
        MPLS Label 8388464
        Tunnel-Endpoint 192.0.2.5
"
```

Log 99 shows that spoke SDP 32765:4294967293 is operationally down:

```
184 2021/03/05 09:57:39.472 CET MINOR: SVCMGR #2313 Base
"Status of SDP Bind 32765:4294967293 in service 1000 (customer 1) peer PW status bits
 changed to pwNotForwarding "

183 2021/03/05 09:57:37.207 CET MINOR: SVCMGR #2306 Base
"Status of SDP Bind 32765:4294967293 in service 1000 (customer 1) changed to admin=up oper=
down flags=evpnRouteConflict "

182 2021/03/05 09:57:37.207 CET MINOR: SVCMGR #2326 Base
"Status of SDP Bind 32765:4294967293 in service 1000 (customer 1) local PW status bits
 changed to pwNotForwarding "
```

Spoke SDP 32765:4294967293 is the spoke SDP toward PE-5 and it is kept down:

```
[/]
A:admin@PE-3# show service id 1000 base
---snip---

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                                Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sdp:32765:4294967293 SB(192.0.2.5)        BgpAd        0       8978    Up   Down
sdp:32766:4294967294 SB(192.0.2.4)        BgpAd        0       8978    Up   Up
sdp:32767:4294967295 SB(192.0.2.2)        BgpAd        0       8978    Up   Up
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

The reason why the spoke SDP toward PE-5 is down is an EVPN route conflict:

```
[/]
A:admin@PE-3# show service id 1000 sdp 32765:4294967293 detail | match Flag post-lines 1
Flags              : PWPeerFaultStatusBits
                     EvpnRouteConflict
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

An EVPN destination to PE-5 is created:

```
[/]
A:admin@PE-3# show service id 1000 evpn-mpls

===============================================================================
BGP EVPN-MPLS Dest
===============================================================================
TEP Address     Egr Label     Num. MACs   Mcast           Last Change
                Transport:Tnl                             Sup BCast Domain
-------------------------------------------------------------------------------
192.0.2.5       524279        1           bum             03/05/2021 09:57:37
                ldp:65539                                 No
-------------------------------------------------------------------------------
Number of entries : 1
-------------------------------------------------------------------------------
===============================================================================
---snip---
```

5. Repeat Steps 1 to 3 for the rest of the PEs (PE-2, PE-4).

   The same process is repeated in all the PEs, node by node. The service impact for the I-VPLS 1001 is minimal.

6. (Optional) Remove the ISID policy.

   When all the PEs in the B-VPLS 1000 are migrated, the ISID policy can optionally be removed, node by node. This forces the B-VPLS instance to start using the MFIB to send I-VPLS BUM traffic to the remote nodes. This has no effect on Epipes (traffic is always unicast for Epipes).

   Before removing the ISID policy and starting to use the MFIB, it is recommended to check that the Inclusive Multicast routes for an ISID to the remote PEs are all active. Otherwise, connectivity for BUM traffic could be interrupted if any of the expected routes are not active. This is illustrated for PE-3.

```
[/]
A:admin@PE-3# show service id 1000 evpn-mpls

===============================================================================
BGP EVPN-MPLS Dest
===============================================================================
TEP Address     Egr Label     Num. MACs   Mcast           Last Change
                Transport:Tnl                             Sup BCast Domain
-------------------------------------------------------------------------------
192.0.2.2       524282        1           bum             03/05/2021 10:00:32
                ldp:65537                                 No
192.0.2.4       524282        1           bum             03/05/2021 10:00:33
                ldp:65538                                 No
192.0.2.5       524279        1           bum             03/05/2021 09:57:37
                ldp:65539                                 No
-------------------------------------------------------------------------------
Number of entries : 3
-------------------------------------------------------------------------------
===============================================================================
---snip---
```

The routes for ISID 1001 are valid and used by BGP (flags **u*>i**):

```
[/]
A:admin@PE-3# show router bgp routes evpn incl-mcast tag 1001
===============================================================================
 BGP Router ID:192.0.2.3        AS:64500       Local AS:64500
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Inclusive-Mcast Routes
===============================================================================
Flag  Route Dist.         OrigAddr
      Tag                 NextHop
-------------------------------------------------------------------------------
u*>i  64500:1000          192.0.2.2
      1001                192.0.2.2

u*>i  64500:1000          192.0.2.4
      1001                192.0.2.4

u*>i  64500:1000          192.0.2.5
      1001                192.0.2.5

-------------------------------------------------------------------------------
Routes : 3
===============================================================================
```

There are no entries in the MFIB:

```
[/]
A:admin@PE-3# show service id 1000 mfib

===============================================================================
Multicast FIB, Service 1000
===============================================================================
Source Address  Group Address         Port Id                     Svc Id   Fwd
                                                                         Blk
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Number of entries: 0
===============================================================================
```

The ISID policy is removed as follows:

```
# on all PEs:
configure {
    service {
        vpls "B-VPLS 1000" {
            delete isid-policy
```

After removing the ISID-policy, the MFIB is populated with entries for the ISID 1001 group BMAC to the three remote PEs where ISID 1001 is defined:

```
[/]
A:admin@PE-3# show service id 1000 mfib

===============================================================================
Multicast FIB, Service 1000
===============================================================================
Source Address  Group Address         Port Id                     Svc Id   Fwd
                                                                         Blk
-------------------------------------------------------------------------------
*               01:1e:83:00:03:e9     b-mpls:192.0.2.2:524282     Local    Fwd
                                      b-mpls:192.0.2.4:524282     Local    Fwd
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                                    b-mpls:192.0.2.5:524279      Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
```

7. (Optional) Remove the BGP-AD configuration.

   The BGP-AD configuration can stay in the B-VPLS services. However, when the entire network is migrated to PBB-EVPN, all the spoke-SDPs will be operationally down and, even if they are not forwarding traffic, they consume resources in the system. Consider removing the BGP-AD configuration and, therefore, the spoke-SDPs.

   The following example shows the removal of BGP-AD in PE-4. Be aware that when BGP-AD is removed from the configuration, if the RD/RT was derived from the VPLS ID (as in this example), a new RD/RT must be auto-derived for the service. Therefore, new updates will be sent for all the EVPN NLRIs, as shown in the following output.

```
[/]
A:admin@PE-4# show service id 1000 bgp

===============================================================================
BGP Information
===============================================================================
Bgp Instance         : 1
Vsi-Import           : None
Vsi-Export           : None
Route Dist           : None
Oper Route Dist      : 64500:1000
Oper RD Type         : derivedVpls
Rte-Target Import    : None                 Rte-Target Export: None
Oper RT Imp Origin   : derivedVpls          Oper RT Import   : 64500:1000
Oper RT Exp Origin   : derivedVpls          Oper RT Export   : 64500:1000

PW-Template Id       : 1                    PW-Template SHG  : CORE
Oper Group           : None
Mon Oper Group       : None
BFD Template         : None
BFD-Enabled          : no                   BFD-Encap        : ipv4
Import Rte-Tgt       : None
-------------------------------------------------------------------------------
===============================================================================
```

   On all PEs, the BGP-AD configuration and the PW template binding are removed, as follows:

```
# on all PEs:
configure {
    service {
        vpls "B-VPLS 1000" {
            delete bgp-ad
            bgp 1 {
                delete pw-template-binding "PW1"
```

   After BGP-AD is disabled, the spoke SDP bindings are deleted. The following messages are logged in log 99 on PE-4:

```
179 2021/03/05 10:05:41.204 CET MAJOR: SVCMGR #2319 Base
"Dynamic bgp-l2vpn SDP 32765 (192.0.2.5) was deleted."

178 2021/03/05 10:05:41.204 CET MINOR: SVCMGR #2303 Base
"Status of SDP 32765 changed to admin=down oper=down"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
177 2021/03/05 10:05:41.204 CET MAJOR: SVCMGR #2320 Base
"Service Id 1000, Dynamic bgp-l2vpn SDP Bind Id 32765:4294967292 was deleted."

176 2021/03/05 10:05:41.194 CET MINOR: SVCMGR #2306 Base
"Status of SDP Bind 32765:4294967292 in service 1000 (customer 1) changed to admin=down
 oper=down flags="
```

Initially, the RD/RT was derived from the VPLS ID (64500:1000). After the BGP-AD configuration is removed, a new RD and RT must be auto-derived from the EVI:

```
[/]
A:admin@PE-4# show service id 1000 bgp

===============================================================================
BGP Information
===============================================================================
Bgp Instance         : 1
Vsi-Import           : None
Vsi-Export           : None
Route Dist           : None
Oper Route Dist      : 192.0.2.4:1000
Oper RD Type         : derivedEvi
Rte-Target Import    : None                Rte-Target Export: None
Oper RT Imp Origin   : derivedEvi          Oper RT Import   : 64500:1000
Oper RT Exp Origin   : derivedEvi          Oper RT Export   : 64500:1000

PW-Template Id       : None
-------------------------------------------------------------------------------
===============================================================================
```

In this case, the system picks up the RD in the following order:

a. Manual RD or auto-RD always take precedence when configured.

b. If no manual/auto-RD, the RD is derived from the **bgp-ad vpls-id**.

c. If no manual/auto-rd/vpls-id configuration, the RD is derived from the **bgp evpn evi**.

d. If no manual/auto-rd/vpls-id/evi configuration, there will be no RD and the service will fail.

If in the migration from BGP-AD to BGP-EVPN, the advertisement of new updates is not needed, the initial configuration must include manual/auto-RDs. If manual/auto-RDs were not included, disabling BGP-AD would not cause the change of RD and the consequent BGP updates.

PBB-EVPN multi-homing

This section provides configuration guidelines for PBB-EVPN multi-homing. In the same way that EVPN-MPLS supports single-active and all-active multi-homing, PBB-EVPN can also be configured to support both modes. The same Ethernet segment that is used for regular EVPN-MPLS service SAPs and spoke-SDPs can be shared with I-VPLS/Epipe SAPs and spoke-SDPs.

Figure 102: PBB-EVPN multi-homing shows the example topology used in this section.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 102: PBB-EVPN multi-homing

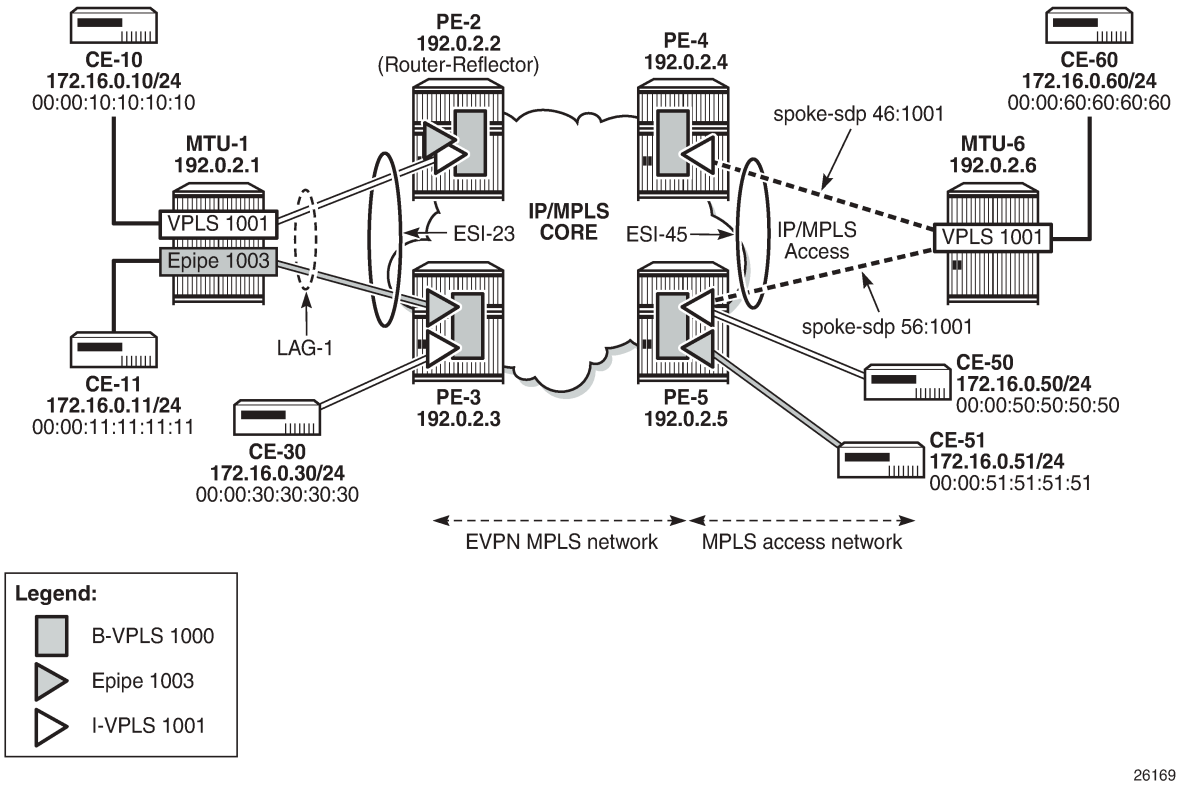

MTU-1 and MTU-6 have been added to the network (compared to Figure 100: PBB-EVPN network without multi-homing). I-VPLS 1001 has two new sites that are multi-homed to the PBB-EVPN network. MTU-1 uses all-active multi-homing, whereas MTU-6 is connected to a single-active ES. As with EVPN-MPLS, all-active multi-homing is only supported when a LAG is used at the access. Single-active multi-homing can be supported with regular Ethernet ports (that can form an independent LAG per PE) or SDPs.

RFC 7623 describes two types of system BMAC assignments that a PE can implement in a B-VPLS when ESs are present:

- Shared BMAC addresses that can be used for all the single-homed CEs and a number of multi-homed CEs connected to Ethernet-segments.
- Dedicated BMAC addresses per Ethernet segment.

In this chapter and in SR OS terminology:

- A shared BMAC address (in IETF) is a source BMAC address as configured in **service>(b)vpls>pbb>source-bmac**. All the I-VPLS/Epipe traffic coming from single-homed CEs is sent encapsulated in a PBB packet with that source BMAC address.
- A dedicated-BMAC per ES (in IETF) is an ES BMAC address as activated in **service>(b)vpls>pbb>source-bmac>use-es-bmac-lsb** and generated from the combination of **vpls>pbb>source-bmac** plus **ethernet-segment>pbb>source-bmac-lsb**. If configured, any I-VPLS/Epipe traffic coming from an ES is encapsulated in a PBB packet with the ES-BMAC address as the source BMAC address.

The system allows the following user choices per B-VPLS and ES:

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- A dedicated ES BMAC address per ES can be used. In that case, the **pbb>source-bmac>use-es-bmac-lsb** command is configured in the B-VPLS. In all-active multi-homing, all the PEs that are part of the ES source the PBB packets with the same source ES BMAC address; single-active multi-homing requires the use of a different ES BMAC address per PE.
- A non-dedicated source BMAC address can be used (this is only possible in single-active multi-homing). In this case, the user does not configure **pbb>source-bmac>use-es-bmac-lsb** and the regular source BMAC address is used for the traffic. A different source BMAC address has to be advertised per PE.

As discussed, single-active multi-homing can use source BMAC addresses or ES BMAC addresses. Using one type or another has a different impact on CMAC flushing, as illustrated in Figure 103: The use of ES BMAC to minimize CMAC flush.

- If ES BMAC addresses are used, as shown on the right-hand side of Figure 103: The use of ES BMAC to minimize CMAC flush, a less-impacting CMAC flush is achieved, therefore minimizing the flooding after ES failures. In the case of ES failure, PE-1 withdraws the ES BMAC address 00:12 and the remote PE-3 only flushes the CMACs associated with that ES BMAC address (only the CMAC addresses behind the CE are flushed).
- If source BMAC addresses are used, as shown on the left-hand side of Figure 103: The use of ES BMAC to minimize CMAC flush, in the case of ES failure, a BGP update with higher sequence number is issued by PE-1 and the remote PE-3 flushes all the CMAC addresses associated with the source BMAC address. Therefore, all the CMAC addresses behind the B-VPLS of the PEs will be flushed, as opposed to only the CMAC addresses behind the CE of the Ethernet Service Instances (ESIs).

Figure 103: The use of ES BMAC to minimize CMAC flush

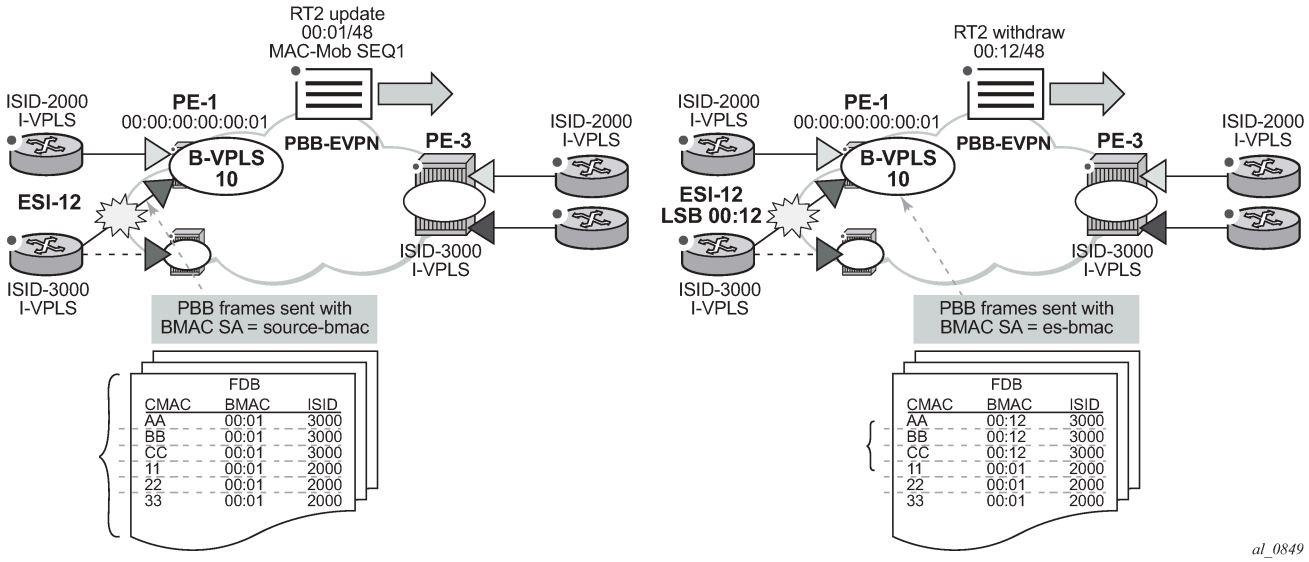

Table 7: PBB-EVPN multi-homing supported combinations in SR OS shows the PBB-EVPN multi-homing combinations supported in the current release in the topology of Figure 102: PBB-EVPN multi-homing.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Table 7: PBB-EVPN multi-homing supported combinations in SR OS

CE Connectivity	PE Connectivity	PE Redundancy	BMAC Assignment	I-VPLS Support	Epipe Support
LAG (LACP optional)	LAG SAP	EVPN MH all-active	use-es-bmac-lsb (shared BMAC)	Yes	Yes
Ethernet ports (no LAG)	LAG SAP or port SAP	EVPN MH single-active	use-es-bmac-lsb (dedicated per PE)	Yes	No
Ethernet ports (no LAG)	LAG SAP or port SAP	EVPN MH single-active	source-bmac address (dedicated per PE)	Yes	No
MPLS	spoke-SDP	EVPN MH single-active	source-bmac address (dedicated per PE)	Yes	No
MPLS	spoke-SDP	EVPN MH single-active	use-es-bmac-lsb (dedicated per PE)	Yes	No

As an example, the configurations of the first, and last two, rows (LAG SAP all-active, MPLS source-BMAC, and MPLS ES-BMAC, respectively) will be discussed in the following three sections.

PBB-EVPN all-active multi-homing for I-VPLS and Epipes

Figure 102: PBB-EVPN multi-homing shows a PBB-EVPN network where ESI-23 is configured as an all-active multi-homing ES on PE-2 and PE-3. Two services are using ESI-23: I-VPLS 1001 and Epipe 1003. The following output shows the relevant configuration in PE-2:

```
# on PE-2:
configure {
    service {
        pbb {
            mac "PE-5" {
                address 00:00:00:00:00:05
            }
        }
        system {
            bgp {
                evpn {
                    ethernet-segment "ESI-23" {
                        admin-state enable
                        esi 01:00:00:00:00:23:00:00:00:01
                        multi-homing-mode all-active
                        df-election {
                            es-activation-timer 3
                        }
                        association {
                            lag "lag-1" {
                            }
                        }
                        pbb {
                            source-bmac-lsb 23-23
                        }
                    }
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
        }
        vpls "B-VPLS 1000" {
            admin-state enable
            service-id 1000
            customer "1"
            service-mtu 2000
            pbb-type b-vpls
            pbb {
                source-bmac {
                    address 00:00:00:00:00:02
                    use-es-bmac-lsb true
                }
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 1000
                mpls 1 {
                    admin-state enable
                    split-horizon-group "CORE"
                    ecmp 2
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            split-horizon-group "CORE" {
            }
        }
        vpls "I-VPLS 1001" {
            admin-state enable
            service-id 1001
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS 1000" {
                    isid 1001
                }
            }
            sap lag-1:1001 {
            }
        }
        epipe "Epipe 1003" {
            admin-state enable
            service-id 1003
            customer "1"
            pbb {
                tunnel {
                    backbone-vpls-service-name "B-VPLS 1000"
                    isid 1003
                    backbone-dest-mac-name "PE-5"
                }
            }
            sap lag-1:1003 {
            }
        }
```

The following output shows the relevant configuration in PE-3:

```
# on PE-3:
configure {
    service {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        pbb {
            mac "PE-5" {
                address 00:00:00:00:00:05
            }
        }
        system {
            bgp {
                evpn {
                    ethernet-segment "ESI-23" {
                        admin-state enable
                        esi 01:00:00:00:00:23:00:00:00:01
                        multi-homing-mode all-active
                        df-election {
                            es-activation-timer 3
                        }
                        association {
                            lag "lag-1" {
                            }
                        }
                        pbb {
                            source-bmac-lsb 23-23
                        }
                    }
                }
            }
        }
        vpls "B-VPLS 1000" {
            admin-state enable
            service-id 1000
            customer "1"
            service-mtu 2000
            pbb-type b-vpls
            pbb {
                source-bmac {
                    address 00:00:00:00:00:03
                    use-es-bmac-lsb true
                }
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 1000
                mpls 1 {
                    admin-state enable
                    split-horizon-group "CORE"
                    ecmp 2
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            split-horizon-group "CORE" {
            }
        }
        vpls "I-VPLS 1001" {
            admin-state enable
            service-id 1001
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS 1000" {
                    isid 1001
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            sap 1/2/1:1001 {
            }
            sap lag-1:1001 {
            }
        }
        epipe "Epipe 1003" {
            admin-state enable
            service-id 1003
            customer "1"
            pbb {
                tunnel {
                    backbone-vpls-service-name "B-VPLS 1000"
                    isid 1003
                    backbone-dest-mac-name "PE-5"
                }
            }
            sap lag-1:1003 {
            }
        }
```

The preceding configuration shows that Epipe 1003 has a PBB tunnel pointing at the PE-5 source-BMAC. Epipe 1003 has the following configuration in PE-5 (the PBB tunnel points at the ESI-23 ES-BMAC):

```
# on PE-5:
configure {
    service {
        pbb {
            mac "ES-MAC-23" {
                address 00:00:00:00:23:23
            }
        }
        epipe "Epipe 1003" {
            admin-state enable
            service-id 1003
            customer "1"
            pbb {
                tunnel {
                    backbone-vpls-service-name "B-VPLS 1000"
                    isid 1003
                    backbone-dest-mac-name "ES-MAC-23"
                }
            }
            sap 1/2/1:1003 {
            }
        }
```

Source-BMAC addresses and ES-BMAC addresses are distributed in BGP-EVPN. PE-2 and PE-3 will each advertise their own source-BMAC in a MAC route with ESI-0 and the shared ES-BMAC address with ESI-MAX (as per the RFC 7623). The ES-BMAC address that each PE uses in a B-VPLS is derived from the configured **service>(b)vpls>pbb>source-bmac** (four high-order bytes) and the ESI-23 configured **source-bmac-lsb**. In this example, PE-2 and PE-3 will both derive ES-BMAC 00:00:00:00:23:23. For both PEs to derive the required same ES-BMAC address, the four high-order bytes of the source-BMAC address must match on both PEs.

The **es-bmac-table-size** parameter modifies the default value (8) for the maximum number of ES-BMAC addresses that can be associated with the Ethernet segment across different B-VPLS services. When **source-bmac-lsb** is configured, the associated **es-bmac-table-size** is reserved out of the total FDB space.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The following outputs show the source-BMAC addresses and ES-BMAC address and how they are advertised and installed in the B-VPLS FDB.

```
[/]
A:admin@PE-2# show service system bgp-evpn ethernet-segment name "ESI-23" | match BMAC
Source BMAC LSB         : 23-23
```

The following output shows that ES-BMAC is used and that the operational source-BMAC is 00:00:00:00:00:02.

```
[/]
A:admin@PE-2# show service id 1000 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 1000               Vpn Id            : 0
Service Type      : b-VPLS
---snip---
Oper Backbone Src : 00:00:00:00:00:02
Use SAP B-MAC     : Disabled
i-Vpls Count      : 1
Epipe Count       : 1
Use ESI B-MAC     : Enabled
---snip---
```

The source BMAC LSB is configured with the same value on PE-2 and PE-3. The two low-order bytes of the ES-BMAC will be 23:23.

```
[/]
A:admin@PE-3# show service system bgp-evpn ethernet-segment name "ESI-23" | match BMAC
Source BMAC LSB         : 23-23
```

On PE-3, ES-BMAC is used and the operational source BMAC is 00:00:00:00:00:03, as follows:

```
[/]
A:admin@PE-3# show service id 1000 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 1000               Vpn Id            : 0
Service Type      : b-VPLS
---snip---
Oper Backbone Src : 00:00:00:00:00:03
Use SAP B-MAC     : Disabled
i-Vpls Count      : 1
Epipe Count       : 2
Use ESI B-MAC     : Enabled
---snip---
```

On PE-2, the FDB for B-VPLS 1000 has an entry for each of the other PEs. PEs do not show their own system BMAC addresses in the FDB:

```
[/]
A:admin@PE-2# show service id 1000 fdb detail

===============================================================================
Forwarding Database, Service 1000
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
1000       00:00:00:00:00:03 mpls:                   EvpnS:P  03/05/21 10:05:43
                             192.0.2.3:524282
           ldp:65537
1000       00:00:00:00:00:04 mpls:                   EvpnS:P  03/05/21 10:05:41
                             192.0.2.4:524282
           ldp:65538
1000       00:00:00:00:00:05 mpls:                   EvpnS:P  03/05/21 10:05:42
                             192.0.2.5:524279
           ldp:65539
-------------------------------------------------------------------------------
No. of MAC Entries: 3
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

On PE-4, the FDB for B-VPLS 1000 has an entry for each of the other PEs and an entry for the ES-BMAC address of ES "ESI-23":

```
[/]
A:admin@PE-4# show service id 1000 fdb detail

===============================================================================
Forwarding Database, Service 1000
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
1000       00:00:00:00:00:02 mpls:                   EvpnS:P  03/05/21 10:05:44
                             192.0.2.2:524282
           ldp:65538
1000       00:00:00:00:00:03 mpls:                   EvpnS:P  03/05/21 10:05:43
                             192.0.2.3:524282
           ldp:65537
1000       00:00:00:00:00:05 mpls:                   EvpnS:P  03/05/21 10:05:42
                             192.0.2.5:524279
           ldp:65539
1000       00:00:00:00:23:23 eES:                    EvpnS:P  03/05/21 10:07:25
                             MAX-ESI
-------------------------------------------------------------------------------
No. of MAC Entries: 4
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

On PE-4, there are two BGP routes for ES-BMAC address 00:00:00:00:23:23: one with next hop PE-2 and the other with next hop PE-3, as follows:

```
[/]
A:admin@PE-4# show router bgp routes evpn mac mac-address 00:00:00:00:23:23
===============================================================================
 BGP Router ID:192.0.2.4        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
BGP EVPN MAC Routes
===============================================================================
Flag  Route Dist.         MacAddr           ESI
      Tag                 Mac Mobility      Label1
                          Ip Address
                          NextHop
-------------------------------------------------------------------------------
u*>i  192.0.2.2:1000      00:00:00:00:23:23 ESI-MAX
      0                   Static            LABEL 524282
                          n/a
                          192.0.2.2

u*>i  192.0.2.3:1000      00:00:00:00:23:23 ESI-MAX
      0                   Static            LABEL 524282
                          n/a
                          192.0.2.3

-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

PBB-EVPN all-active multi-homing is based on the same concepts as EVPN-MPLS all-active multi-homing: DF election, split-horizon, and aliasing.

Designated forwarder (DF) election

Only the DF PE for an ISID will send multicast traffic to the ES. The following command shows that PE-3 is the DF PE in ES "ESI-23" for ISID 1003:

```
[/]
A:admin@PE-3# show service system bgp-evpn ethernet-segment name "ESI-23"
                                                                    isid isid-2 1003

===============================================================================
ISID DF and Candidate List
===============================================================================
Isid          SvcId         Actv Timer Rem      DF  DF Last Change
-------------------------------------------------------------------------------
1003          1003          0                   yes 03/05/2021 10:07:45
===============================================================================
---snip---
```

The following command shows the DF and DF candidate list in the ES for all EVIs and all ISIDs:

```
[/]
A:admin@PE-3# show service system bgp-evpn ethernet-segment name "ESI-23" all

===============================================================================
Service Ethernet Segment
===============================================================================
Name                    : ESI-23
Eth Seg Type            : None
Admin State             : Enabled            Oper State         : Up
ESI                     : 01:00:00:00:00:23:00:00:00:01
Multi-homing            : allActive          Oper Multi-homing  : allActive
---snip---

===============================================================================
ISID Information
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
ISID                SvcId               Actv Timer Rem      DF
-------------------------------------------------------------------------------
1001                1001                0                   yes
1003                1003                0                   yes
-------------------------------------------------------------------------------
Number of entries: 2
===============================================================================

-------------------------------------------------------------------------------
DF Candidate list
-------------------------------------------------------------------------------
ISID                                    DF Address
-------------------------------------------------------------------------------
1001                                    192.0.2.2
1001                                    192.0.2.3
1003                                    192.0.2.2
1003                                    192.0.2.3
-------------------------------------------------------------------------------
Number of entries: 4
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
---snip---
```

The DF PE identifies multicast traffic by looking at either the destination BMAC or the EVPN label (which can be unicast or multicast).

In the case of Epipes, there are also DF and non-DF PEs. However, traffic is usually unicast (sent to the PBB tunnel backbone-destination BMAC). The non-DF PE will usually not discard Epipe traffic to the ES, unless the packet comes with an EVPN multicast label. To avoid packet duplication at the CE for Epipes, it is recommended to either:

- configure **discard-unknown** on all the B-VPLS instances where there are PBB-Epipes. This will prevent the ingress PE from flooding Epipe traffic if the PBB tunnel BMAC is unknown in the FDB.
- configure **ingress-replication-bum-label true** so that, when the PBB tunnel BMAC is unknown in the FDB, the ingress PE sends traffic with a multicast label. The non-DF will discard traffic identified as multicast at Epipes.

Ethernet segment split-horizon

In PBB-EVPN all-active multi-homing, the split-horizon function is not based in the ESI label but in a source BMAC check. When BUM traffic is received on an I-VPLS, the PE will encapsulate it in PBB using the ES-BMAC as source BMAC and the group BMAC for the ISID. When the DF PE for the ISID receives that packet, it will not send it back to the ES if the packet is identified as being originated from the ES itself (based on the ES-BMAC shared between the PEs).

Aliasing

Aliasing is based on the advertisement of the same ES-BMAC with MAX-ESI from the PEs part of the same ES. PE-2 and PE-3 advertise the ES-BMAC 00:00:00:00:23:23 with MAX-ESI (ESI = all FFs, as per the RFC 7623) and as Static (protected). When the remote PEs, PE-4, and PE-5, receive the two routes for the same BMAC and MAX-ESI, they will create a single EVPN-MPLS destination that will give more than one next-hop (in this case 2), as long as ECMP > 1:

```
[/]
A:admin@PE-4# show service id 1000 evpn-mpls
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
BGP EVPN-MPLS Dest
===============================================================================
TEP Address     Egr Label     Num. MACs   Mcast           Last Change
                Transport:Tnl                             Sup BCast Domain
-------------------------------------------------------------------------------
192.0.2.2       524282        1           bum             03/05/2021 10:05:44
                ldp:65538                                 No
192.0.2.3       524282        1           bum             03/05/2021 10:05:43
                ldp:65537                                 No
192.0.2.5       524279        1           bum             03/05/2021 10:05:42
                ldp:65539                                 No
-------------------------------------------------------------------------------
Number of entries : 3
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN-MPLS Ethernet Segment Dest
===============================================================================
Eth SegId                       Num. Macs               Last Change
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================

===============================================================================
BGP EVPN-MPLS ES BMAC Dest
===============================================================================
ES BMAC Addr                            Last Change
-------------------------------------------------------------------------------
00:00:00:00:23:23                       03/05/2021 10:07:45
-------------------------------------------------------------------------------
Number of entries: 1
-------------------------------------------------------------------------------
===============================================================================
```

The EVPN-MPLS ES BMAC destination has two next hops: PE-2 and PE-3.

```
[/]
A:admin@PE-4# show service id 1000 evpn-mpls es-bmac ieee-address 00:00:00:00:23:23

===============================================================================
BGP EVPN-MPLS ES BMAC Dest
===============================================================================
ES BMAC Addr                            Last Change
-------------------------------------------------------------------------------
00:00:00:00:23:23                       03/05/2021 10:07:45
===============================================================================

===============================================================================
BGP EVPN-MPLS ES BMAC Dest TEP Info
===============================================================================
TEP Address               Egr Label               Last Change          Tunnel-
                          Transport                                    Id
-------------------------------------------------------------------------------
192.0.2.2                 524282                  03/05/2021 10:07:25  65538
                          ldp
192.0.2.3                 524282                  03/05/2021 10:07:45  65537
                          ldp
-------------------------------------------------------------------------------
Number of entries : 2
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
```

A similar output will be obtained in PE-5. Unicast traffic entering I-VPLS 1001 in either PE-4 or PE-5 will be hashed and load-balanced to PE-2 and PE-3 if the destination CMAC lookup yields an **es-bmac-dest**:

```
[/]
A:admin@PE-5# show service id 1001 fdb detail pbb

===============================================================================
Forwarding Database, i-Vpls Service 1001
===============================================================================
MAC               Source-Identifier     B-Svc     b-Vpls MAC        Type/Age
 Transport:Tnl-Id
-------------------------------------------------------------------------------
00:00:10:10:10:10 eES-BMAC:             1000      00:00:00:00:23:23 L/0
                  00:00:00:00:23:23
00:00:30:30:30:30 b-mpls:               1000      00:00:00:00:00:03 L/0
                  192.0.2.3:524282
00:00:50:50:50:50 sap:1/2/1:1001        1000      N/A               L/0
00:00:60:60:60:60 sdp:56:1001           1000      N/A               L/0
===============================================================================
```

Verify the FDB of I-VPLS 1001 for ES BMAC destination 00:00:00:00:23:23 as follows:

```
[/]
A:admin@PE-5# show service id 1001 fdb evpn-mpls es-bmac-dest 00:00:00:00:23:23

===============================================================================
Forwarding Database, Service 1001
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
1001       00:00:10:10:10:10 eES-BMAC:               L/30     03/05/21 10:30:52
                             00:00:00:00:23:23
-------------------------------------------------------------------------------
No. of Entries: 1
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

If a failure occurs in the ES, the PE will withdraw the ES-BMAC and the remote PEs will remove one next-hop of the ES-BMAC EVPN-MPLS destination.

For PBB-Epipes, aliasing will also work, as long as shared-queuing or policing are enabled on the ingress PE Epipe. In Figure 102: PBB-EVPN multi-homing, Epipe 1003 on PE-5 requires shared-queuing or policing at the ingress SAP. Otherwise, the traffic will be sent to only one PE of the ES (usually to the lower-IP PE).

For more information about the configuration of the Ethernet segment and its parameters, see the EVPN for MPLS Tunnels chapter.

PBB-EVPN single-active multi-homing for I-VPLS with source BMAC addresses

ESI-45 is a single-active Ethernet-segment (see Figure 102: PBB-EVPN multi-homing) with SDPs linked to it. As indicated in Table 7: PBB-EVPN multi-homing supported combinations in SR OS , only I-VPLS services can be used in this configuration. As described in section PBB-EVPN multi-homing, single-

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

active ES and B-VPLS services can be configured to either use source-BMAC addresses or ES-BMAC addresses. The following configuration shows the former option on PE-4:

```
# on PE-4:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "ESI-45" {
                        admin-state enable
                        esi 01:00:00:00:00:45:00:00:00:01
                        multi-homing-mode single-active
                        df-election {
                            es-activation-timer 3
                        }
                        association {
                            sdp 46 {
                            }
                        }
                        pbb {
                            source-bmac-lsb 45-04
                        }
                    }
                }
            }
        }
        sdp 46 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.6
            }
        }
        vpls "B-VPLS 1000" {
            admin-state enable
            service-id 1000
            customer "1"
            service-mtu 2000
            pbb-type b-vpls
            pbb {
                source-bmac {
                    address 00:00:00:00:00:04
                }
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 1000
                mpls 1 {
                    admin-state enable
                    split-horizon-group "CORE"
                    ecmp 2
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            split-horizon-group "CORE" {
            }
        }
        vpls "I-VPLS 1001" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            admin-state enable
            service-id 1001
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS 1000" {
                    isid 1001
                }
            }
            spoke-sdp 46:1001 {
            }
            sap 1/2/1:1001 {
            }
        }
```

The configuration on PE-5 is similar:

```
# on PE-5:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "ESI-45" {
                        admin-state enable
                        esi 01:00:00:00:00:45:00:00:00:01
                        multi-homing-mode single-active
                        df-election {
                            es-activation-timer 3
                        }
                        association {
                            sdp 56 {
                            }
                        }
                        pbb {
                            source-bmac-lsb 45-05
                        }
                    }
                }
            }
        }
        sdp 56 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.6
            }
        }
        vpls "B-VPLS 1000" {
            admin-state enable
            service-id 1000
            customer "1"
            service-mtu 2000
            pbb-type b-vpls
            pbb {
                source-bmac {
                    address 00:00:00:00:00:05
                }
            }
            bgp 1 {
            }
            bgp-evpn {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                evi 1000
                mpls 1 {
                    admin-state enable
                    split-horizon-group "CORE"
                    ecmp 2
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            split-horizon-group "CORE" {
            }
        }
        vpls "I-VPLS 1001" {
            admin-state enable
            service-id 1001
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS 1000" {
                    isid 1001
                }
            }
            spoke-sdp 56:1001 {
            }
            sap 1/2/1:1001 {
            }
        }
```

With the preceding configuration, PE-4 and PE-5 will not advertise ES-BMAC addresses with MAX-ESI. Therefore, all the remote BMACs on PE-2 and PE-3 are associated with regular backbone EVPN-MPLS destinations. The CMAC addresses will be learned in the data plane associated with local SAP/SDP-bindings or remote BMAC addresses. An example for the I-VPLS and B-VPLS FDB in PE-2 follows:

```
[/]
A:admin@PE-2# show service id 1001 fdb detail pbb

===============================================================================
Forwarding Database, i-Vpls Service 1001
===============================================================================
MAC               Source-Identifier     B-Svc     b-Vpls MAC        Type/Age
 Transport:Tnl-Id
-------------------------------------------------------------------------------
00:00:10:10:10:10 sap:lag-1:1001        1000      N/A               L/60
00:00:60:60:60:60 b-mpls:               1000      00:00:00:00:00:05 L/60
                  192.0.2.5:524279
===============================================================================
```

The B-VPLS FDB on PE-2 looks as follows:

```
[/]
A:admin@PE-2# show service id 1000 fdb detail

===============================================================================
Forwarding Database, Service 1000
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
1000       00:00:00:00:00:03 mpls:                   EvpnS:P  03/05/21 10:05:43
                             192.0.2.3:524282
           ldp:65537
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
1000       00:00:00:00:00:04 mpls:                    EvpnS:P  03/05/21 10:05:41
                             192.0.2.4:524282
           ldp:65538
1000       00:00:00:00:00:05 mpls:                    EvpnS:P  03/05/21 10:05:42
                             192.0.2.5:524279
           ldp:65539
-------------------------------------------------------------------------------
No. of MAC Entries: 3
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

In the preceding example, the DF for ISID 1001 is PE-5. With a failure event on the SDP to MTU-6, PE-5 will not withdraw the advertised source BMAC (because it is still being used as source BMAC for other services and even CEs within the same service). PE-5 will send an update of the same source BMAC instead, increasing the sequence number in the MAC mobility extended community. That will be a *flush-all-from-me* indication for the remote PEs (they will flush all the CMACs associated with the updated source BMAC, regardless of the service).

When the former DF (PE-5) comes back up, PE-4 will become non-DF and will send a CMAC flush indication using the same mechanism as described above.

The following example shows a failure of SDP 56 in PE-5 and the corresponding DF switchover and CMAC flush.

```
# on PE-5:
221 2021/03/05 10:34:56.487 CET MINOR: SVCMGR #2095 Base
"Ethernet Segment:ESI-45, ISID:1001, Designated Forwarding state changed to:false"

219 2021/03/05 10:34:56.486 CET MINOR: SVCMGR #2303 Base
"Status of SDP 56 changed to admin=up oper=down"
```

PE-5 sends a BGP update with the same source BMAC, increasing the sequence number in the MAC mobility extended community—CMAC flush:

```
# on PE-5:
97 2021/03/05 10:34:56.487 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 89
    Flag: 0x90 Type: 14 Len: 44 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.5
        Type: EVPN-MAC Len: 33 RD: 192.0.2.5:1000 ESI: ESI-0, tag: 0, mac len: 48
                       mac: 00:00:00:00:00:05, IP len: 0, IP: NULL, label1: 8388464
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:1000
        bgp-tunnel-encap:MPLS
        mac-mobility:Seq:2/Static
"
```

Individual SAP or spoke-SDP failures do not trigger any MAC flush or DF re-election. This is as per RFC 7623. In EVPN-MPLS, individual SAP/spoke-SDP failures are captured by the AD per-EVI withdrawal, which triggers a DF switchover.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

PBB-EVPN single-active multi-homing for I-VPLS with ES-BMACs

As discussed throughout this chapter, the use of ES-BMACs for single-active multi-homing can minimize the number of CMACs flushed in a network. A simple change is necessary: configure the **use-es-bmac-lsb true** command and ensure that the generated ES-BMACs in PE-4 and PE-5 are different (the **source-bmac-lsb** in the previous configuration had different values for PE-4 and PE-5 already):

```
# on PE-4 and PE-5:
configure {
    service {
        vpls "B-VPLS 1000" {
            pbb {
                source-bmac {
                    use-es-bmac-lsb true
                }
            }
```

On PE-4, the source BMAC LSB in ESI-45 is configured with a value of 45-04:

```
[/]
A:admin@PE-4# show service system bgp-evpn ethernet-segment name "ESI-45" | match BMAC
Source BMAC LSB         : 45-04
```

On PE-5, the source BMAC LSB in ESI-45 is configured with a value of 45-05:

```
[/]
A:admin@PE-5# show service system bgp-evpn ethernet-segment name "ESI-45" | match BMAC
Source BMAC LSB         : 45-05
```

The remote PEs (such as PE-2 in the following output) will receive additional BGP EVPN-MAC routes for the ES-BMACs. The following FDB on PE-2 shows the source BMAC addresses of PE-4 and PE-5 and the ES BMAC address of DF PE-5.

```
[/]
A:admin@PE-2# show service id 1000 fdb detail

===============================================================================
Forwarding Database, Service 1000
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
1000       00:00:00:00:00:03 mpls:                   EvpnS:P  03/05/21 10:05:43
                             192.0.2.3:524282
           ldp:65537
1000       00:00:00:00:00:04 mpls:                   EvpnS:P  03/05/21 10:05:41
                             192.0.2.4:524282
           ldp:65538
1000       00:00:00:00:00:05 mpls:                   EvpnS:P  03/05/21 10:05:42
                             192.0.2.5:524279
           ldp:65539
1000       00:00:00:00:45:05 mpls:                   EvpnS:P  03/05/21 10:37:14
                             192.0.2.5:524279
           ldp:65539
-------------------------------------------------------------------------------
No. of MAC Entries: 4
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
```

The benefit is that in case of a failure in ESI-45 (as before) the ES-BMAC is withdrawn and the remote PEs will only flush the CMACs associated with the remote ESI-45, as opposed to all the CMACs associated with PE-5.

PBB-EVPN single-active multi-homing for Epipes

In the network in Figure 102: PBB-EVPN multi-homing, Epipes can only support single-homing or all-active multi-homing but not single-active. For non-local-switching PBB-Epipes (there is a single SAP per Epipe), only all-active multi-homing is supported. Single-active multi-homing for local-switching enabled PBB-Epipes (two SAPs are defined within the PBB-Epipe instance) is only supported in the following scenarios.

Figure 104: PBB-EVPN single-active support for Epipes

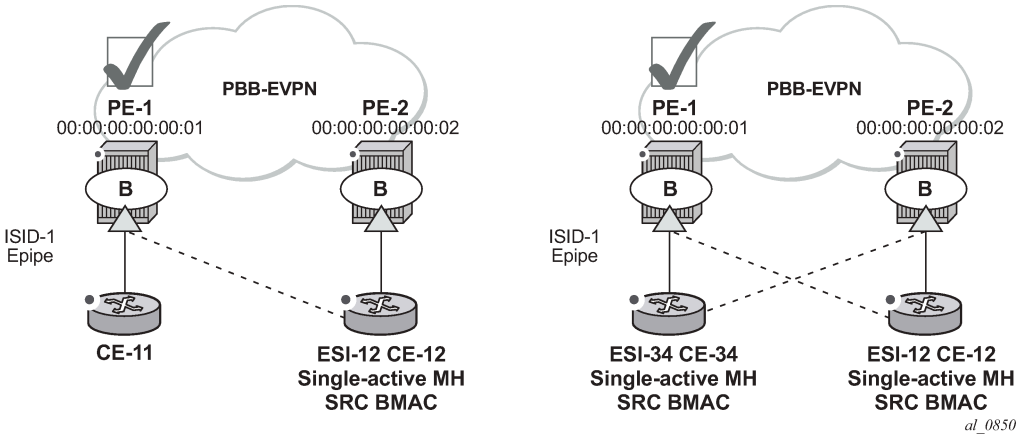

Single-active multi-homing is supported for redundancy in a two-node, three or four SAP, scenario, as displayed in Figure 104: PBB-EVPN single-active support for Epipes. In these two cases, the Epipe PBB tunnel will be configured with the source BMAC of the remote PE node. When two SAPs are active in the same Epipe, local-switching is used to exchange frames between the CEs.

All-active multi-homing is not supported for redundancy in this scenario because the PE-1 PBB tunnel cannot point at a locally defined ES-BMAC.

PBB-EVPN multi-homing operation

See the EVPN for MPLS Tunnels chapter for the commands to operate Ethernet segments. Consider that there are no AD routes in PBB-EVPN. Also, the DF election algorithm will be based on the ISID values as opposed to EVIs.

Troubleshooting and debug commands

When troubleshooting PBB-EVPN networks, most of the troubleshooting commands discussed in EVPN for MPLS Tunnels can be used in the B-VPLS service and the base **service>system>bgp-evpn** instance. Some examples of useful commands are:

- show redundancy bgp-evpn-multi-homing

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- show router bgp routes evpn (and filters)
- show service evpn-mpls [<TEP ip-address>]
- show service id bgp-evpn
- show service id evpn-mpls (and modifiers)
- show service id fdb pbb (and modifiers)
- show service system bgp-evpn
- show service system bgp-evpn ethernet-segment (and modifiers)
- debug router bgp update (in classic CLI)
- log-id "99"

In addition, the following **tools dump** commands also discussed in EVPN for MPLS Tunnels can help too:

- tools dump service evpn usage
- tools dump service system bgp-evpn ethernet-segment <name> isid <isid> df (Note: **isid** is used instead of **evi**.)

There are two aspects that are specific to PBB-EVPN and not EVPN:

1. Consumption of virtual BMAC addresses in the system— source BMACs, SAP BMACs, SDP BMACs, and ES BMACs are system BMACs that use FDB space but are not shown in the FDB together with the rest of the learned MAC addresses. The following command provides information about the virtual system MAC addresses consumed in the system.

```
[/]
A:admin@PE-3# tools dump redundancy src-bmac-lsb
Src-bmac-lsb:      3 (00-03) User: B-Vpls - 1 service(s)
Src-bmac-lsb:   8995 (23-23) User: Evpn Mpls

Total Src-bmac-lsbs = 2
```

2. Consumption of MFIBs — when ISIDs are not using the default-multicast list in the B-VPLS context for sending BUM traffic, an MFIB is consumed per ISID. The following command provides information about the consumption of MFIBs per system and per B-VPLS.

```
[/]
A:admin@PE-3# tools dump service vpls-pbb-mfib-stats detail

Service Manager VPLS PBB MFIB statistics at 03/05/2021 10:39:28:

Usage per Service
   ServiceId    MFIB User      Count
  ------------+--------------+-------
   1000         Evpn               1
  ------------+--------------+-------
                       Total       1

MMRP
  Current Usage     :     0
  System Limit      :  8191 Full, 40959 ESOnly
  Per Service Limit :  2048 Full,  8192 ESOnly

SPB
  Current Usage     :     0
  System Limit      :  8191
  Per Service Limit :  8191
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Evpn
  Current Usage     :     1
  System Limit      : 40959
  Per Service Limit :  8191
```

Conclusion

In addition to a full RFC 7432 EVPN-MPLS implementation, SR OS supports PBB-EVPN as per RFC 7623 for large Layer 2 deployments, including single-active and all-active multi-homing. This example has shown how to configure and operate a PBB-EVPN network focusing on the specific aspects of PBB-EVPN compared to EVPN-MPLS.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

EVPN for VXLAN Tunnels (Layer 2)

This chapter provides information about Ethernet Virtual Private Network (EVPN) for Virtual eXtensible Local Area Network (VXLAN) tunnels in VPLS services.

Topics in this chapter include:

Applicability

This chapter is applicable to SR OS and was initially written for SR OS Release 12.0.R4. The MD-CLI in the current edition is based on SR OS Release 21.2.R1. Ethernet Virtual Private Network (EVPN) is a control plane technology and does not have line card hardware dependencies.

Overview

SR OS supports the EVPN control plane with Virtual eXtensible Local Area Network (VXLAN) data plane in VPLS services.

EVPN (RFC 7432) is an IETF technology that uses a dedicated BGP address family which allows VPLS services to be operated in a similar way to IP-VPNs, where the MAC addresses, IP addresses, and the information to set up the flooding tree are distributed by BGP. EVPN can be used as the control plane for different data plane encapsulations, such as VXLAN and MPLS.

VXLAN (RFC 7348) is an overlay IP tunneling technology used to carry Ethernet traffic over any IP network, and it is becoming the de facto standard for overlay data centers and networks. Compared to other IP overlay tunneling technologies, such as GRE, VXLAN supports multi-tenancy and multi-pathing:

- A tenant identifier, the VXLAN Network Identifier (VNI), is encoded in the VXLAN header and allows each tenant to have an isolated Layer 2 domain.
- VXLAN supports multi-pathing scalability through ECMP. VXLAN uses the outer source UDP port as an entropy field that can be used by the core IP routers to balance the load across different paths.

In SR OS, EVPN and VXLAN can be enabled in VPLS or R-VPLS services. In this chapter, EVPN-VXLAN services will refer to VPLS or R-VPLS services with EVPN and VXLAN enabled. These services can terminate/originate VXLAN tunnels and may have SAPs and/or SDP bindings at the same time. Some other SR OS implementation-specific considerations are the following:

- VXLAN is only supported on network or hybrid ports on Ethernet/LAG/POS/APS interfaces.
- VXLAN packets are originated/terminated with the system IPv4 address, in other words, a system originating VXLAN packets will use the system IP address as source outer IPv4 address and systems will only process VXLAN packets if their destination outer IPv4 address matches its own system IP address.
- Data plane MAC learning is not supported over VXLAN bindings. Only the control plane (EVPN) will be used for populating the FDB with MAC addresses associated with VXLAN bindings.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- EVPN provides support for the following features that are described in this chapter:
  - The BGP advertisement of the MAC addresses learned on SAPs, SDP-bindings and conditional static MACs to the remote BGP peers. The advertisement of MAC addresses in BGP can optionally be disabled.
  - The optional advertisement of an unknown MAC route, that allows the remote EVPN PEs or Network Virtualization Edge devices (NVEs) to suppress the unknown unicast flooding and send any unknown unicast frame to the owner of the unknown MAC route.
  - Ingress replication of Broadcast, Unknown unicast, and Multicast (BUM) packets over VXLAN.
  - A proxy-ARP table per service populated by the MAC-IP pairs received in BGP MAC advertisements. When an ARP request is received on a SAP or SDP-binding, the system will perform a lookup on this table and will reply to the ARP request if the lookup yields a valid result.
  - MAC mobility and static-MAC protection as described in RFC 7432, as well as MAC duplication detection.
- Multi-homing redundancy for SAPs and SDP-bindings in EVPN-VXLAN services is supported through BGP Multi-homing (L2VPN BGP address family). Only one BGP-MH site is supported in an EVPN-VXLAN service.

One of the main applications for EVPN-VXLAN services in SR OS is the Data Center Gateway (DC GW) function. In such an application, EVPN and VXLAN are expected to be used within the data center and VPLS SDP-bindings or SAPs are expected to be used for the connectivity to the WAN. When the system is used as a DC GW, a VPLS service is configured per Layer 2 domain that has to be extended to the WAN. In those VPLS services, BGP EVPN automatically sets up the VXLAN auto-bindings that connect the DC GW to the data center Network Virtual Edge devices (NVEs). The WAN connectivity is based on regular VPLS constructs where SAPs (null, dot1q, and QinQ), spoke-SDPs (FEC type 128 and 129, BGP-VPLS), and mesh-SDPs are supported. B-VPLS or I-VPLS services are not supported.

Although the DC GW application is one of the most common uses for this feature, this chapter focuses on the configuration and operation of EVPN-VXLAN for Layer 2 services in general, and its integration with regular VPLS services in MPLS networks.

Configuration

This section describes the configuration of EVPN-VXLAN on SR OS as well as the available troubleshooting and show commands. This example focuses on the following configuration aspects:

- Enabling EVPN and VXLAN in a VPLS service, including the use of BGP-EVPN, BGP Auto-discovery (BGP-AD), and BGP-Multi-homing (BGP-MH) in the same VPLS instance.
- Scaling BGP-MH resiliency with the use of operational groups (oper-groups).
- Use of proxy-ARP in EVPN-VXLAN services
- MAC mobility, MAC duplication, and MAC protection in EVPN-VXLAN services.

The configuration will be shown for PE-1, PE-2, and PE-3 only; the PEs in Overlay-Network-2 (Figure 105: EVPN-VXLAN example topology) have an equivalent configuration.

Enabling EVPN-VXLAN in a VPLS service

Figure 105: EVPN-VXLAN example topology shows the topology used in this example.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 105: EVPN-VXLAN example topology

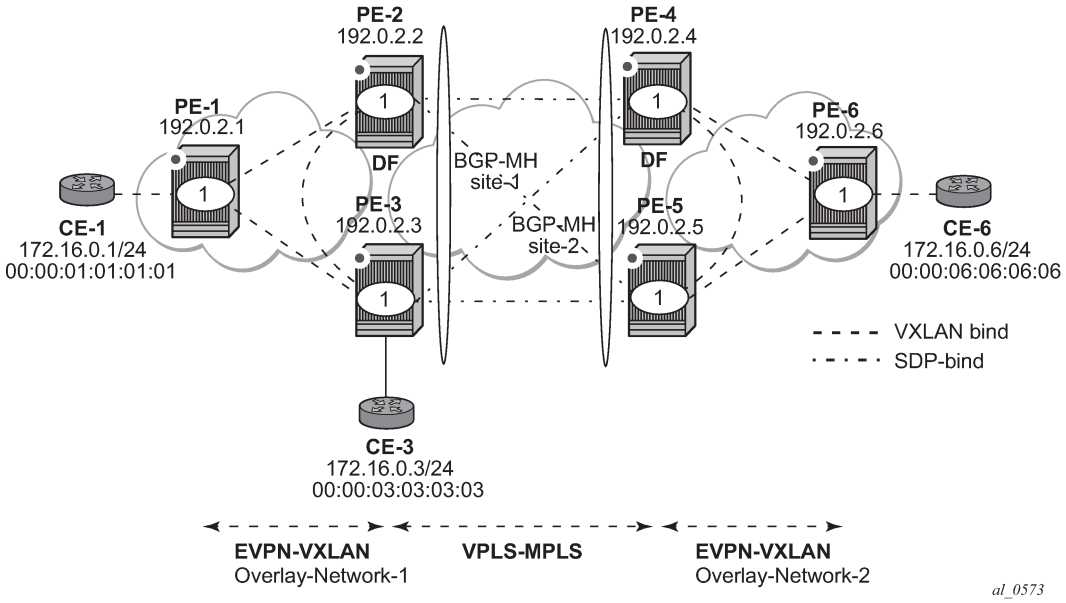

The example topology shows two overlay (VXLAN) networks interconnected by an MPLS network:

- PE-1, PE-2, and PE-3 are part of Overlay-Network-1
- PE-4, PE-5, and PE-6 are part of Overlay-Network-2

CE-1, CE-3, and CE-6 belong to the same IP subnet, therefore, Layer 2 connectivity must be provided to them.

The example topology can illustrate a Data Center Interconnect (DCI) example, where Overlay-Network-1 and Overlay-Network-2 are two data centers interconnected through an MPLS WAN. In this application, CE-1, CE-3, and CE-6 simulate virtual machines or appliances, PE-2/3/4/5 act as DC GWs and PE-1/6 as NVEs (or virtual PEs running on compute infrastructure).

The following protocols and objects are configured beforehand:

- The ports interconnecting the six PEs in Figure 105: EVPN-VXLAN example topology are configured as network ports (or hybrid) and have router network interfaces defined on them. Only the ports connected to the CEs are configured as access ports.
- The six PEs shown in the Figure 105: EVPN-VXLAN example topology are running IS-IS for the global routing table with the four core PEs interconnected using IS-IS Level-2 point-to-point interfaces and each overlay network is using IS-IS Level-1 point-to-point interfaces.
- LDP is used as the MPLS protocol to signal transport tunnel labels among PE-2, PE-3, PE-4, and PE-5. There is no LDP running in the two overlay networks.
- The network port MTU (in all the ports sending/receiving VXLAN packets) must be at least 50 bytes (54 if dot1q encapsulation is used) greater than the service MTU in order to accommodate the size of the VXLAN header.

Once the IGP infrastructure and LDP are enabled in the core, BGP has to be configured. In this example, two BGP families have to be enabled: EVPN within each overlay network for the exchange of MAC/IP

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

addresses and setting up the flooding domains, and L2-VPN for the use of BGP-MH and BGP-AD in the VPLS-MPLS network.

As an example, the following CLI output shows the relevant BGP configuration of PE-1, which only needs the EVPN family. PE-6 would have a similar BGP configuration. The use of route reflectors (RRs) in this type of scenarios is common. Although this example does not use RRs, an EVPN RR could have been used in Overlay-Network-1 and Overlay-Network-2 and an L2-VPN RR could have been used in the core VPLS-MPLS network.

```
# on PE-1:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            rapid-update {
                evpn true
            }
            group "DC" {
                peer-as 64500
                family {
                    evpn true
                }
            }
            neighbor "192.0.2.2" {
                group "DC"
            }
            neighbor "192.0.2.3" {
                group "DC"
            }
        }
```

The BGP configuration on PE-2 is as follows:

```
# on PE-2:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            rapid-update {
                l2-vpn true
                evpn true
            }
            group "DC" {
                peer-as 64500
                family {
                    l2-vpn true
                    evpn true
                }
            }
            group "WAN" {
                peer-as 64500
                family {
                    l2-vpn true
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
            neighbor "192.0.2.1" {
                group "DC"
            }
            neighbor "192.0.2.3" {
                group "DC"
            }
            neighbor "192.0.2.4" {
                group "WAN"
            }
            neighbor "192.0.2.5" {
                group "WAN"
            }
```

The BGP configuration on PE-3 is as follows:

```
# on PE-3:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            rapid-update {
                l2-vpn true
                evpn true
            }
            group "DC" {
                peer-as 64500
                family {
                    l2-vpn true
                    evpn true
                }
            }
            group "WAN" {
                peer-as 64500
                family {
                    l2-vpn true
                }
            }
            neighbor "192.0.2.1" {
                group "DC"
            }
            neighbor "192.0.2.2" {
                group "DC"
            }
            neighbor "192.0.2.4" {
                group "WAN"
            }
            neighbor "192.0.2.5" {
                group "WAN"
            }
```

The BGP configuration on PE-4 and PE-5 is equivalent.

Figure 106: BGP adjacencies and enabled families shows the BGP peering sessions among the PEs and the enabled BGP families. PE-1 will only establish an EVPN peering session with its peers (only the EVPN family is enabled on PE-1), even though PE-2 and PE-3 have EVPN and L2-VPN families configured.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 106: BGP adjacencies and enabled families

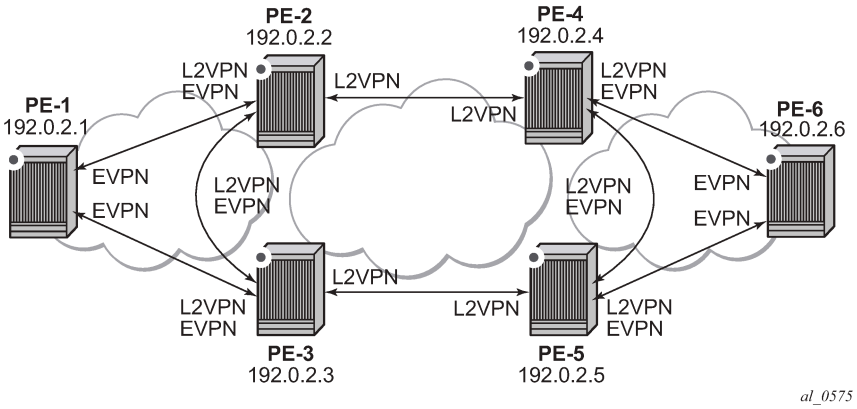

Once the network infrastructure is running properly, the actual service configuration can be carried out. The following CLI outputs show the configuration of VPLS 1 in PE-1, PE-2, and PE-3 as per the topology illustrated in Figure 105: EVPN-VXLAN example topology.

VPLS 1 in those three PEs are interconnected using VXLAN bindings, whereas PE-2 and PE-3 are connected to the remote PEs by means of BGP-AD SDP-bindings. Although BGP-AD SDP-bindings are used in this example for the connectivity of the EVPN-VXLAN PEs to a regular VPLS network, SAPs, BGP-VPLS spoke-SDPs, manual spoke-SDPs, or mesh-SDPs could have been used instead.

VPLS 1 is configured on PE-1, as follows:

```
# on PE-1:
configure {
    service {
        vpls "VPLS1" {
            admin-state enable
            service-id 1
            customer "1"
            vxlan {
                instance 1 {
                    vni 1
                }
            }
            bgp 1 {
                route-distinguisher "192.0.2.1:1"
                route-target {
                    export "target:64500:12"
                    import "target:64500:12"
                }
            }
            bgp-evpn {
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
            sap 1/2/1:1 {
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

EVPN-VXLAN is enabled by the configuration of a valid VXLAN Network Identifier (VNI) and the **bgp-evpn>vxlan>admin-state enable** command. These two commands, along with the required BGP Route Distinguisher (RD) and Route Target (RT) information, are the minimum mandatory attributes:

- The VNI is a 24-bit identifier with valid values in the [1..16777215] range. This defines the VNI that SR OS will use in the EVPN routes generated for the VPLS service, and therefore the VNI that the system expects to see in the VXLAN packets destined to that particular VPLS service. The configured VNI determines the VNI that has to be received in the packets for the VPLS service, but not the VNI that will be sent in VXLAN packets to remote PEs for the service. In other words, in this example, VPLS 1 is configured with VNI=1 in all the PEs; however, each PE could have used a different VNI. The VNI is a system-wide significant value and two VPLS services cannot be configured with the same VNI.
- The **bgp-evpn>vxlan>admin-state enable** command enables the use of EVPN for VXLAN. It requires the previous configuration of the VNI, RD, and RT. As soon as this command is executed, EVPN will advertise an inclusive multicast route to all the BGP EVPN peers (regardless of the existing SAP/SDP-binding operational status). The exchange of inclusive multicast routes allows the establishment of the VXLAN bindings among the PEs.

Upon the reception of the EVPN inclusive multicast routes from PE-2 and PE-3, PE-1 will automatically set up its VXLAN bindings for VPLS 1. A VXLAN binding is represented by an (egress VTEP, egress VNI) pair, where VTEP is a VXLAN Termination End Point. This can be shown with the following show commands on PE-1:

```
[/]
A:admin@PE-1# show service vxlan

===============================================================================
VXLAN Tunnel Endpoints (VTEPs)
===============================================================================
VTEP Address                            VXLAN Dest        ES Dest
-------------------------------------------------------------------------------
192.0.2.2                               1                 0
192.0.2.3                               1                 0
-------------------------------------------------------------------------------
Number of VTEPs: 2
-------------------------------------------------------------------------------
===============================================================================
```

```
[/]
A:admin@PE-1# show service id 1 vxlan instance 1 destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
Instance    VTEP Address                        Egress VNI  EvpnStatic Num
 Mcast       Oper State                          L2 PBR     SupBcasDom MACs
-------------------------------------------------------------------------------
1           192.0.2.2                           1           evpn       1
 BUM         Up                                  No          No
1           192.0.2.3                           1           evpn       1
 BUM         Up                                  No          No
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 2
-------------------------------------------------------------------------------
===============================================================================
---snip---
```

To actually see this output, the VPLS service needs to be configured on all PEs, with import and export policy "vsi-policy-1" defined on the core PEs; see further. As can be seen in the CLI output, PE-1 has two

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

VXLAN bindings: one to PE-2 and one to PE-3. Both use egress VNI=1 (the actual VNI used in its egress VXLAN packets) and both are part of the flooding multicast list (BUM) for VPLS 1 and are up. There is no layer 2 Policy-Based Routing (L2 PBR).

- The **Mcast= BUM** entry is set when the proper inclusive multicast route is received from the remote VTEP. The VXLAN binding will be used to flood BUM packets.
- The **Oper State** is based on the existence of the VTEP in the global routing table.

The VPLS 1 configuration of PE-2 and PE-3 is as follows:

```
# on PE-2:
configure {
    service {
        pw-template "PW1" {
            pw-template-id 1
        }
        vpls "VPLS1" {
            admin-state enable
            service-id 1
            customer "1"
            vxlan {
                instance 1 {
                    vni 1
                }
            }
            bgp 1 {
                route-distinguisher "192.0.2.2:1"
                vsi-import ["vsi-policy-1"]
                vsi-export ["vsi-policy-1"]
                pw-template-binding "PW1" {
                    split-horizon-group "CORE"
                }
            }
             bgp-ad {
                admin-state enable
                vpls-id "64500:1"
            }
            bgp-evpn {
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
            bgp-mh-site "site-1" {
                admin-state enable
                id 1
                shg-name "CORE"
            }
        }
```

```
# on PE-3:
configure {
    service {
        pw-template "PW1" {
            pw-template-id 1
        }
        vpls "VPLS1" {
            admin-state enable
            service-id 1
            customer "1"
            vxlan {
                instance 1 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                vni 1
            }
        }
        bgp 1 {
            route-distinguisher "192.0.2.3:1"
            vsi-import ["vsi-policy-1"]
            vsi-export ["vsi-policy-1"]
            pw-template-binding "PW1" {
                split-horizon-group "CORE"
            }
        }
        bgp-ad {
            admin-state enable
            vpls-id "64500:1"
        }
        bgp-evpn {
            vxlan 1 {
                admin-state enable
                vxlan-instance 1
            }
        }
        bgp-mh-site "site-1" {
            admin-state enable
            id 1
            shg-name "CORE"
            }
        }
        sap 1/2/1:1 {
        }
```

In addition to the VNI and **bgp-evpn>vxlan>admin-state enable** commands for enabling EVPN-VXLAN in VPLS 1, PE-2 and PE-3 require the configuration of BGP-AD for the discovery and establishment of FEC129 spoke-SDPs to the remote PEs in the core, as well as BGP-MH for redundancy. As described in Figure 105: EVPN-VXLAN example topology, there are two BGP-MH sites defined in the network: site-1 is used on PE-2/PE-3 and site-2 is used on PE-4/PE-5. Only one of the two gateway PEs in each overlay network will be the Designated Forwarder (DF) for VPLS 1, and only the DF will send/receive traffic for VPLS 1 in the overlay network. The following considerations must be taken into account when configuring the connectivity of EVPN-VXLAN services to regular VPLS objects:

- As discussed, in this example, BGP-AD spoke-SDPs are used, but SAPs, BGP-VPLS spoke-SDPs, manual spoke-SDPs, or mesh-SDPs are also supported.
- In this example, BGP-AD spoke-SDPs are auto-instantiated using **pw-template-binding "PW1" split-horizon-group “CORE”**.
  - This requires the creation of the pw-template "PW1" (**config>service>pw-template "PW1"**).
- The split-horizon group CORE is added to the BGP-MH site "site-1”. This statement will ensure that all the spoke-SDPs automatically established to the remote PEs are part of the BGP-MH site.
- Although the route targets for the overlay network and the VPLS-MPLS network can have the same value for the same VPLS service, they are usually different. This example assumes the use of RT-DC-1 in Overlay-Network-1 and RT-WAN-1 in the VPLS-MPLS core for VPLS 1. The "vsi-policy-1” allows the system to export and import the right RTs for VPLS 1 on the core PEs:

```
# on PE-2 and PE-3:
configure {
    policy-options {
        community "RT-DC-1" {
            member "target:64500:12" { }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        community "RT-WAN-1" {
            member "target:64500:11" { }
        }
        policy-statement "vsi-policy-1" {
            entry 10 {          # to import all the EVPN routes with RT-DC-1
                from {
                    family [evpn]
                    community {
                        name "RT-DC-1"
                    }
                }
                action {
                    action-type accept
                }
            }
            entry 20 {           # to import all the BGP-AD/MH routes from the WAN
                from {
                    family [l2-vpn]
                    community {
                        name "RT-WAN-1"
                    }
                }
                action {
                    action-type accept
                }
            }
            entry 30 {          # to export all the EVPN routes with "RT-DC-1"
                from {
                    family [evpn]
                }
                action {
                    action-type accept
                    community {
                        add ["RT-DC-1"]
                    }
                }
            }
            entry 40 {           # to export all the BGP-AD/MH routes with "RT-WAN-1"
                from {
                    family [l2-vpn]
                }
                action {
                    action-type accept
                    community {
                        add ["RT-WAN-1"]
                    }
                }
            }
            default-action {
                action-type reject
            }
        }
```

Once PE-2 and PE-3 are configured as shown, they will set up the spoke-SDPs and will run the DF election algorithm to determine the operational status of those spoke-SDPs. See chapters LDP VPLS Using BGP Auto-Discovery and BGP Multi-Homing for VPLS Networks for more information about the use of BGP-AD and BGP-MH.

In the configuration for VPLS 1, both gateway PEs, PE-2 and PE-3, will attempt to establish two parallel Layer 2 paths between each other (a BGP-AD spoke-SDP and an EVPN VXLAN binding). Because that would create a Layer 2 loop, the SR OS implementation gives priority to the EVPN path and only the VXLAN binding will be active. In other words, when a VXLAN (egress VTEP, VNI) and a spoke-SDP are

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

attempted to be set up to the same far-end IP address at the same time, the VXLAN path will prevail and the spoke-SDP will be kept down. The spoke-SDP will only be brought up if the VXLAN (egress VTEP, VNI) goes down.

This behavior can be easily observed in this setup by using the following **show** commands. In PE-2, the spoke-SDP to far-end PE-3 will be down with an **EvpnRouteConflict** Flag. The (egress VTEP, VNI) = (192.0.2.3, 1) VXLAN bind will be up.

```
[/]
A:admin@PE-2# show service id 1 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 1                   Vpn Id            : 0
Service Type      : VPLS
---snip---

Admin State       : Up                  Oper State        : Up
MTU               : 1514
SAP Count         : 0                   SDP Bind Count    : 3
---snip---

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                             Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sdp:32765:4294967293 SB(192.0.2.5)      BgpAd        0       8978    Up   Up
sdp:32766:4294967294 SB(192.0.2.4)      BgpAd        0       8978    Up   Up
sdp:32767:4294967295 SB(192.0.2.3)      BgpAd        0       8978    Up   Down
===============================================================================
```

```
[/]
A:admin@PE-2# show service id 1 sdp 32767 detail | match 'Flag' post-lines 1
Flags              : PWPeerFaultStatusBits
                     EvpnRouteConflict
```

```
[/]
A:admin@PE-2# show service id 1 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
Instance    VTEP Address                          Egress VNI  EvpnStatic Num
 Mcast       Oper State                            L2 PBR     SupBcasDom MACs
-------------------------------------------------------------------------------
1           192.0.2.1                             1           evpn       1
 BUM         Up                                    No          No
1           192.0.2.3                             1           evpn       1
 BUM         Up                                    No          No
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 2
-------------------------------------------------------------------------------
===============================================================================
---snip---
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

At the non-DF, PE-3, all the spoke-SDPs will be down due to BGP-MH, but for the SDP 32767 toward PE-2, an additional reason is an EVPN route conflict:

```
[/]
A:admin@PE-3# show service id 1 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 1                   Vpn Id            : 0
Service Type      : VPLS
---snip---

Admin State       : Up                  Oper State        : Up
MTU               : 1514
SAP Count         : 1                   SDP Bind Count    : 3
---snip---

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                             Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/2/1:1                            q-tag        1518    1518    Up   Up
sdp:32765:4294967293 SB(192.0.2.5)     BgpAd        0       8978    Up   Down
sdp:32766:4294967294 SB(192.0.2.4)     BgpAd        0       8978    Up   Down
sdp:32767:4294967295 SB(192.0.2.2)     BgpAd        0       8978    Up   Down
===============================================================================
```

```
[/]
A:admin@PE-3# show service id 1 sdp 32767 detail | match 'Flag' post-lines 2
Flags              : StandbyForMHProtocol
                     PWPeerFaultStatusBits
                     EvpnRouteConflict
```

MAC learning and unknown-mac

Once the VPLS service (VPLS 1) is configured, the network allows the CEs to exchange unicast and BUM traffic over the overlay and VPLS-MPLS service infrastructure. BUM traffic sent by CE-1 will be ingress-replicated by PE-1 to PE-2 and PE-3, and propagated by PE-2 (the DF) to the remote network. From this point on, MAC addresses will be learned on active SAPs and spoke-SDPs and advertised in EVPN MAC routes. No data plane MAC learning is carried out on VXLAN bindings. MACs associated with (egress VTEP, VNI) bindings will always be learned through EVPN.

The following CLI output shows the reception of an EVPN MAC route on PE-1 and how the (CE-3) MAC address appears in the FDB for VPLS 1.

```
# on PE-1:
11 2021/02/10 15:48:52.094 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 88
    Flag: 0x90 Type: 14 Len: 44 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.3
        Type: EVPN-MAC Len: 33 RD: 192.0.2.3:1 ESI: ESI-0, tag: 0, mac len: 48
                       mac: 00:00:03:03:03:03, IP len: 0, IP: NULL, label1: 1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x80 Type: 4 Len: 4 MED: 0
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:12
        bgp-tunnel-encap:VXLAN
"
```

```
[/]
A:admin@PE-1# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
1          00:00:01:01:01:01 sap:1/2/1:1             L/120    02/10/21 15:47:35
1          00:00:03:03:03:03 vxlan-1:                Evpn     02/10/21 15:48:52
                             192.0.2.3:1
1          00:00:06:06:06:06 vxlan-1:                Evpn     02/10/21 15:52:31
                             192.0.2.2:1
-------------------------------------------------------------------------------
No. of MAC Entries: 3
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

When a frame destined to 00:00:03:03:03:03 enters SAP 1/2/1:1, it is encapsulated into a VXLAN packet with outer destination IP 192.0.2.3 and VNI 1, and sent on the wire.

In virtualized data center networks where all the MACs are known beforehand (all the virtual machine and appliance MACs are distributed by EVPN before any traffic flows), unknown MAC addresses are always outside the data center. If that is the case, the DC GWs can make use of the **unknown-mac true** so that the DC NVEs supporting the concept of this route send the unknown unicast traffic only to the DC GW. This minimizes the flooding within the data center, as described in draft-ietf-bess-dci-evpn-overlay.

In this example, the unknown MAC route is configured in the gateway PEs (in Overlay-Network-1: PE-2 and PE-3) in the following way:

```
# on PE-2, PE-3:
configure {
    service {
        vpls "VPLS1" {
            bgp-evpn {
                routes {
                    mac-ip {
                        unknown-mac true
                    }
                }
            }
```

```
47 2021/02/10 16:00:36.068 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 88
    Flag: 0x90 Type: 14 Len: 44 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        Type: EVPN-MAC Len: 33 RD: 192.0.2.2:1 ESI: ESI-0, tag: 0, mac len: 48
                       mac: 00:00:00:00:00:00, IP len: 0, IP: NULL, label1: 1
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x80 Type: 4 Len: 4 MED: 0
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:12
        bgp-tunnel-encap:VXLAN
"
```

Note that:

- Although SR OS can generate the unknown MAC route, it will never honor it and normal flooding applies when an unknown unicast packet arrives at an ingress SAP/SDP-binding.
- When **unknown-mac true** is configured, it will only be generated when: a) no BGP-MH site is configured within the same VPLS service or b) a site is configured and the site is DF in the PE. If the site becomes a non-DF site, the unknown MAC route will be withdrawn.
- If the **unknown-mac true** is used in the DC GW and all the NVEs in the DC understand it, the advertisement of MAC addresses can be disabled with the **routes>mac-ip>advertise false** command. If so, SR OS will only advertise the unknown MAC route.

```
# on DC GWs PE-2 and PE-3:
configure {
    service {
        vpls "VPLS1" {
            bgp-evpn {
                routes {
                    mac-ip {
                        advertise false
                        unknown-mac true
                    }
                }
            }
```

Scaling BGP-MH resiliency with the use of operational groups

In Figure 105: EVPN-VXLAN example topology, VPLS 1 in PE-2 and PE-3 is configured with a BGP-MH site that controls which of the two PEs forwards the traffic to the remote PEs (in this case, PE-2 is the DF and the GW responsible for forwarding packets to the remote PEs).

When new VPLS services are required in PE-2 and PE-3, the same BGP-MH configuration can be used. However, if the number of VPLS services grows significantly, the use of individual BGP-MH sites per service will not scale. Because all the services in these two PEs share the same physical topology, the use of operational groups can provide a simple and scalable way of providing resiliency to as many services as the user needs (up to the maximum number of VPLS services per system).

The way operational groups can be used to scale this type of deployments is the following (using the network topology in Figure 105: EVPN-VXLAN example topology and focusing on Overlay-Network-1):

- A control-VPLS service is defined in PE-2 and PE-3. For instance, VPLS 1.
  - This service is configured with a BGP-MH site in both PEs.
  - An oper-group “control-vpls-1” is created and associated with the pw-template-binding 1 in VPLS 1.
- Data VPLS services are defined in both PEs. For instance: VPLS 2, VPLS 3,... VPLS 999.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

  - In all these services, the pw-template-binding is configured with **monitor-oper-group "control-vpls-1".**
  - The status of the spoke-SDPs in the data VPLS services depends on the status of the operational group. If there is a DF switchover in VPLS 1 and VPLS 1 spoke-SDPs go down on PE-2, all the spoke-SDPs in all the data VPLS services controlled by "control-vpls-1" in PE-2 will go down too. In the same way, the spoke-SDPs in PE-3 will come up.
- To allow per-service load balancing, a second control-VPLS service with a different BGP-MH site should be configured.
  - For instance, VPLS 1 may have PE-2 as the DF and VPLS 1000 may be a second control-VPLS service with PE-3 as the DF.
  - Each control-VPLS would control a group of data VPLS services based on the definition and association of a second operational group.

The following example shows the modification of VPLS 1 as the control-VPLS and the configuration of VPLS 2 as a data-VPLS on PE-2. VPLS 1 controls the VPLS 2 spoke-SDP status.

```
# on PE-2:
configure {
    service {
        oper-group "control-vpls-1" {
        }
        vpls "VPLS1" {
            description "control-VPLS"
            bgp 1 {
                pw-template-binding "PW1" {
                    oper-group "control-vpls-1"
                }
            }
        }
        vpls "VPLS2" {
            admin-state enable
            description "data-VPLS"
            service-id 2
            customer "1"
            vxlan {
                instance 1 {
                    vni 2
                }
            }
            bgp 1 {
                route-distinguisher "192.0.2.2:2"
                vsi-import ["vsi-policy-2"]
                vsi-export ["vsi-policy-2"]
                pw-template-binding "PW1" {
                    monitor-oper-group "control-vpls-1"
                }
            }
            bgp-ad {
                admin-state enable
                vpls-id "64500:2"
            }
            bgp-evpn {
                routes {
                    mac-ip {
                        unknown-mac true
                    }
                }
                vxlan 1 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    admin-state enable
                    vxlan-instance 1
                }
            }
        }
```

Use of proxy-ARP in EVPN-VXLAN services

EVPN-VXLAN services support proxy-ARP functionality that is enabled by the **proxy-arp admin-state** command. By default, proxy-ARP is disabled. When proxy-ARP is enabled, the following applies:

- MAC and IP addresses contained in the received valid EVPN MAC routes are populated in the proxy-ARP table.
- ARP-request messages received on SAPs and SDP-bindings are intercepted and the target IP address is looked up. If the IP address is found, an ARP reply will be issued based on the information found in the proxy-ARP table, otherwise the ARP request would be flooded in the VPLS service (except for the source SAP/SDP binding).
- ARP-reply messages received on SAPs and SDP-bindings are also intercepted and sent to the CPM. These ARP-reply messages are re-injected in the data plane and forwarded based on the FDB information to the destination MAC address. If the destination MAC address is not in the FDB, the ARP-reply message will be flooded in the VPLS service (except for the source SAP/SDP binding).

The following CLI output shows the proxy-ARP configuration in PE-3 and a received valid MAC route that includes the MAC address 00:00:01:01:01:01 and IP address 172.16.0.1 of CE-1. This MAC-IP pair is installed in the proxy-ARP table for VPLS 1.

```
# on PE-3:
configure {
    service {
        vpls "VPLS1" {
            proxy-arp {
                admin-state enable
            }
```

```
120 2021/02/10 16:12:53.542 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.1
"Peer 1: 192.0.2.1: UPDATE
Peer 1: 192.0.2.1 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 127
    Flag: 0x90 Type: 14 Len: 83 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.1
        Type: EVPN-MAC Len: 37 RD: 192.0.2.1:1 ESI: ESI-0, tag: 0, mac len: 48
                       mac: 00:00:01:01:01:01, IP len: 4, IP: 172.16.0.1, label1: 1
        Type: EVPN-MAC Len: 33 RD: 192.0.2.1:1 ESI: ESI-0, tag: 0, mac len: 48
                       mac: 00:00:01:01:01:01, IP len: 0, IP: NULL, label1: 1
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x80 Type: 4 Len: 4 MED: 0
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:12
        bgp-tunnel-encap:VXLAN
"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

This MAC-IP pair is installed in the proxy-ARP table for VPLS 1 on PE-3, as follows:

```
[/]
A:admin@PE-3# show service id 1 proxy-arp detail
-------------------------------------------------------------------------------
Proxy Arp
-------------------------------------------------------------------------------
Admin State       : enabled
Dyn Populate      : disabled
Age Time          : disabled            Send Refresh      : disabled
Table Size        : 250                 Total             : 1
Static Count      : 0                   EVPN Count        : 1
Dynamic Count     : 0                   Duplicate Count   : 0

Dup Detect
-------------------------------------------------------------------------------
Detect Window     : 3 mins              Num Moves         : 5
Hold down         : 9 mins
Anti Spoof MAC    : None

EVPN
-------------------------------------------------------------------------------
Garp Flood        : enabled             Req Flood         : enabled
Static Black Hole : disabled
EVPN Route Tag    : 0
-------------------------------------------------------------------------------

===============================================================================
VPLS Proxy Arp Entries
===============================================================================
IP Address          Mac Address         Type      Status    Last Update
-------------------------------------------------------------------------------
172.16.0.1          00:00:01:01:01:01   evpn      active    02/10/2021 16:12:54
-------------------------------------------------------------------------------
Number of entries : 1
===============================================================================
```

SR OS does not include a host IP address in any EVPN MAC advertisement for a MAC learned on a SAP or SDP-binding. Host IP addresses are only included in the EVPN MAC advertisements corresponding to R-VPLS IP interfaces. When deployed as DC GW in a Nuage architecture, the Nuage Networks Virtual Services Controller (VSC) or Virtual Services Gateway (VSG) will send virtual machine and host MAC/IP pairs in EVPN MAC routes. See the Nokia Nuage documentation for more information about the Nuage DC architecture. The 7x50 DC GW will populate the proxy-ARP tables with those MAC/IP pairs.

In the preceding CLI excerpt, assume that PE-1 is replaced by a Nuage VSC that sends the pair <172.16.0.1, 00:00:01:01:01:01> in an EVPN MAC route. PE-3 receives the advertisement and adds the entry to its proxy-ARP table for VPLS 1.

The proxy-ARP feature was significantly improved in SR OS Release 13.0; see the EVPN for MPLS Tunnels chapter.

MAC mobility, MAC duplication, and MAC protection in EVPN

MAC mobility, duplication and protection are fully supported as specified in RFC 7432. Figure 107: EVPN MAC mobility illustrates the concept of mobility (Virtual Machine VM-1 moves from PE-1 to PE-3).

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 107: EVPN MAC mobility

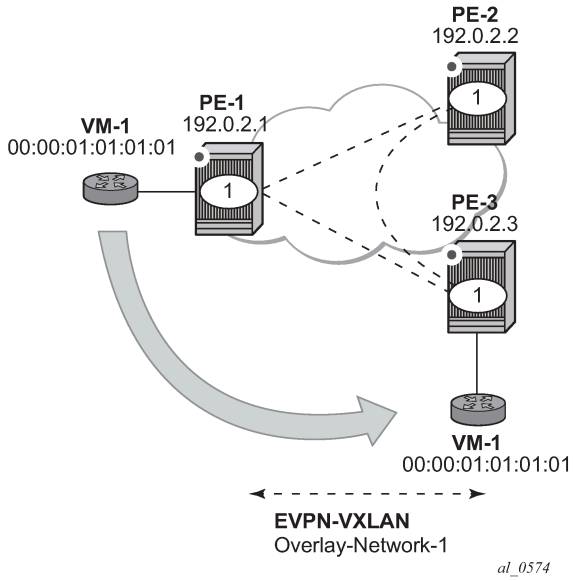

MAC mobility is handled in EVPN by the use of sequence numbers in the MAC routes. When 00:00:01:01:01:01 moves from PE-1 to PE-3, SR OS will gracefully handle it in this way:

- 00:00:01:01:01:01 moves to PE-3 SAP 1/2/1:1
- PE-3 advertises 00:00:01:01:01:01 using a higher sequence number (the first time a MAC is advertised, EVPN uses sequence number 0).
- PE-2 at this point has two valid MAC routes for 00:00:01:01:01:01. It picks up the one coming from PE-3 because the sequence number is higher.
- PE-1 receives the MAC route, and because the sequence number is higher than the one for its own route, it updates the FDB and withdraws its own MAC route.

However, if MAC 00:00:01:01:01:01 is constantly learned on the PE-1 and PE-3 SAPs, the preceding process causes an endless exchange of MAC route advertisements and withdraws that has a negative impact on all the PEs in the EVPN network. This issue is known as *MAC duplication* and is originated by a loop at the access or a duplicated MAC address in two hosts of the same service. SR OS solves this issue through the use of the MAC duplication detection feature. MAC duplication is always enabled with the following default settings:

```
[ex:/configure service vpls "VPLS1" bgp-evpn]
A:admin@PE-3# info detail | match 'mac-duplication' post-lines 7
    mac-duplication {
        retry 9
        blackhole false
        detect {
            num-moves 5
            window 3
        }
    }
```

Where:

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

num-moves

Identifies the number of MAC moves in a VPLS service. The counter is incremented when a MAC is locally relearned in the FDB or flushed from the FDB because of the reception of a better remote EVPN route for that MAC. When the threshold is reached for a MAC address, this MAC address is put in hold-down state (this 'hold-down' state is described below). Range: <3..10>. Default value: 5.

window

Identifies the timer within which a MAC is considered duplicate if it reaches the configured num-moves. Range: <1..15> minutes. Default value: 3 minutes.

retry

The timer after which the MAC in hold-down state is automatically flushed and the mac-duplication process starts again. This value is expected to be equal to two times or more than the window. If no retry is configured, this implies that, once MAC duplication is detected, MAC updates for that MAC will be held down until the user intervenes or a network event (that flushes the MAC) occurs. Range: <2..60> minutes. Default value: 9 minutes.

blackhole

If enabled and a duplicate MAC address is detected, the router adds the MAC address to the duplicate MAC list and it programs the MAC in the FDB as a protected MAC associated with a black-hole (with type EvpnD:P and source ID "black-hole")

When a MAC address is considered a duplicate or in the hold-down state, no further BGP advertisements are issued for this MAC and an alarm is triggered (by the first MAC address in hold-down state). The following CLI output shows how PE-3 detects that MAC 00:00:01:01:01:01 is a duplicate (after reaching the **num-moves** in **window**) and the corresponding alarm.

```
# on PE-3:
144 2021/02/10 16:16:44.974 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.1
"Peer 1: 192.0.2.1: UPDATE
Peer 1: 192.0.2.1 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 96
    Flag: 0x90 Type: 14 Len: 44 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.3
        Type: EVPN-MAC Len: 33 RD: 192.0.2.3:1 ESI: ESI-0, tag: 0, mac len: 48
                       mac: 00:00:01:01:01:01, IP len: 0, IP: NULL, label1: 1
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x80 Type: 4 Len: 4 MED: 0
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:12
        bgp-tunnel-encap:VXLAN
        mac-mobility:Seq:5
"
```

Log 99 on PE-3 shows the following message when EVPN has detected a duplicate MAC address in VPLS 1:

```
# on PE-3:
154 2021/02/10 16:18:58.902 UTC MINOR: SVCMGR #2331 Base
"VPLS Service 1 has MAC(s) detected as duplicates by EVPN mac-duplication detection."
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The **show service id bgp-evpn** command shows the MAC duplication settings and the list of duplicate MAC addresses on hold-down.

```
[/]
A:admin@PE-3# show service id 1 bgp-evpn

===============================================================================
BGP EVPN Table
===============================================================================
MAC Advertisement  : Enabled            Unknown MAC Route  : Enabled
CFM MAC Advertise  : Disabled
Creation Origin    : manual
MAC Dup Detn Moves : 3                  MAC Dup Detn Window: 1
MAC Dup Detn Retry : 2                  Number of Dup MACs : 1
MAC Dup Detn BH    : Disabled
IP Route Advert    : Disabled
Sel Mcast Advert   : Disabled

EVI                : n/a
Ing Rep Inc McastAd: Enabled
Accept IVPLS Flush : Disabled

-------------------------------------------------------------------------------
Detected Duplicate MAC Addresses            Time Detected
-------------------------------------------------------------------------------
00:00:01:01:01:01                           02/10/2021 16:18:58
-------------------------------------------------------------------------------
===============================================================================
---snip---
```

SR OS stops sending and processing any BGP MAC advertisement routes for that MAC address until:

- The MAC is flushed because of a local event (SAP/SDP-binding associated with the MAC fails) or the reception of a remote withdraw for the MAC (because of a MAC flush at the remote 7x50).
- The **retry** <*in_minutes*> timer expires, which flushes the MAC and restart the process.

When the last duplicate MAC address is removed from the duplicate list, log 99 on PE-3 will show the following message:

```
155 2021/02/10 16:21:58.885 UTC MINOR: SVCMGR #2332 Base
"VPLS Service 1 no longer has MAC(s) detected as duplicates by EVPN mac-duplication
detection."
```

EVPN also provides a mechanism to protect specific MAC addresses that do not move for which connectivity must be guaranteed. These addresses must be protected in case there is an attempt to dynamically learn them in a different place in the EVPN-VXLAN VPLS service (on the same or different PE).

The protected MAC addresses are configured in SR OS as conditional static MAC addresses. A conditional static MAC address defined in an EVPN-VXLAN VPLS service is advertised by BGP-EVPN as a static address. An example of the configuration of a conditional static MAC address is as follows:

```
# on PE-1:
configure {
    service {
        vpls "VPLS1"
            fdb {
                static-mac {
                    mac 00:00:05:05:05:05 {
                        sap 1/2/1:1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        monitor forward-status
                    }
                }
            }
        }
```

The protected MAC addresses advertised in EVPN are shown in the receiving BGP RIB as Static (MAC mobility extended community with Sequence 0 and sticky bit set) and **EvpnS:P** (Evpn Static: Protected) in the FDB. The advertising PE shows the protected MAC as **CStatic:P** (Conditional Static: Protected) in the FDB:

On the advertising PE:

```
[/]
A:admin@PE-1# show service id 1 fdb mac 00:00:05:05:05:05

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
1          00:00:05:05:05:05 sap:1/2/1:1             CStatic: 02/10/21 16:31:03
                                                     P
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

On the receiving PE:

```
[/]
A:admin@PE-3# show service id 1 fdb mac 00:00:05:05:05:05

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
1          00:00:05:05:05:05 vxlan-1:                EvpnS:P  02/10/21 16:31:03
                             192.0.2.1:1
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

```
[/]
A:admin@PE-3# show router bgp routes evpn mac hunt mac-address 00:00:05:05:05:05
===============================================================================
 BGP Router ID:192.0.2.3       AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN MAC Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Network        : n/a
Nexthop        : 192.0.2.1
From           : 192.0.2.1
Res. Nexthop   : 192.168.13.1
Local Pref.    : 100                    Interface Name : int-PE-3-PE-1
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : 0
AIGP Metric    : None                   IGP Cost       : 10
Connector      : None
Community      : target:64500:12 bgp-tunnel-encap:VXLAN
                 mac-mobility:Seq:0/Static
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.1
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
EVPN type      : MAC
ESI            : ESI-0
Tag            : 0
IP Address     : n/a
Route Dist.    : 192.0.2.1:1
Mac Address    : 00:00:05:05:05:05
MPLS Label1    : VNI 1                  MPLS Label2    : n/a
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
Add Paths Send : Default
Last Modified  : 00h02m32s

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

The following procedures are supported to protect the configured static MAC addresses:

- All the SAP/SDP-bindings are internally configured as MAC protect restrict-protected-src as soon as BGP-EVPN is enabled in the VPLS service.
- Local static MAC addresses or remote EVPN static MAC addresses are considered as protected.
- If a frame with a source MAC address matching one of the protected MAC addresses is received on a different SAP/SDP-binding than the owner of the protected MAC address, the frame is discarded and an alarm triggered. This MAC protection is not performed for frames received on VXLAN bindings.
- The same throttled alarm mechanism used in MAC protect for restrict-protected-src with discard-frame is used here: the offending frames are captured to a list to be polled by the CPM every ~10min.

In this example, PE-3 has 00:00:05:05:05:05 in its FDB as EvpnS. If SAP 1/2/1:1 on PE-3 receives a frame with source MAC address 00:00:05:05:05:05, the frame is discarded and an alarm triggered. The following is logged in log 99 on PE-3:

```
164 2021/02/10 16:44:03.736 UTC MINOR: SVCMGR #2208 Base Slot 1
"Protected MAC 00:00:05:05:05:05 received on SAP 1/2/1:1 in service 1. "
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Debug and show commands

In addition to the previously mentioned **show service id vxlan destinations**, **show service id bgp-evpn** and **show service id fdb detail** commands, the following commands provide valuable information when troubleshooting an EVPN-VXLAN VPLS service.

The **show router bgp routes evpn** command supports filtering by route type as well as many other route fields.

```
[/]
A:admin@PE-1# show router bgp routes evpn ?

 auto-disc             - Display BGP EVPN Auto-Disc Routes
 eth-seg               - Display BGP EVPN Eth-Seg Routes
 incl-mcast            - Display BGP EVPN Inclusive-Mcast Routes
 ip-prefix             - Display BGP EVPN IPv4-Prefix Routes
 ipv6-prefix           - Display BGP EVPN IPv6-Prefix Routes
 mac                   - Display BGP EVPN Mac Routes
 mcast-join-synch      - Display BGP EVPN Mcast Join Sync Routes
 mcast-leave-synch     - Display BGP EVPN Mcast Leave Sync Routes
 smet                  - Display BGP EVPN Smet Routes
 spmsi-ad              - Display BGP EVPN Spmsi AD Routes
```

```
[/]
A:admin@PE-1# show router bgp routes evpn mac ?

 mac [<keyword>] [rd <string>] [next-hop <string>] [mac-address <string>] [community
 <string>] [tag <string>]
  [aspath-regex <string>]

 [hunt-detail] <keyword>
 <keyword>  - (hunt|detail)

    keywords

 [hunt-detail]         - keywords
 aspath-regex          - string '<1..80 characters>'
 community             - <as-number1:comm-val1>|<ext-comm>|, <well-known-comm>,
                         ext-comm -  <type>:{<ip-address:comm-val1>|,
                         <as-number1:comm-val2>|,<as-number2:comm-val1>},
                         as-number1 - [0..65535], comm-val1 - [0..65535],
                         type  - target|origin, ip-address - a.b.c.d,
                         comm-val2 - [0..4294967295], as-number2  - [0..4294967295],
                         well-known-comm - null|no-export|no-export-subconfed|,
                         no-advertise
 mac-address           - string '<0..255 characters>'
 next-hop              - Attribute next-hop for mac
 rd                    - {<ip-addr:comm-val>|, <2byte-asnumber:ext-comm-val>|,
                                               <4byte-asnumber:comm-val>}
 tag                   - Attribute tag for mac
```

```
[/]
A:admin@PE-3# show router bgp routes evpn mac tag 0
===============================================================================
 BGP Router ID:192.0.2.3        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
BGP EVPN MAC Routes
===============================================================================
Flag  Route Dist.         MacAddr           ESI
      Tag                 Mac Mobility      Label1
                          Ip Address
                          NextHop
-------------------------------------------------------------------------------
u*>i  192.0.2.1:1         00:00:05:05:05:05 ESI-0
      0                   Static            VNI 1
                          n/a
                          192.0.2.1

u*>i  192.0.2.1:1         04:09:ff:00:03:3a ESI-0
      0                   Static            VNI 1
                          n/a
                          192.0.2.1

u*>i  192.0.2.2:1         00:00:00:00:00:00 ESI-0
      0                   Seq:0             VNI 1
                          n/a
                          192.0.2.2

-------------------------------------------------------------------------------
Routes : 3
===============================================================================
```

The **tools dump service id vxlan** displays the number of times a service could not add a VXLAN binding or <VTEP, Egress VNI> because of the following limits:

- The per system VTEP limit has been reached
- The per system (egress VTEP, egress VNI) limit has been reached
- The per service (egress VTEP, egress VNI) limit has been reached
- The per system Bind limit: Total bind limit or VXLAN bind limit has been reached.

```
[/]
A:admin@PE-1# tools dump service id 1 vxlan

VTEP, Egress VNI Failure statistics at 02/10/2021 17:03:07:

statistics last cleared at 02/10/2021 10:43:55:

Failures: None
```

```
[/]
A:admin@PE-1# tools dump service id 1 evpn usage

Evpn Tunnel Interface IP Next Hop: N/A
```

Tools dump service evpn usage displays the consumed resources in the system, as follows:

```
[/]
A:admin@PE-1# tools dump service evpn usage

vxlan-evpn-mpls usage statistics at 02/10/2021 17:03:07:

MPLS-TEP                                        :             0
VXLAN-TEP                                       :             2
Total-TEP                                       :      2/ 16383
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Mpls Dests (TEP, Egress Label + ES + ES-BMAC)   :             0
Mpls Etree Leaf Dests                           :             0
Vxlan Dests (TEP, Egress VNI + ES)              :             2
Total-Dest                                      :      2/196607

Sdp Bind +  Evpn Dests                          :      2/245759
ES L2/L3 PBR                                    :      0/ 32767
Evpn Etree Remote BUM Leaf Labels               :             0
```

Conclusion

SR OS supports the EVPN control plane for VXLAN tunnels terminated in VPLS services. VXLAN is an overlay IP tunneling mechanism that is being used in data centers, data center interconnect, and other applications. EVPN is a scalable and flexible control plane that provides control over the MAC addresses being learned and advertised, as well as other mechanisms to optimize Layer 2 services such as proxy-ARP, MAC mobility, MAC duplication detection, and MAC protection. SR OS provides a resilient and scalable EVPN-VXLAN solution for Layer 2 services, including interoperability to existing VPLS networks. This chapter showed all of those functions and how they are configured and operated.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

EVPN for VXLAN Tunnels (Layer 3)

This chapter provides information about EVPN for VXLAN tunnels (Layer 3).

Topics in this chapter include:

Applicability

This chapter is applicable to SR OS and was initially written for Release 12.0.R4. The MD-CLI in the current edition is based on SR OS Release 21.10.R3. Ethernet Virtual Private Network (EVPN) is a control plane technology and does not have line card hardware dependencies.

Chapter EVPN for VXLAN Tunnels (Layer 2) is prerequisite reading.

Overview

As discussed in the EVPN for VXLAN Tunnels (Layer 2) chapter, EVPN and VXLAN can be enabled on VPLS or R-VPLS services in SR OS. Where that chapter focuses on the use of EVPN-VXLAN layer 2 services, in other words, how EVPN-VXLAN is configured in VPLS services, this chapter describes how EVPN-VXLAN can be used to provide inter-subnet forwarding in R-VPLS and VPRN services. Inter-subnet forwarding can be provided by regular R-VPLS and VPRN services. However, EVPN provides an efficient and unified way to populate Forwarding Databases (FDBs), Address Resolution Protocol (ARP) tables, and routing tables using a single BGP address family. Inter-subnet forwarding in overlay networks would otherwise require data plane learning and the use of routing protocols on a per VPRN basis.

The SR OS solution for inter-subnet forwarding using EVPN is based on building blocks described in *draft-ietf-bess-evpn-inter-subnet-forwarding* and the use of the EVPN IP-prefix routes (route type 5) as described in RFC 9136. This example describes three supported common scenarios and provides the CLI configuration and required tools to troubleshoot EVPN-VXLAN in each case. The scenarios configured and described are:

- EVPN-VXLAN in R-VPLS services
- EVPN-VXLAN in Integrated Routing Bridging (IRB) backhaul R-VPLS services
- EVPN-VXLAN in EVPN tunnel R-VPLS services

In all these scenarios, redundant PEs are usually deployed. If that is the case, the interaction of EVPN, IP-VPN, and the Routing Table Manager (RTM) may lead to some routing loop situations that must be avoided by using routing policies (this also may happen in traditional IP-VPN deployments when eBGP and MP-BGP interact to populate VPRN routing tables in multi-homed networks). This chapter describes when those routing loops can happen and how to avoid them.

The term IRB interface refers to an R-VPLS service bound to a VPRN IP interface. The terms IRB interface and R-VPLS interface are used interchangeably throughout this chapter.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Configuration

This section describes the configuration of EVPN-VXLAN for Layer 3 services on SR OS, as well as the available troubleshooting and show commands. The three scenarios described in the overview are analyzed independently.

EVPN-VXLAN in an R-VPLS service

Figure 108: EVPN-VXLAN for R-VPLS services shows the topology used in the first scenario.

Figure 108: EVPN-VXLAN for R-VPLS services

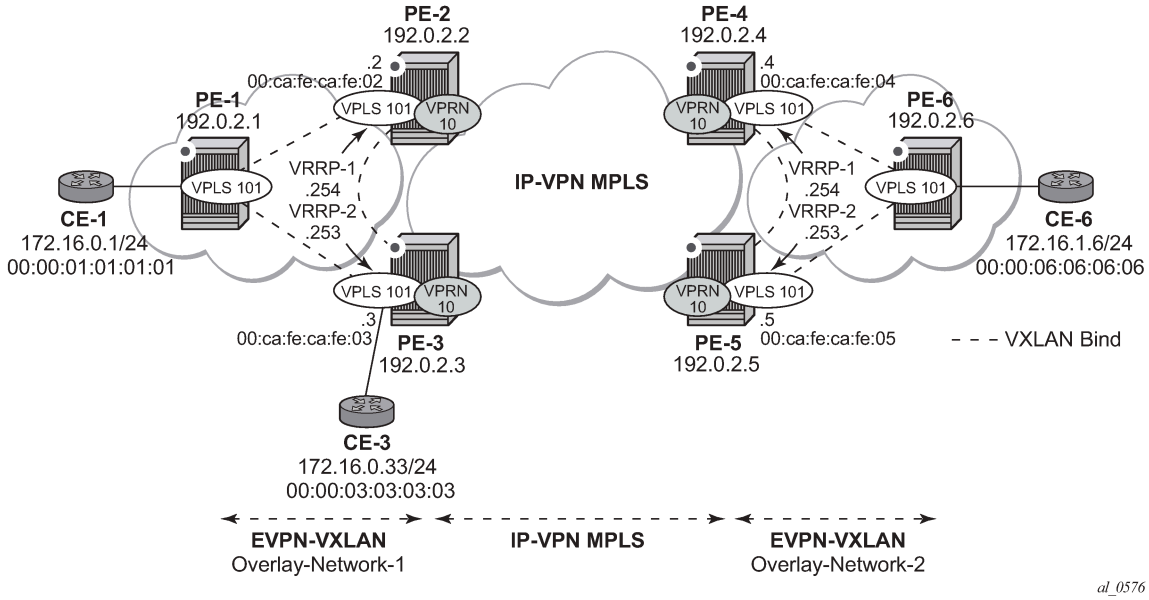

The network topology shows two overlay (VXLAN) networks interconnected by an MPLS network:

- PE-1, PE-2, and PE-3 are part of Overlay-Network-1
- PE-4, PE-5, and PE-6 are part of Overlay-Network-2

A Layer 2/Layer 3 service is provided to a customer to connect CE-1, CE-3, and CE-6. In this scenario, Layer 2 connectivity is provided within each overlay network and inter-subnet connectivity (Layer 3) is provided between the overlay networks. VPLS "evi-101" is defined within each overlay network and VPRN "VPRN10" connects both Layer 2 services through an IP-VPN MPLS network.

This topology can illustrate a Data Center Interconnect (DCI) example, where Overlay-Network-1 and Overlay-Network-2 are two data centers interconnected through an MPLS WAN. In this application, CE-1, CE-3, and CE-6 simulate virtual machines or appliances, PE-2/3/4/5 act as Data Center Gateways (DC GWs) and PE-1/6 as Network Virtualization Edge devices (or virtual PEs running on a compute infrastructure).

The following protocols and objects are configured beforehand:

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- The ports interconnecting the six PEs in Figure 108: EVPN-VXLAN for R-VPLS services are configured as network or hybrid ports and have router network interfaces defined in them. Only the ports connected to the CEs are configured as access ports.
- The six PEs are running IS-IS for the global routing table with the four core PEs interconnected using IS-IS Level-2 point-to-point interfaces and each overlay network using IS-IS Level-1 point-to-point interfaces.
- LDP is used as the MPLS protocol to signal transport tunnel labels among PE-2, PE-3, PE-4, and PE-5. There is no LDP running within the overlay networks.
- The network port MTU (in all the ports sending/receiving VXLAN packets) must be at least 50 bytes (54 if dot1q encapsulation is used) greater than the service MTU to accommodate the size of the VXLAN header.

Once the IGP infrastructure and LDP in the core are enabled, BGP is configured. In this scenario, two BGP families must be enabled: EVPN within each overlay network for the exchange of MAC/IP addresses and setting up the flooding domains, and VPN-IPv4 among the four core PEs so that IP prefixes can be exchanged and resolved to MPLS tunnels in the core.

The following MD-CLI shows the BGP configuration of PE-1, which only needs the EVPN family. PE-6 has a similar BGP configuration, that is, only EVPN family is configured for its peers. The use of Route Reflectors (RRs) in these scenarios is common. Although this scenario does not use RRs, an EVPN RR could have been used in Overlay-Network-1 and Overlay-Network-2 and a separate VPN-IPv4 RR could have been used in the core IP-VPN MPLS network.

```
# on PE-1:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            rapid-update {
                evpn true
            }
            group "DC" {
                peer-as 64500
                family {
                    evpn true
                }
            }
            neighbor "192.0.2.2" {
                group "DC"
            }
            neighbor "192.0.2.3" {
                group "DC"
            }
        }
```

The BGP configuration on the DC GWs is as follows:

```
# on PE-2:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            rapid-withdrawal true
            peer-ip-tracking true
            rapid-update {
                evpn true
            }
            group "DC" {
                peer-as 64500
                family {
                    vpn-ipv4 true
                    evpn true
                }
            }
            group "WAN" {
                peer-as 64500
                family {
                    vpn-ipv4 true
                }
            }
            neighbor "192.0.2.1" {
                group "DC"
            }
            neighbor "192.0.2.3" {
                group "DC"
            }
            neighbor "192.0.2.4" {
                group "WAN"
            }
            neighbor "192.0.2.5" {
                group "WAN"
            }
        }
```

```
# on PE-3:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            rapid-update {
                evpn true
            }
            group "DC" {
                peer-as 64500
                family {
                    vpn-ipv4 true
                    evpn true
                }
            }
            group "WAN" {
                peer-as 64500
                family {
                    vpn-ipv4 true
                }
            }
            neighbor "192.0.2.1" {
                group "DC"
            }
            neighbor "192.0.2.2" {
                group "DC"
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            neighbor "192.0.2.4" {
                group "WAN"
            }
            neighbor "192.0.2.5" {
                group "WAN"
            }
        }
```

The DC GWs PE-4 and PE-5 have an equivalent BGP configuration.

Figure 109: BGP adjacencies and enabled families shows the BGP peering sessions among the PEs and the enabled BGP families. PE-1 and PE-6 only establish an EVPN peering session with their peers (only the EVPN family is enabled on PE-1 and PE-6, even if the peer PEs are VPN-IPv4 capable as well).

Figure 109: BGP adjacencies and enabled families

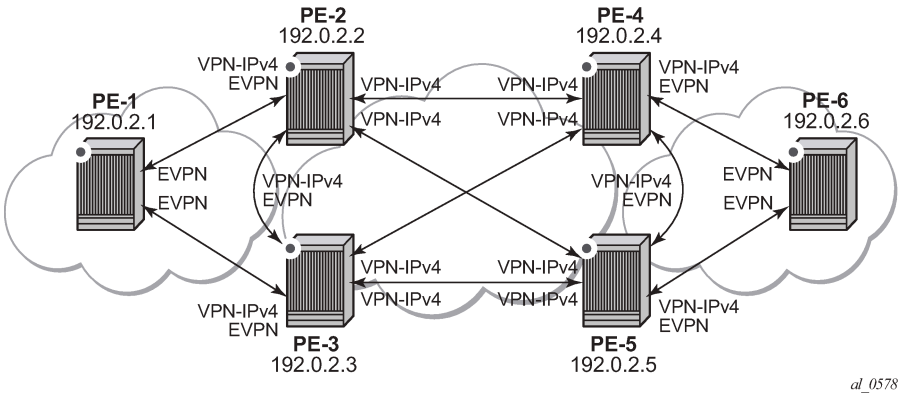

When the network infrastructure is running properly, the actual service configuration, as illustrated in Figure 108: EVPN-VXLAN for R-VPLS services, can be carried out. The following MD-CLI shows the configuration for VPLS "evi-101" and VPRN "VPRN10" in PE-1, PE-2, and PE-3. The other overlay network has a similar configuration.

```
# on PE-1:
configure {
    service {
        vpls "evi-101" {
            admin-state enable
            service-id 101
            customer "1"
            vxlan {
                instance 1 {
                    vni 101
                }
            }
            proxy-arp {
                admin-state enable
            }
            bgp 1 {
                route-distinguisher "192.0.2.1:101"
                route-target {
                    export "target:64500:101"
                    import "target:64500:101"
                }
            }
            bgp-evpn {
                vxlan 1 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    admin-state enable
                    vxlan-instance 1
                }
            }
            sap 1/2/1:101 {
            }
        }
```

Proxy-ARP is disabled (default) on PE-2, as well as on the other core PEs:

```
# on PE-2:
configure {
    service {
        vpls "evi-101" {
            admin-state enable
            service-id 101
            customer "1"
            vxlan {
                instance 1 {
                    vni 101
                }
            }
            routed-vpls {
            }
            bgp 1 {
                route-distinguisher "192.0.2.2:101"
                route-target {
                    export "target:64500:101"
                    import "target:64500:101"
                }
            }
            bgp-evpn {
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
        }
        vprn "VPRN10" {
            admin-state enable
            service-id 10
            customer "1"
            ecmp 2
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "192.0.2.2:10"
                    vrf-target {
                        community "target:64500:10"
                    }
                    auto-bind-tunnel {
                        resolution filter
                        resolution-filter {
                            ldp true
                        }
                    }
                }
            }
            interface "int-1" {
                mac 00:ca:fe:ca:fe:02
                ipv4 {
                    primary {
                        address 172.16.0.2
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    prefix-length 24
                }
                vrrp 1 {
                    backup [172.16.0.254]
                    mac 00:ca:fe:ca:fe:54
                    priority 254
                    ping-reply true
                    traceroute-reply true
                }
                vrrp 2 {
                    backup [172.16.0.253]
                    mac 00:ca:fe:ca:fe:53
                    ping-reply true
                    traceroute-reply true
                }
            }
            vpls "evi-101" {
            }
        }
    }
```

```
# on PE-3:
configure {
    service {
        vpls "evi-101" {
            admin-state enable
            service-id 101
            customer "1"
            vxlan {
                instance 1 {
                    vni 101
                }
            }
            routed-vpls {
            }
            bgp 1 {
                route-distinguisher "192.0.2.3:101"
                route-target {
                    export "target:64500:101"
                    import "target:64500:101"
                }
            }
            bgp-evpn {
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
            sap 1/2/1:101 {
            }
        }
        vprn "VPRN10" {
            admin-state enable
            service-id 10
            customer "1"
            ecmp 2
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "192.0.2.3:10"
                    vrf-target {
                        community "target:64500:10"
                    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    auto-bind-tunnel {
                        resolution filter
                        resolution-filter {
                            ldp true
                        }
                    }
                }
            }
            interface "int-1" {
                mac 00:ca:fe:ca:fe:03
                ipv4 {
                    primary {
                        address 172.16.0.3
                        prefix-length 24
                    }
                    vrrp 1 {
                        backup [172.16.0.254]
                        mac 00:ca:fe:ca:fe:54
                        ping-reply true
                        traceroute-reply true
                    }
                    vrrp 2 {
                        backup [172.16.0.253]
                        mac 00:ca:fe:ca:fe:53
                        priority 254
                        ping-reply true
                        traceroute-reply true
                    }
                }
                vpls "evi-101" {
                }
            }
        }
```

For details about the EVPN and VXLAN configuration in VPLS "evi-101" on PE-1, see chapter EVPN for VXLAN Tunnels (Layer 2). The configuration of VPLS "evi-101" on PE-2 and PE-3 has the following important aspects:

- The **routed-vpls** command is required so that the R-VPLS can be bound to VPRN "VPRN10".
- The service name "evi-101" is configured when the service is created and cannot be modified afterward. The service name must match the name configured in the VPLS interface in VPRN "VPRN10".
- Even though EVPN and VXLAN are properly configured, proxy-ARP cannot be enabled in VPLS "evi-101". In an R-VPLS with EVPN-VXLAN, proxy-ARP is not supported and the VPRN ARP table is used instead. When an EVPN MAC route that includes an IP address is received in an R-VPLS, the MAC-IP pair encoded in the route is added to the ARP table of the VPRN, as opposed to the proxy-ARP table.

```
*[ex:/configure service vpls "evi-101" proxy-arp]
A:admin@PE-2# commit
MINOR: MGMT_CORE #4001: configure service vpls "evi-101" proxy-arp admin-state -
 configuration not supported on routed-vpls - configure service vpls "evi-101" routed-vpls
```

When configuring VPRN "VPRN10" on PE-2 and PE-3, the following considerations must be taken into account:

- When trying to enable existing VPRN features on interfaces linked to EVPN-VXLAN R-VPLS interfaces, the **radius-auth-policy** command is not supported:

```
*[ex:/configure service vprn "VPRN10" interface "int-1"]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
A:admin@PE-2# radius-auth-policy "authPol1"

*[ex:/configure service vprn "VPRN10" interface "int-1"]
A:admin@PE-2# commit
MINOR: MGMT_CORE #4001: configure service vprn "VPRN10" interface "int-1" radius-auth-policy
 - Combination of radius authentication policy and vpls binding not supported - configure
 service vprn "VPRN10" interface "int-1"
```

- Dynamic routing protocols such as IS-IS, RIP, or OSPF are not supported.
- In general, no SR OS control plane generated packets are sent to the egress VXLAN bindings except for ARP, VRRP, ICMP, BFD, and Eth-CFM.
- As shown in Figure 108: EVPN-VXLAN for R-VPLS services and in the CLI excerpts, VRRP can be configured on the VPLS interfaces in VPRN "VPRN10" to provide default gateway redundancy to the hosts connected to VPLS "evi-101". Two VRRP instances are configured so that VPLS "evi-101" upstream traffic can be load-balanced to PE-2 and PE-3. With VRRP on EVPN-VXLAN R-VPLS interfaces:
  - **Ping-reply** and **traceroute-reply** can be configured and are supported. BFD is also supported to speed up the fault detection.
  - **standby-forwarding**, even if it were configured for VRRP, would not have any effect in this configuration: the standby PE will never see any flooded traffic sent to it, so this command is not applicable to this scenario.
- When a VPRN "VPRN10" VPLS interface is bound to VPLS "evi-101", EVPN advertises all the IP addresses configured for that VPLS interface as MAC routes with a static MAC indication. For the remote EVPN peers, that means that those MAC addresses linked to remote IP interfaces are protected. VRRP virtual IP/MACs are also advertised by EVPN as "static" and so protected. In the example of Figure 108: EVPN-VXLAN for R-VPLS services, the VPLS "evi-101" FDB in PE-1 shows the IP interface MAC addresses and VRRP MAC addresses as **EvpnS:P** (Static and protected MAC) as shown in the following output:

```
[/]
A:admin@PE-1# show service id 101 fdb detail

===============================================================================
Forwarding Database, Service 101
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
101        00:00:01:01:01:01 sap:1/2/1:101           L/0      03/02/22 11:34:55
101        00:00:03:03:03:03 vxlan-1:                Evpn     03/02/22 11:35:37
                             192.0.2.3:101
101        00:ca:fe:ca:fe:02 vxlan-1:                EvpnS:P  03/02/22 11:35:05
                             192.0.2.2:101
101        00:ca:fe:ca:fe:03 vxlan-1:                EvpnS:P  03/02/22 11:35:37
                             192.0.2.3:101
101        00:ca:fe:ca:fe:53 vxlan-1:                EvpnS:P  03/02/22 11:35:40
                             192.0.2.3:101
101        00:ca:fe:ca:fe:54 vxlan-1:                EvpnS:P  03/02/22 11:35:08
                             192.0.2.2:101
-------------------------------------------------------------------------------
No. of MAC Entries: 6
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The VPRN "VPRN10" VRRP instances on PE-2 are the following:

```
[/]
A:admin@PE-2# show router 10 vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                  VR Id Own Adm  State       Base Pri   Msg Int
                                IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-1                           1     No  Up   Master       254       1
                                IPv4      Up   n/a         254        No
  Backup Addr: 172.16.0.254
int-1                           2     No  Up   Backup       100       1
                                IPv4      Up   n/a         100        No
  Backup Addr: 172.16.0.253
-------------------------------------------------------------------------------
Instances : 2
===============================================================================
```

The ARP entries for PE-2 are the following:

```
[/]
A:admin@PE-2# show router 10 arp

===============================================================================
ARP Table (Service: 10)
===============================================================================
IP Address      MAC Address       Expiry    Type   Interface
-------------------------------------------------------------------------------
172.16.0.2      00:ca:fe:ca:fe:02 00h00m00s Oth[I] int-1
172.16.0.3      00:ca:fe:ca:fe:03 00h00m00s Evp[I] int-1
172.16.0.253    00:ca:fe:ca:fe:53 00h00m00s Oth    int-1
172.16.0.254    00:ca:fe:ca:fe:54 00h00m00s Oth[I] int-1
-------------------------------------------------------------------------------
No. of ARP Entries: 4
===============================================================================
```

EVPN-VXLAN in IRB backhaul R-VPLS services

Figure 110: EVPN-VXLAN for IRB backhaul R-VPLS services illustrates the second inter-subnet forwarding scenario, where Layer 3 connectivity must be provided not only between the overlay networks but also within each overlay network. In the example shown in Figure 110: EVPN-VXLAN for IRB backhaul R-VPLS services, a customer (tenant) has different subnets and connectivity must be provided across all of them (CE-1, CE-3, and CE-6 must be able to communicate), bearing in mind that EVPN-VXLAN is enabled in each overlay network and IP-VPN MPLS is used to interconnect both overlay networks. VPLS "evi-201" is an IRB Backhaul R-VPLS service because it provides connectivity to the VPRN instances.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 110: EVPN-VXLAN for IRB backhaul R-VPLS services

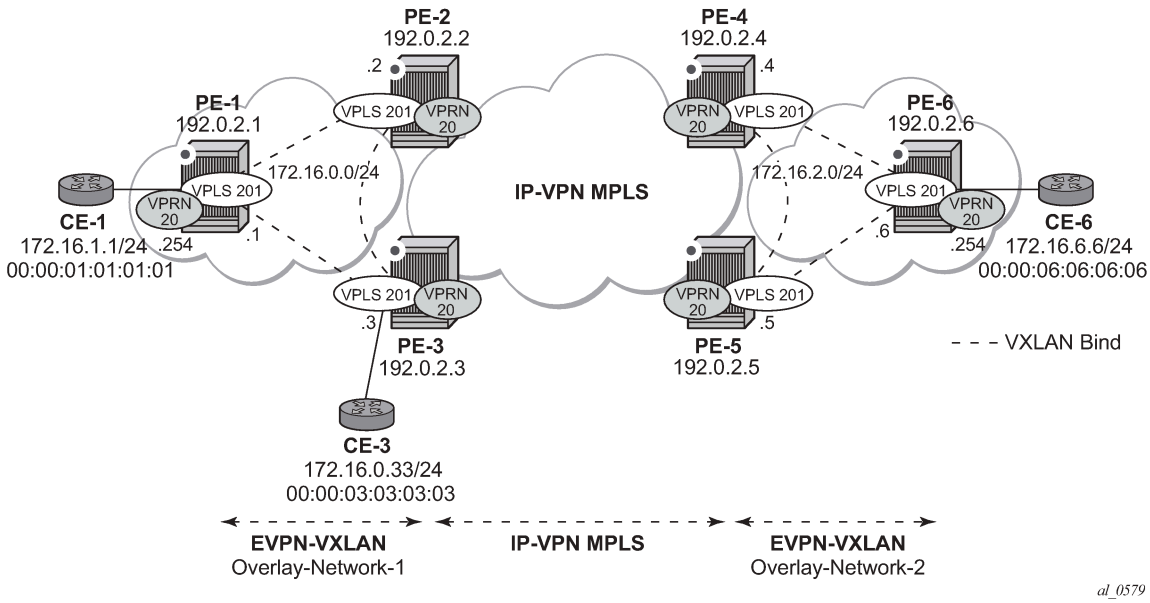

From a BGP peering perspective, there is no change in this scenario compared to the previous one: PE-1 and PE-6 only support the EVPN address family. However, in this scenario, CE-1 is now connected to an R-VPLS directly linked to the VPRN instances in PE-2/PE-3. As a result of that, IP prefixes must be exchanged between PE-1 and PE-2/PE-3. EVPN can advertise not only MAC routes and Inclusive Multicast routes, but also IP prefix routes that contain IP prefixes that can be installed in the attached VPRN routing table.

As an example, the VPRN "VPRN20" and VPLS "evi-201" configurations on PE-1, PE-2, and PE-3 are shown. Similar configurations are needed in PE-4, PE-5, and PE-6.

On PE-1, VPRN "VPRN20" and VPLS "evi-201" are configured as follows:

```
# on PE-1:
configure {
    service {
        vpls "evi-201" {
            admin-state enable
            service-id 201
            customer "1"
            vxlan {
                instance 1 {
                    vni 201
                }
            }
            routed-vpls {
            }
            bgp 1 {
                route-distinguisher "192.0.2.1:201"
                route-target {
                    export "target:64500:201"
                    import "target:64500:201"
                }
            }
            bgp-evpn {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                routes {
                    ip-prefix {
                        advertise true
                    }
                }
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
        }
        vprn "VPRN20" {
            admin-state enable
            service-id 20
            customer "1"
            interface "int-PE-1-CE-1" {
                ipv4 {
                    primary {
                        address 172.16.1.254
                        prefix-length 24
                    }
                }
                sap 1/2/1:20 {
                }
            }
            interface "int-evi-201" {
                ipv4 {
                    primary {
                        address 172.16.0.1
                        prefix-length 24
                    }
                }
                vpls "evi-201" {
                }
            }
        }
```

On PE-2, VPRN "VPRN20" and VPLS "evi-201" are configured as follows:

```
# on PE-2:
configure {
    service {
        vpls "evi-201" {
            admin-state enable
            service-id 201
            customer "1"
            vxlan {
                instance 1 {
                    vni 201
                }
            }
            routed-vpls {
            }
            bgp 1 {
                route-distinguisher "192.0.2.2:201"
                route-target {
                    export "target:64500:201"
                    import "target:64500:201"
                }
            }
            bgp-evpn {
                routes {
                    ip-prefix {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        advertise true
                    }
                }
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
        }
        vprn "VPRN20" {
            admin-state enable
            service-id 20
            customer "1"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "192.0.2.2:20"
                    vrf-target {
                        community "target:64500:20"
                    }
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            interface "int-evi-201" {
                ipv4 {
                    primary {
                        address 172.16.0.2
                        prefix-length 24
                    }
                }
                vpls "evi-201" {
                }
            }
        }
```

On PE-3, VPRN "VPRN20" and VPLS "evi-201" are configured as follows:

```
# on PE-3:
configure {
    service {
        vpls "evi-201" {
            admin-state enable
            service-id 201
            customer "1"
            vxlan {
                instance 1 {
                    vni 201
                }
            }
            routed-vpls {
            }
            bgp 1 {
                route-distinguisher "192.0.2.3:201"
                route-target {
                    export "target:64500:201"
                    import "target:64500:201"
                }
            }
            bgp-evpn {
                routes {
                    ip-prefix {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    advertise true
                }
            }
            vxlan 1 {
                admin-state enable
                vxlan-instance 1
            }
        }
        sap 1/2/1:20 {
        }
    }
    vprn "VPRN20" {
        admin-state enable
        service-id 20
        customer "1"
        bgp-ipvpn {
            mpls {
                admin-state enable
                route-distinguisher "192.0.2.3:20"
                vrf-target {
                    community "target:64500:20"
                }
                auto-bind-tunnel {
                    resolution any
                }
            }
        }
        interface "int-evi-201" {
            ipv4 {
                primary {
                    address 172.16.0.3
                    prefix-length 24
                }
            }
            vpls "evi-201" {
            }
        }
    }
```

As shown in the CLI excerpt, the configuration in the three nodes (PE-1, PE-2, and PE-3) for VPLS "evi-201" and VPRN "VPRN20" is very similar. The main difference is the **auto-bind-tunnel** command in VRPN 20 on PE-2 and PE-3. This command allows the VPRN "VPRN20" on PE-2 and PE-3 to receive IP-VPN routes from the core and resolve them to MPLS tunnels. VPRN "VPRN20" on PE-1 does not require such command because all its IP prefixes are resolved to local interfaces or to EVPN peers.

The **routes>ip-prefix>advertise** command enables:

- The advertisement of IP prefixes in EVPN, in routes type 5. All the existing IP prefixes in the attached VPRN "VPRN20" routing table are advertised in EVPN within the VPLS "evi-201" context (except for the ones associated with VPLS "evi-201" itself).
- The installation of IP prefixes in the attached VPRN "VPRN20" routing table with a preference of 169 (BGP-VPN routes for IP-VPN have a preference of 170) and a next-hop of the gateway IP (GW IP) address included in the EVPN IP prefix route.

For instance, the following output shows that PE-1 advertises the IP prefix 172.16.1.0/24 as an EVPN route to PE-3 (a similar route is sent to PE-2), captured by a **//debug router bgp update** session.

```
# on PE-1:
44 2022/03/02 11:38:45.956 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Send BGP UPDATE:
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    Withdrawn Length = 0
    Total Path Attr Length = 82
    Flag: 0x90 Type: 14 Len: 45 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.1
        Type: EVPN-IP-PREFIX Len: 34 RD: 192.0.2.1:201, tag: 0,
              ip_prefix: 172.16.1.0/24 gw_ip 172.16.0.1 Label: 201 (Raw Label: 0xc9)
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:201
        bgp-tunnel-encap:VXLAN
"
```

The VPRN "VPRN20" routing table in PE-1 includes two EVPN Interface-ful (EVPN-IFF) routes with preference 169, as follows:

```
[/]
A:admin@PE-1# show router 20 route-table

===============================================================================
Route Table (Service: 20)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
172.16.0.0/24                                 Local   Local     00h22m22s  0
       int-evi-201                                                 0
172.16.1.0/24                                 Local   Local     00h22m22s  0
       int-PE-1-CE-1                                               0
172.16.2.0/24                                 Remote  EVPN-IFF  00h01m41s  169
       172.16.0.2                                                  0
172.16.6.0/24                                 Remote  EVPN-IFF  00h01m41s  169
       172.16.0.2                                                  0
-------------------------------------------------------------------------------
No. of Routes: 4
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

The subnet 172.16.0.0/24 is used on the interfaces "int-evi-201" in overlay network 1 and subnet 172.16.2.0/24 is used on similar interfaces in overlay network 2. CE-1 has an IP address in subnet 172.16.1.0/24 and CE-6 has an IP address in subnet 172.16.6.0/24. The next hop to reach 172.16.2.0/24 (overlay network 2) or CE-6, is 172.16.0.2 (PE-2), but it could have been PE-3.

There is redundancy in the example setup and therefore, loops can occur. To avoid loops, routing policies need to be configured on the core PEs (PE-2, PE-3, PE-4, and PE-5). These policies are described in the Use of routing policies to avoid routing loops in redundant PEs section for routing loop use case 1.

The routing table on PE-2 shows a EVPN-IFF route toward CE-1 (subnet 172.16.1.0/24) via PE-1. The route toward CE-6 uses a tunnel toward PE-4 in overlay network 2.

```
[/]
A:admin@PE-2# show router 20 route-table

===============================================================================
Route Table (Service: 20)
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
172.16.0.0/24                                   Local   Local     00h20m36s  0
       int-evi-201                                                     0
172.16.1.0/24                                   Remote  EVPN-IFF  00h02m17s  169
       172.16.0.1                                                      0
172.16.2.0/24                                   Remote  BGP VPN   00h01m43s  170
       192.0.2.4 (tunneled)                                            10
172.16.6.0/24                                   Remote  BGP VPN   00h01m43s  170
       192.0.2.4 (tunneled)                                            10
-------------------------------------------------------------------------------
No. of Routes: 4
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

The routing table on PE-3 is as follows:

```
[/]
A:admin@PE-3# show router 20 route-table

===============================================================================
Route Table (Service: 20)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
172.16.0.0/24                                   Local   Local     00h09m20s  0
       int-evi-201                                                     0
172.16.1.0/24                                   Remote  EVPN-IFF  00h02m46s  169
       172.16.0.1                                                      0
172.16.2.0/24                                   Remote  BGP VPN   00h02m20s  170
       192.0.2.4 (tunneled)                                            10
172.16.6.0/24                                   Remote  BGP VPN   00h01m53s  170
       192.0.2.4 (tunneled)                                            10
-------------------------------------------------------------------------------
No. of Routes: 4
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

When checking the operation of EVPN in this scenario, it is important to observe that the right next hops and prefixes are successfully installed in the routing table of VPRN "VPRN20":

- EVPN IP prefixes are sent using a GW IP address matching the primary IP interface address of the R-VPLS for which the routes are sent. For instance, as shown above, IP prefix 172.16.1.0/24 is advertised from PE-1 with GW IP address 172.16.0.1, which is the IP address configured for the VPLS interface in VPRN "VPRN20" on PE-1. In the routing tables of VPRN "VPRN20" on PE-2 and PE-3, IP prefix 172.16.1.0/24 is installed with next hop 172.16.0.1. Traffic arriving at PE-2 or PE-3 on VPRN "VPRN20" with IP Destination Address (DA) in the 172.16.1.0/24 subnet matches the mentioned routing table entry. As usual, the next-hop is resolved by the ARP table to a MAC address and the MAC address resolved by the FDB table to an egress VTEP, VNI.
- IP prefixes in the routing table of VPRN "VPRN20" are advertised in IP-VPN to the remote IP-VPN MPLS peers. Received IP-VPN prefixes are installed in the routing table of VPRN "VPRN20" using the

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

remote PE system IP address as the next hop, as usual. For instance, 172.16.6.0/24 is installed in the routing table of VPRN "VPRN20" on PE-2 with next-hop (tunneled) 192.0.2.4 and preference 170.

The following considerations of how the routing table manager (RTM) handles EVPN and IP-VPN prefixes must be taken into account:

- Only VPRN interface primary addresses are advertised as GW IP in EVPN IP prefix routes. Secondary addresses are never sent as GW IP addresses.
- EVPN IP prefixes are advertised by default as soon as the **routes>ip-prefix>advertise** command is enabled and there are active IP prefixes in the attached VPRN routing table.
- If the same IP prefix is received on a PE via EVPN and IP-VPN at the same time for the same VPRN, by default, the EVPN prefix is selected because its preference (169) is better than the IP-VPN preference (170).
- Because EVPN has a better preference compared to IP-VPN, when the VPRNs on redundant PEs are attached to the same R-VPLS service, routing loops may occur. The use case described here is an example where routing loops can occur. Check the Use of routing policies to avoid routing loops in redundant PEs section to avoid routing loops in redundant PEs for more information.
- When the command **routes>ip-prefix>advertise** is enabled, the subnet IP prefixes are advertised in EVPN but not the host IP prefixes (/32 prefixes associated with the local interfaces). If the user wants to advertise the host IP prefixes as well, the **routes>ip-prefix>include-direct-interface-host** command must be configured. The following example illustrates this.

```
[/]
A:admin@PE-1# show router 20 route-table

===============================================================================
Route Table (Service: 20)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
172.16.0.0/24                                 Local   Local     00h10m12s  0
       int-evi-201                                                 0
172.16.1.0/24                                 Local   Local     00h10m12s  0
       int-PE-1-CE-1                                               0
172.16.2.0/24                                 Remote  EVPN-IFF  00h02m51s  169
       172.16.0.2                                                  0
172.16.6.0/24                                 Remote  EVPN-IFF  00h02m51s  169
       172.16.0.2                                                  0
-------------------------------------------------------------------------------
No. of Routes: 4
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

The host routes can be shown with the **show router route-table all** command:

```
[/]
A:admin@PE-1# show router 20 route-table all

===============================================================================
Route Table (Service: 20)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
      Next Hop[Interface Name]                          Active     Metric
-------------------------------------------------------------------------------
172.16.0.0/24                                 Local   Local     00h10m49s  0
       int-evi-201                                     Y           0
172.16.0.1/32                                 Local   Host      00h10m49s  0
       int-evi-201                                     Y           0
172.16.1.0/24                                 Local   Local     00h10m49s  0
       int-PE-1-CE-1                                   Y           0
172.16.1.254/32                               Local   Host      00h10m49s  0
       int-PE-1-CE-1                                   Y           0
172.16.2.0/24                                 Remote  EVPN-IFF  00h03m28s  169
       172.16.0.2                                      Y           0
172.16.6.0/24                                 Remote  EVPN-IFF  00h03m28s  169
       172.16.0.2                                      Y           0
-------------------------------------------------------------------------------
No. of Routes: 6
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
       E = Inactive best-external BGP route
===============================================================================
```

When the **routes>ip-prefix>include-direct-interface-host** command is enabled on VPLS "evi-201" on PE-1, PE-1 advertises the host routes as well and these are installed in the routing tables on the remote PEs.

```
# on PE-1:
configure {
    service {
        vpls "evi-201" {
            bgp-evpn {
                routes {
                    ip-prefix {
                            advertise true
                            include-direct-interface-host true
                        }
                    }
                }
```

```
[/]
A:admin@PE-2# show router 20 route-table

===============================================================================
Route Table (Service: 20)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
172.16.0.0/24                                 Local   Local     00h11m40s  0
       int-evi-201                                             0
172.16.1.0/24                                 Remote  EVPN-IFF  00h04m59s  169
       172.16.0.1                                              0
172.16.1.254/32                               Remote  EVPN-IFF  00h00m11s  169
       172.16.0.1                                              0
172.16.2.0/24                                 Remote  BGP VPN   00h04m27s  170
       192.0.2.4 (tunneled)                                    10
172.16.6.0/24                                 Remote  BGP VPN   00h04m27s  170
       192.0.2.4 (tunneled)                                    10
-------------------------------------------------------------------------------
No. of Routes: 5
Flags: n = Number of times nexthop is repeated
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

- ECMP is fully supported for the VPRN for EVPN IP prefix routes coming from different GW IP next-hops. However, ECMP is not supported for IP prefixes routes belonging to different owners (EVPN and IP-VPN). ECMP is enabled in VPRN "VPRN20" on PE-1, as follows:

  ```
  # on PE-1:
  configure {
      service {
          vprn "VPRN20" {
              ecmp 2
  ```

  When policies are applied that prevent routing loops, as described in section Use of routing policies to avoid routing loops in redundant PEs, both PE-2 and PE-3 have IP-VPN tunnels for IP prefixes 172.16.2.0/24 and 172.16.6.0/24. In that case, an additional route with a different GW IP as next hop is installed in the routing table for these IP prefixes:

  ```
  [/]
  A:admin@PE-1# show router 20 route-table

  ===============================================================================
  Route Table (Service: 20)
  ===============================================================================
  Dest Prefix[Flags]                            Type    Proto     Age        Pref
        Next Hop[Interface Name]                                    Metric
  -------------------------------------------------------------------------------
  172.16.0.0/24                                 Local   Local     00h12m39s  0
        int-evi-201                                                   0
  172.16.1.0/24                                 Local   Local     00h12m39s  0
        int-PE-1-CE-1                                                 0
  172.16.2.0/24                                 Remote  EVPN-IFF  00h00m08s  169
        172.16.0.2                                                    0
  172.16.2.0/24                                 Remote  EVPN-IFF  00h00m08s  169
        172.16.0.3                                                    0
  172.16.6.0/24                                 Remote  EVPN-IFF  00h00m08s  169
        172.16.0.2                                                    0
  172.16.6.0/24                                 Remote  EVPN-IFF  00h00m08s  169
        172.16.0.3                                                    0
  -------------------------------------------------------------------------------
  No. of Routes: 6
  Flags: n = Number of times nexthop is repeated
         B = BGP backup route available
         L = LFA nexthop available
         S = Sticky ECMP requested
  ===============================================================================
  ```

EVPN-VXLAN in EVPN tunnel R-VPLS services

The previous scenario shows how to use EVPN-VXLAN to provide inter-subnet forwarding for a tenant, where R-VPLS services can contain hosts and also offer transit services between VPRN instances. For example, in the use case shown in Figure 110: EVPN-VXLAN for IRB backhaul R-VPLS services, VPLS "evi-201" in Overlay-Network-1 is an R-VPLS that can provide intra-subnet connectivity to all the hosts in subnet 172.16.0.0/24 (for example, CE-3 belongs to this subnet) but it can also provide transit or

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

backhaul connectivity to hosts in subnet 172.16.1.0/24 (for example, CE-1) sending packets to subnets 172.16.2.0/24 or 172.16.6.0/24.

In some cases, the R-VPLS where EVPN-VXLAN is enabled does not need to provide intra-subnet connectivity and it is purely a transit or backhaul service where VPRN IRB interfaces are connected. Figure 111: EVPN-VXLAN in EVPN-tunnel R-VPLS services illustrates this use case.

Figure 111: EVPN-VXLAN in EVPN-tunnel R-VPLS services

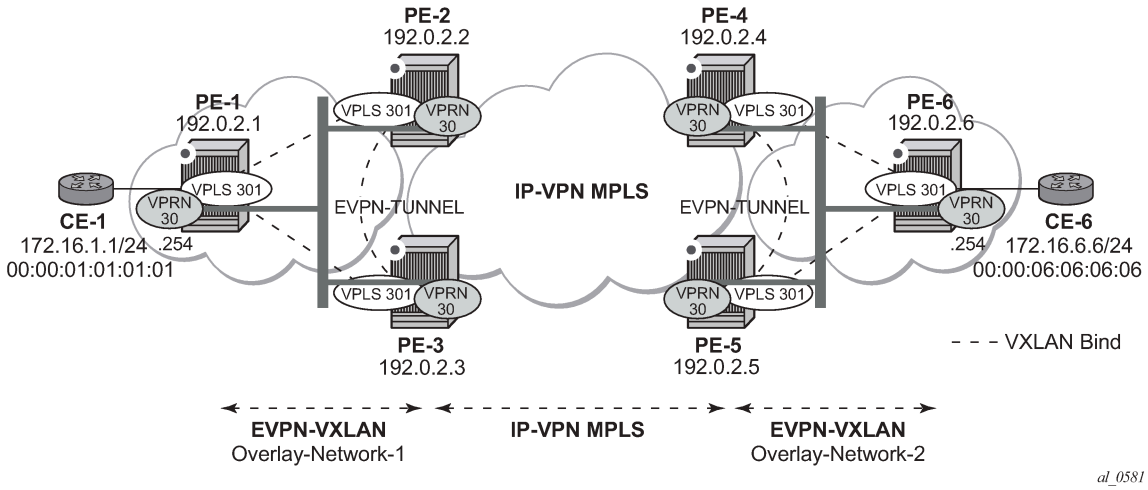

Compared to the preceding use case in Figure 110: EVPN-VXLAN for IRB backhaul R-VPLS services, in this case the R-VPLS connecting the IRB interfaces in Overlay-Network-1 (VPLS "evi-301") does not have any connected host. If that is the case, VPLS "evi-301" can be configured as an EVPN tunnel.

EVPN tunnels are enabled using the **evpn-tunnel** command under the R-VPLS interface configured on the VPRN. EVPN tunnels bring the following benefits to EVPN-VXLAN IRB backhaul R-VPLS services:

- Easier and simpler provisioning of the tenant service: if an EVPN tunnel is configured in an IRB backhaul R-VPLS, there is no need to provision the IRB IP addresses in the VPRN. This makes the provisioning easier to automate and saves IP addresses from the tenant IP space.
- Higher scalability of the IRB backhaul R-VPLS: if EVPN tunnels are enabled, BUM traffic is suppressed in the EVPN-VXLAN IRB backhaul R-VPLS service (it is not required). As a result, the number of VXLAN bindings in IRB backhaul R-VPLS services with EVPN tunnels can be much higher.

As an example, the VPRN "VPRN30" and VPLS "evi-301" configurations on PE-1, PE-2, and PE-3 are shown. Similar configurations are needed in PE-4, PE-5, and PE-6.

```
# on PE-1:
configure {
    service {
        vpls "evi-301" {
            admin-state enable
            service-id 301
            customer "1"
            vxlan {
                instance 1 {
                    vni 301
                }
            }
            routed-vpls {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
            bgp 1 {
                route-distinguisher "192.0.2.1:301"
                route-target {
                    export "target:64500:301"
                    import "target:64500:301"
                }
            }
            bgp-evpn {
                routes {
                    ip-prefix {
                        advertise true
                    }
                }
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
        }
        vprn "VPRN30" {
            admin-state enable
            service-id 30
            customer "1"
            interface "int-PE-1-CE-1" {
                ipv4 {
                    primary {
                        address 172.16.1.254
                        prefix-length 24
                    }
                }
                sap 1/2/1:30 {
                }
            }
            interface "int-evi-301" {
                vpls "evi-301" {
                    evpn-tunnel {
                    }
                }
            }
        }
```

```
# on PE-2:
configure {
    service {
        vpls "evi-301" {
            admin-state enable
            service-id 301
            customer "1"
            vxlan {
                instance 1 {
                    vni 301
                }
            }
            routed-vpls {
            }
            bgp 1 {
                route-distinguisher "192.0.2.2:301"
                route-target {
                    export "target:64500:301"
                    import "target:64500:301"
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            bgp-evpn {
                routes {
                    ip-prefix {
                        advertise true
                    }
                }
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
        }
        vprn "VPRN30" {
            admin-state enable
            service-id 30
            customer "1"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "192.0.2.2:30"
                    vrf-target {
                        community "target:64500:30"
                    }
                    auto-bind-tunnel {
                        resolution filter
                        resolution-filter {
                            ldp true
                        }
                    }
                }
            }
            interface "int-evi-301" {
                vpls "evi-301" {
                    evpn-tunnel {
                    }
                }
            }
        }
```

```
# on PE-3:
configure {
    service {
        vpls "evi-301" {
            admin-state enable
            service-id 301
            customer "1"
            vxlan {
                instance 1 {
                    vni 301
                }
            }
            routed-vpls {
            }
            bgp 1 {
                route-distinguisher "192.0.2.3:301"
                route-target {
                    export "target:64500:301"
                    import "target:64500:301"
                }
            }
            bgp-evpn {
                routes {
                    ip-prefix {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    advertise true
                }
            }
            vxlan 1 {
                admin-state enable
                vxlan-instance 1
            }
        }
    }
    vprn "VPRN30" {
        admin-state enable
        service-id 30
        customer "1"
        bgp-ipvpn {
            mpls {
                admin-state enable
                route-distinguisher "192.0.2.3:30"
                vrf-target {
                    community "target:64500:30"
                }
                auto-bind-tunnel {
                    resolution filter
                    resolution-filter {
                        ldp true
                    }
                }
            }
        }
        interface "int-evi-301" {
            vpls "evi-301" {
                evpn-tunnel {
                }
            }
        }
    }
```

As shown in the preceding output, the configuration in the three nodes (PE-1/2/3) for VPLS "evi-301" and VPRN "VPRN30" is similar to the configuration of VPLS "evi-201" and VPRN "VPRN20" in the preceding scenario. When the **evpn-tunnel** command is added to the VPRN interface, there is no need to configure an IP interface address. The option **evpn-tunnel** can be enabled independently of **routes>ip-prefix>advertise** (although no route type 5 advertisements are sent when **routes>ip-prefix>advertise false** is configured).

A VPRN supports regular IRB backhaul R-VPLS services as well as EVPN tunnel R-VPLS services. A maximum of eight R-VPLS services with **routes>ip-prefix>advertise** enabled per VPRN is supported (in any combination of regular IRB R-VPLS or EVPN tunnel R-VPLS services). EVPN tunnel R-VPLS services do not support SAPs or SDP-bindings. No frames are flooded in an EVPN tunnel R-VPLS service, and, in fact no inclusive multicast routes are exchanged in R-VPLS services that are configured as EVPN tunnels.

The **show service id vxlan destinations** command for an R-VPLS service configured as an EVPN tunnel shows <egress VTEP, VNI> bindings excluded from Mcast, in other words, the VXLAN bindings are not used to flood BUM traffic:

```
[/]
A:admin@PE-2# show service id 301 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
Instance    VTEP Address                          Egress VNI  EvpnStatic Num
 Mcast       Oper State                            L2 PBR     SupBcasDom MACs
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
1           192.0.2.1                               301         evpn       1
 -           Up                                      No          No
1           192.0.2.3                               301         evpn       1
 -           Up                                      No          No
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 2
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN-VXLAN Ethernet Segment Dest
===============================================================================
Instance  Eth SegId                       Num. Macs     Last Change
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

The process followed upon receiving a route type 5 on a regular IRB R-VPLS interface (preceding scenario) differs from the one for an EVPN tunnel type (this scenario):

- IRB backhaul R-VPLS VPRN interface:
  - When a route type 2 that includes an IP address is received and it becomes active, the MAC/IP information is added to the FDB and ARP tables. This can be checked with the **show router 30 arp** command and the **show service id 301 fdb detail** command.
  - When a route type 5 is received on (for instance) PE-2, and becomes active for the R-VPLS service, the IP prefix is added to the VPRN routing table regardless of the existence of a route type 2 that can resolve the GW IP address. If a packet is received from the WAN side and the IP lookup hits an entry for which the GW IP (IP next-hop) does not have an active ARP entry, the system will ARP to get the MAC. If the ARP is resolved but the MAC is unknown in the FDB table, the system floods the ARP message into the R-VPLS multicast list. Routes type 5 can be checked in the routing table with the **show router 30 route-table** command and the **show router 30 fib 1** command.
- EVPN tunnel R-VPLS VPRN interface:
  - When a route type 2 is received and becomes active, the MAC address is added to the FDB (only). This MAC address is normally a GW MAC address.
  - When a route type 5 is received on (for instance) PE-1, the system looks for the GW MAC address. The IP prefix is added to the VPRN routing table with next hop equal to EVPN-tunnel GW MAC; for example, ET-02:13:ff:00:00:6a is an EVPN tunnel with GW MAC 02:13:ff:00:00:6a. The GW MAC address is added from the GW MAC extended community sent along with the route type 5 for prefix 172.16.6.0/24. If a packet is received from CE-1 and the IP lookup hits an entry for which the next hop is an EVPN tunnel: GW MAC, the system looks up the GW MAC address in the FDB. Normally a route type 2 with the GW MAC address has already been received so that the GW MAC address has been added to the FDB. If the GW MAC address is not present in the FDB, the packet will be dropped.
  - The IP prefixes with GW MAC addresses as next hops for the setup in Figure 111: EVPN-VXLAN in EVPN-tunnel R-VPLS services are displayed in the **show router 30 route-table** command, as follows:

```
[/]
A:admin@PE-1# show router 30 route-table

===============================================================================
Route Table (Service: 30)
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                     Metric
-------------------------------------------------------------------------------
172.16.1.0/24                                   Local   Local     00h02m40s  0
       int-PE-1-CE-1                                               0
172.16.6.0/24                                   Remote  EVPN-IFF  00h00m36s  169
       int-evi-301 (ET-02:13:ff:00:00:6a)                          0
-------------------------------------------------------------------------------
No. of Routes: 2
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

The same routing policies are applied on the core PEs to prevent loops; see Use of routing policies to avoid routing loops in redundant PEs.

The **show service id 301 fdb detail** command can be used to look for the forwarding information for a GW MAC address:

```
[/]
A:admin@PE-1# show service id 301 fdb detail

===============================================================================
Forwarding Database, Service 301
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
301        02:0f:ff:00:00:6a cpm                     Intf     03/02/22 11:52:54
301        02:13:ff:00:00:6a vxlan-1:                EvpnS:P  03/02/22 11:53:02
                             192.0.2.2:301
301        02:17:ff:00:00:6a vxlan-1:                EvpnS:P  03/02/22 11:53:09
                             192.0.2.3:301
-------------------------------------------------------------------------------
No. of MAC Entries: 3
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

IP prefix routes sent for EVPN tunnel R-VPLS services do not contain a GW IP address (the GW IP address is zero) but convey a GW MAC address that is used in the peer VPRN routing table. The following output shows PE-2's VPRN "VPRN30" interface MAC address sent to PE-1:

```
[/]
A:admin@PE-2# show router 30 interface "int-evi-301" detail | match "MAC "
MAC Address      : 02:13:ff:00:00:6a    Mac Accounting    : Disabled
```

When **routes>ip-prefix>advertise true** is configured, PE-2 sends route type 5 messages to PE-1, as can be seen in the following BGP update for the route toward subnet 172.16.6.0/24 in overlay network 2, using the MAC address as GW MAC address:

```
# on PE-2:
configure {
    service {
        vpls "evi-301" {
            bgp-evpn {
                routes {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    ip-prefix {
                        advertise true
                    }
                }
            }
```

```
221 2022/03/02 11:54:51.734 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.1
"Peer 1: 192.0.2.1: UPDATE
Peer 1: 192.0.2.1 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 90
    Flag: 0x90 Type: 14 Len: 45 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
        Type: EVPN-IP-PREFIX Len: 34 RD: 192.0.2.2:301, tag: 0,
              ip_prefix: 172.16.6.0/24 gw_ip 0.0.0.0 Label: 301 (Raw Label: 0x12d)
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:301
        mac-nh:02:13:ff:00:00:6a
        bgp-tunnel-encap:VXLAN
"
```

In the routing table of VPRN "VPRN30" on PE-2, IP prefixes are shown with an EVPN tunnel next-hop (GW MAC) as opposed to an IP next-hop, therefore, the user may think that no ARP entries are consumed by VPRN "VPRN30". However, internal ARP entries are still consumed in VPRN "VPRN30". Although not shown in the **show router 30 arp** command, the **summary** option shows the consumption of internal ARP entries for EVPN.

```
[/]
A:admin@PE-2# show router 30 route-table

===============================================================================
Route Table (Service: 30)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
172.16.1.0/24                                 Remote  EVPN-IFF  00h04m28s  169
       int-evi-301 (ET-02:0f:ff:00:00:6a)                          0
172.16.6.0/24                                 Remote  BGP VPN   00h02m40s  170
       192.0.2.4 (tunneled)                                        10
-------------------------------------------------------------------------------
No. of Routes: 2
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

There are no entries in the ARP table:

```
[/]
A:admin@PE-2# show router 30 arp

===============================================================================
ARP Table (Service: 30)
===============================================================================
IP Address      MAC Address       Expiry    Type   Interface
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
No Matching Entries Found
===============================================================================
```

One internal BGP-EVPN ARP entry is consumed, as can be seen as follows:

```
[/]
A:admin@PE-2# show router 30 arp summary

===============================================================
ARP Table Summary (Service: 30)
===============================================================
Local ARP Entries    : 1
Static ARP Entries   : 0
Dynamic ARP Entries  : 0
Managed ARP Entries  : 0
Internal ARP Entries : 0
BGP-EVPN ARP Entries : 1
---------------------------------------------------------------
No. of ARP Entries   : 2
===============================================================
```

The number of BGP-EVPN ARP entries in the **show router 30 arp summary** command matches the number of remote valid GW MAC addresses for VPRN "VPRN30".

Routing policies for IP prefixes in EVPN

Routing policies are supported for IP prefixes imported or exported through BGP EVPN. The default import and export behavior for IP prefixes in EVPN can be modified by using routing policies applied either at peer level (**config>router>bgp>group/neighbor>import/export**) or VPLS level (**config>service>vpls>bgp>vsi-import/vsi-export**).

When applying routing policies to control the distribution of prefixes between EVPN and IP-VPN, the user must take into account that both families are completely separated as far as BGP is concerned and that when prefixes from a family are imported in the RTM, the BGP attributes are lost to the other family. The use of tags allows the controlled distribution of prefixes across the two families.

Figure 112: Routing policies for egress EVPN routes illustrates how VPN-IPv4 routes are imported into the RTM and then passed onto EVPN for its own processing. VPN-IPv4 routes can be tagged at ingress and this tag is preserved throughout the RTM and EVPN processing so that the tag can be matched by the egress BGP routing policy. In this example, egress EVPN routes matching tag 10, are modified to add a site-of-origin (SOO) community origin:64500:1.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 112: Routing policies for egress EVPN routes

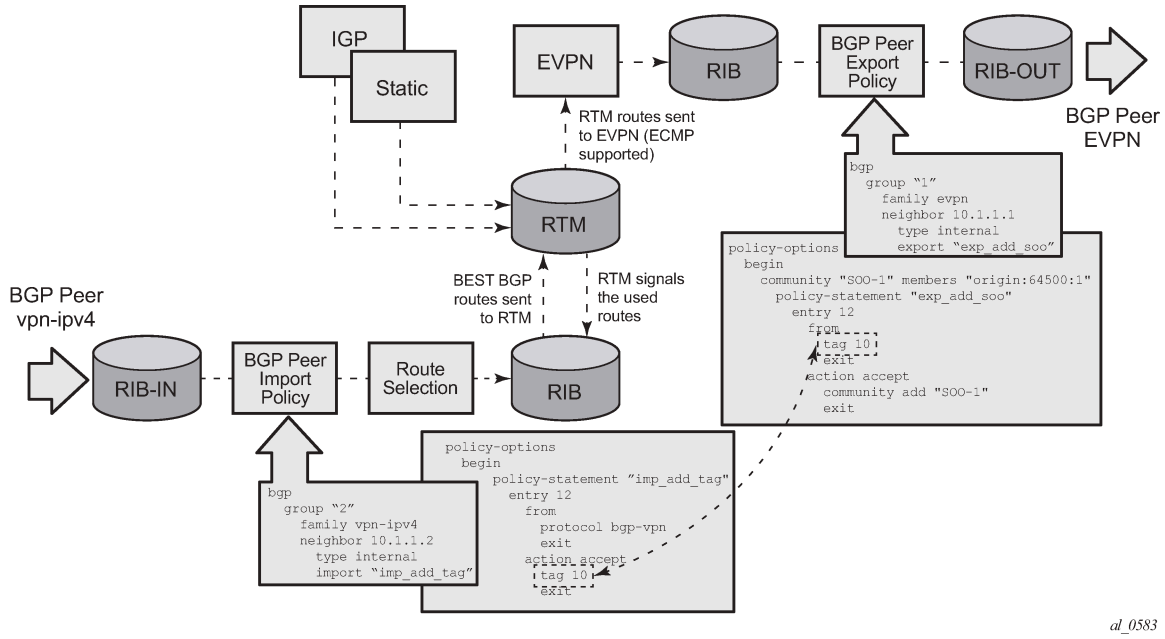

Policy tags can be used to match EVPN IP-prefixes that were learned not only from BGP VPN-IPv4 but also from other routing protocols.

```
*[ex:/configure policy-options policy-statement "test" entry 10 action]
A:admin@PE-1# tag ?

 tag (<number> | <string>)
 <number>  - <1..4294967295>
 <string>  - <1..32 characters>

    OSPF RIP or IS-IS tag applied to routes
```

Figure 113: Routing policies for ingress EVPN routes illustrates the reverse workflow: routes imported from EVPN and exported from RTM to BGP VPN-IPv4. In this example, EVPN routes received with community VM-mob are tagged with tag 200. At the egress VPN-IPv4 peers, only the routes with tag 200 are advertised.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 113: Routing policies for ingress EVPN routes

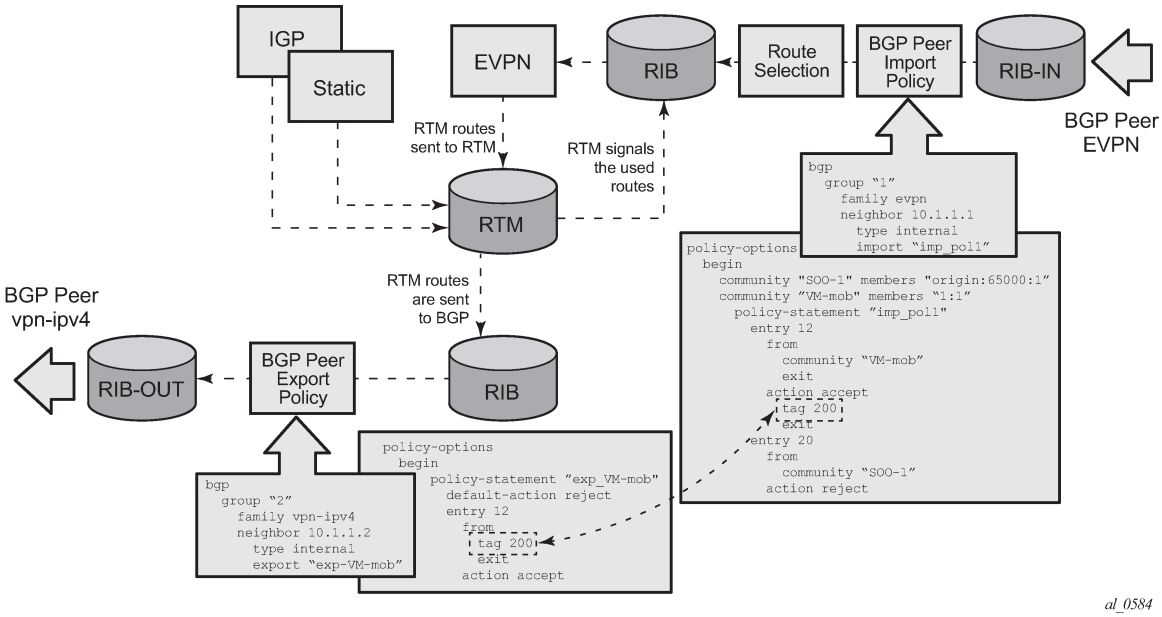

The preceding behavior and the use of tags is also valid for **vsi-import** and **vsi-export** policies. The behavior can be summarized in the following statements:

- For EVPN prefix routes received and imported in RTM:
  - Routes can be matched on communities and tags can be added to them. This works at peer level or vsi-import level.
  - Well-known communities **[no-export | no-export-subconfed | no-advertise]** also require that the routing policies add a tag if the user wants to modify the behavior when exporting to BGP.
  - Routes can be matched based on family EVPN.
  - Routes cannot be matched on prefix list.
- For exporting RTM to EVPN prefix routes:
  - Routes can be matched on tags and based on that, communities added, or routes accepted or rejected, and so on. This works at peer level or vsi-export level.
  - Tags can be added for static routes, RIP, OSPF, IS-IS, and BGP and then be matched in the vsi-export policy for EVPN IP-prefix route advertisement.
  - Tags cannot be added for direct routes.

Use of routing policies to avoid routing loops in redundant PEs

When redundant PE VPRN instances are connected to the same R-VPLS service (IRB backhaul or EVPN tunnel R-VPLS) with the **routes>ip-prefix>advertise true** command enabled, routing loops can occur in two different use cases:

1. Routing loop caused by EVPN and IP-VPN interaction in the RTM.
2. Routing loop caused by EVPN in parallel R-VPLS services.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Policy configuration examples for both cases are provided in the following sections.

Routing loop use-case 1: EVPN and IP-VPN interaction

This use case refers to scenarios with redundant PEs and VPRNs connected to the same R-VPLS with **routes>ip-prefix>advertise true**. The scenarios in Figure 110: EVPN-VXLAN for IRB backhaul R-VPLS services (EVPN-VXLAN for IRB Backhaul R-VPLS services) and Figure 111: EVPN-VXLAN in EVPN-tunnel R-VPLS services (EVPN-VXLAN in EVPN tunnel R-VPLS services) are examples of this use case. In both scenarios, the following process causes a routing loop:

1. PE-4 advertises IP prefix 172.16.6.0/24 with preference 170 (IP-VPN) to PE-2 and PE-3.
2. PE-2 and PE-3 import prefix 172.16.6.0/24 in the VPRN routing table. PE-2 re-advertises prefix 172.16.6.0/24 with preference 169 (EVPN) to PE-1 and PE-3; PE-3 re-advertises the IP prefix in EVPN to PE-1 and PE-2.
3. PE-2 and PE-3 already have the 172.16.6.0/24 prefix in the VPRN routing table with preference 170 (IP-VPN) but because the IP prefix from EVPN has a lower preference (169), the RTM installs the EVPN prefix in the VPRN routing table.
4. PE-2 advertises the EVPN-learned IP prefix 172.16.6.0/24 to all MP-BGP VPN-IPv4 peers, including PE-3; PE-3 advertises the prefix 172.16.6.0/24 to all MP-BGP VPN-IPv4 peers, including PE-2.
5. PE-2 receives the IP prefix 172.16.6.0/24 again from PE-3 and advertises it in EVPN again, creating a routing loop. The same thing happens in PE-3.

This routing loop also happens in traditional multi-homed IP-VPN scenarios where the PE-CE eBGP and MP-BGP VPN-IPv4/v6 protocols interact in the same VPRN RTM, with different router preferences. In either case (EVPN or eBGP interaction with MP-BGP) the issue can be solved by using routing policies and site-of-origin communities.

Routing policies are applied to PE-2 and PE-3 (also to PE-4 and PE-5) and allow the redundant PEs to reject their own generated routes to avoid the loops. These routing policies can be applied at vsi-import/export level or BGP group/neighbor level. The following output shows an example of routing policies applied at BGP neighbor level for PE-2 (similar policies are applied on PE-3, PE-4, and PE-5). Neighbor or group level policies are the preferred way in this kind of use case: a single set of policies is sufficient, as opposed to a set of policies per service (if the policies are applied at vsi-import or vsi-export level).

The following policies are applied in the BGP group or BGP neighbor context on PE-2:

```
# on PE-2:
configure {
    policy-options {
        community "SOO-PE-2" {
            member "origin:2:1" { }
        }
        community "SOO-PE-3" {
            member "origin:3:1" { }
        }
        policy-statement "add-SOO_on_export" {
            entry 10 {
                from {
                    tag 2
                }
                action {
                    action-type accept
                    community {
                        add ["SOO-PE-2"]
                    }
                }
            }
            entry 20 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                from {
                    tag 3
                }
                action {
                    action-type accept
                    community {
                        add ["SOO-PE-3"]
                    }
                }
            }
        }
        policy-statement "reject_based_on_SOO" {
            entry 10 {
                from {
                    community {
                        name "SOO-PE-2"
                    }
                }
                action {
                    action-type reject
                }
            }
            entry 20 {
                from {
                    community {
                        name "SOO-PE-3"
                    }
                }
                action {
                    action-type reject
                }
            }
        }
        policy-statement "add-tag_to_bgp-evpn_routes" {
            entry 10 {
                from {
                    family [evpn]
                }
                action {
                    action-type accept
                    tag 2
                }
            }
        }
        policy-statement "add-tag_to_bgp-vpn_routes" {
            entry 10 {
                from {
                    protocol {
                        name [bgp-vpn]
                    }
                }
                action {
                    action-type accept
                    tag 2
                }
            }
        }
    }
    router "Base" {
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            rapid-update {
                evpn true
            }
            group "DC" {
                peer-as 64500
                family {
                    vpn-ipv4 true
                    evpn true
                }
            }
            group "WAN" {
                peer-as 64500
                family {
                    vpn-ipv4 true
                }
                import {
                    policy ["add-tag_to_bgp-vpn_routes"]
                }
            }
            neighbor "192.0.2.1" {
                group "DC"
                import {
                    policy ["add-tag_to_bgp-evpn_routes"]
                }
            }
            neighbor "192.0.2.3" {
                group "DC"
                import {
                    policy ["reject_based_on_SOO"]
                }
                export {
                    policy ["add-SOO_on_export"]
                }
            }
            neighbor "192.0.2.4" {
                group "WAN"
            }
            neighbor "192.0.2.5" {
                group "WAN"
            }
```

EVPN and MP-BGP routes are tagged at import; on export, a site-of-origin community is added. Routes exchanged between the two redundant PEs are dropped if they are received by a PE with its own site-of-origin.

Routing loop use-case 2: EVPN in parallel R-VPLS services

If a VPRN is connected to more than one R-VPLS with **routes>ip-prefix>advertise true** configured, IP prefixes that belong to one R-VPLS are advertised into the other R-VPLS and vice versa. When redundant PEs are used, a routing loop will occur. Figure 114: EVPN in parallel R-VPLS services illustrates this use case. The example shows R-VPLS with an EVPN tunnel configuration, but the same routing loop occurs for regular IRB backhaul R-VPLS services.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 114: EVPN in parallel R-VPLS services

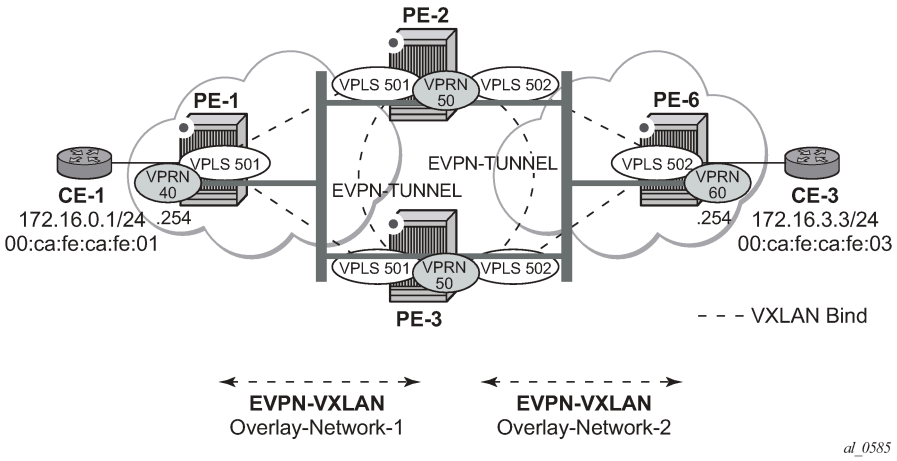

The configuration of VPRN "VPRN50" as well as VPLSs "evi-501" and "evi-502" and the required policies are as follows. For this use case, policies must be applied at vsi-import/export level because more granularity is required when modifying the imported/exported routes.

```
# on PE-2:
configure {
    policy-options {
        community "SOO-PE-2-RVPLS" {
            member "origin:2:11" { }
        }
        community "SOO-PE-3-RVPLS" {
            member "origin:3:11" { }
        }
        community "SOO_PE-3_RVPLS501" {
            member "origin:3:11" { }
            member "target:64500:501" { }
        }
        community "SOO_PE-3_RVPLS502" {
            member "origin:3:11" { }
            member "target:64500:502" { }
        }
        community "exp_RVPLS501" {
            member "origin:2:11" { }
            member "target:64500:501" { }
        }
        community "exp_RVPLS502" {
            member "origin:2:11" { }
            member "target:64500:502" { }
        }
        policy-statement "vsi-export-policy-501" {
            entry 10 {
                from {
                    tag 12
                }
                action {
                    action-type accept
                    community {
                        add ["SOO_PE-3_RVPLS501"]
                    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
        }
        entry 20 {
            action {
                action-type accept
                community {
                    add ["exp_RVPLS501"]
                }
            }
        }
    }
    policy-statement "vsi-export-policy-502" {
        entry 10 {
            from {
                tag 12
            }
            action {
                action-type accept
                community {
                    add ["SOO_PE-3_RVPLS502"]
                }
            }
        }
        entry 20 {
            action {
                action-type accept
                community {
                    add ["exp_RVPLS502"]
                }
            }
        }
    }
    policy-statement "vsi-import-policy-501" {
        entry 10 {
            from {
                community {
                    name "SOO-PE-2-RVPLS"
                }
            }
            action {
                action-type reject
            }
        }
        entry 20 {
            from {
                community {
                    name "SOO_PE-3_RVPLS501"
                }
            }
            action {
                action-type accept
                tag 12
            }
        }
        default-action {
            action-type accept
        }
    }
    policy-statement "vsi-import-policy-502" {
        entry 10 {
            from {
                community {
                    name "SOO-PE-2-RVPLS"
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
            action {
                action-type reject
            }
        }
        entry 20 {
            from {
                community {
                    name "SOO_PE-3_RVPLS502"
                }
            }
            action {
                action-type accept
                tag 12
            }
        }
        default-action {
            action-type accept
        }
    }
}
service {
    vpls "evi-501" {
        admin-state enable
        service-id 501
        customer "1"
        vxlan {
            instance 1 {
                vni 501
            }
        }
        routed-vpls {
        }
        bgp 1 {
            route-distinguisher "192.0.2.2:501"
            vsi-import ["vsi-import-policy-501"]
            vsi-export ["vsi-export-policy-501"]
        }
        bgp-evpn {
            routes {
                ip-prefix {
                    advertise true
                }
            }
            vxlan 1 {
                admin-state enable
                vxlan-instance 1
            }
        }
    }
    vpls "evi-502" {
        admin-state enable
        service-id 502
        customer "1"
        vxlan {
            instance 1 {
                vni 502
            }
        }
        routed-vpls {
        }
        bgp 1 {
            route-distinguisher "192.0.2.2:502"
            vsi-import ["vsi-import-policy-502"]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                vsi-export ["vsi-export-policy-502"]
            }
            bgp-evpn {
                routes {
                    ip-prefix {
                        advertise true
                    }
                }
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
        }
        vprn "VPRN50" {
            admin-state enable
            service-id 50
            customer "1"
            interface "int-evi-501" {
                vpls "evi-501" {
                    evpn-tunnel {
                    }
                }
            }
            interface "int-evi-502" {
                vpls "evi-502" {
                    evpn-tunnel {
                    }
                }
            }
        }
```

Troubleshooting and debug commands

For general information about EVPN and VXLAN troubleshooting and debug commands, see chapter EVPN for VXLAN Tunnels (Layer 2). The following information focuses on specific commands for Layer-3 applications.

When troubleshooting and operating an EVPN-VXLAN scenario with inter-subnet forwarding, it is important to check the IP prefixes and next-hops, as well as ARP tables and FDB tables:

- **show router <..> route-table**
- **show router <..> arp**
- **show service id <..> fdb detail**

ICMP commands can also help checking the connectivity. When **traceroute** is used on EVPN-VXLAN in EVPN tunnel interfaces, EVPN tunnel interface hops in the **traceroute** commands are showing the VPRN loopback address or the other non EVPN-tunnel interface address. In VPRN services where all the interfaces are EVPN tunnels, ICMP packets fail until an IP address is configured. The following output shows a traceroute from VPRN "VPRN30" in PE-1 to CE-6 and from PE-2 to CE-1 (see Figure 111: EVPN-VXLAN in EVPN-tunnel R-VPLS services):

```
[/]
A:admin@PE-1# traceroute 172.16.6.6 router-instance "VPRN30"
traceroute to 172.16.6.6, 30 hops max, 40 byte packets
  1  0.0.0.0  * * *
  2  0.0.0.0  * * *
  3  172.16.6.254 (172.16.6.254)    4.79 ms  4.65 ms  4.80 ms
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
  4  172.16.6.6 (172.16.6.6)    5.32 ms  5.12 ms  4.86 ms
```

```
[/]
A:admin@PE-2# traceroute 172.16.1.1 router-instance "VPRN30"
traceroute to 172.16.1.1, 30 hops max, 0 byte packets
No route to destination. Address: 172.16.1.1, Router Instance:VPRN30
```

When troubleshooting R-VPLS services, specifically R-VPLS services configured as EVPN tunnels, the limit of peer PEs per EVPN tunnel service is much higher than for a regular R-VPLS service because the egress <VTEP, VNI> bindings do not have to be added to the multicast flooding list. For this reason, the following **tools dump** command has been added to check the consumed/total EVPN tunnel next hops. The number of EVPN tunnel next hops matches the number of remote GW MAC addresses per EVPN tunnel R-VPLS service.

```
[/]
A:admin@PE-1# tools dump service id 501 evpn usage

Evpn Tunnel Interface IP Next Hop: 2/8189
```

Finally, when troubleshooting EVPN routes and routing policies, the **show router bgp routes evpn** command and its filters can help:

- Check that the expected routes are received, properly imported, and communities/tags added/replaced/removed.
- Check that the expected routes are sent, properly exported, and communities added/replaced/removed.

Examples of EVPN IP prefix routes including communities and tags are the following.

```
[/]
A:admin@PE-2# show router bgp routes evpn ?

 auto-disc             - Display BGP EVPN Auto-Disc Routes
 eth-seg               - Display BGP EVPN Eth-Seg Routes
 incl-mcast            - Display BGP EVPN Inclusive-Mcast Routes
 ip-prefix             - Display BGP EVPN IPv4-Prefix Routes
 ipv6-prefix           - Display BGP EVPN IPv6-Prefix Routes
 mac                   - Display BGP EVPN Mac Routes
 mcast-join-synch      - Display BGP EVPN Mcast Join Sync Routes
 mcast-leave-synch     - Display BGP EVPN Mcast Leave Sync Routes
 smet                  - Display BGP EVPN Smet Routes
 spmsi-ad              - Display BGP EVPN Spmsi AD Routes
```

```
[/]
A:admin@PE-2# show router bgp routes evpn ip-prefix ?

 ip-prefix [[hunt-detail] <keyword>] [rd <string>] [prefix <string>]
           [community <string>] [tag <string>] [next-hop <string>]
           [aspath-regex <string>]

 [hunt-detail] <keyword>
 <keyword>  - (hunt|detail)

 [hunt-detail]         - keywords
 aspath-regex          - string '<1..80 characters>'
 community             - <as-number1:comm-val1>|<ext-comm>|, <well-known-comm>,
                         ext-comm - <type>:{<ip-address:comm-val1>|,
                         <as-number1:comm-val2>|,  <as-number2:comm-val1>},
                         as-number1 - [0..65535], comm-val1 - [0..65535],
                         type  - target|origin, ip-address - a.b.c.d,
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                          comm-val2 - [0..4294967295], as-number2  - [0..4294967295],
                          well-known-comm - null|no-export|no-export-subconfed|,
                          no-advertise
 next-hop              - Attribute next-hop for ip-prefix
 prefix                - ip-prefix/ip-prefix-length
 rd                    - {<ip-addr:comm-val>|, <2byte-asnumber:ext-comm-val>|,
                          <4byte-asnumber:comm-val>}
 tag                   - Attribute tag for ip-prefix
```

Routing policy "vsi-export-policy-502" adds community "origin:2:11 target:64500:502" to the outgoing routes, as can be verified as follows:

```
[/]
A:admin@PE-2# show router bgp routes evpn ip-prefix hunt prefix 172.16.1.0/24
===============================================================================
 BGP Router ID:192.0.2.2        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN IP-Prefix Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
---snip---
-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
Network        : n/a
Nexthop        : 192.0.2.2
Path Id        : None
To             : 192.0.2.1
Res. Nexthop   : n/a
Local Pref.    : 100                    Interface Name : NotAvailable
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : n/a
Connector      : None
Community      : origin:2:11 target:64500:502
                 mac-nh:02:13:ff:00:01:33 bgp-tunnel-encap:VXLAN
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.1
Origin         : IGP
AS-Path        : No As-Path
EVPN type      : IP-PREFIX
ESI            : n/a
Tag            : 0
Gateway Address: 02:13:ff:00:01:33
Prefix         : 172.16.1.0/24
Route Dist.    : 192.0.2.2:502
MPLS Label     : VNI 502
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
---snip---
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

On PE-2, policy "add-tag_to_bgp-evpn_routes" adds route tag 2 to all BGP EVPN routes, as can be verified in the following output:

```
[/]
A:admin@PE-2# show router bgp routes evpn ip-prefix prefix 172.16.1.0/24 detail
===============================================================================
 BGP Router ID:192.0.2.2       AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN IP-Prefix Routes
===============================================================================
Original Attributes

Network        : n/a
Nexthop        : 192.0.2.1
Path Id        : None
From           : 192.0.2.1
Res. Nexthop   : 192.168.12.1
Local Pref.    : 100                    Interface Name : int-PE-2-PE-1
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : 10
Connector      : None
Community      : target:64500:201 bgp-tunnel-encap:VXLAN
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.1
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
EVPN type      : IP-PREFIX
ESI            : n/a
Tag            : 0
Gateway Address: 172.16.0.1
Prefix         : 172.16.1.0/24
Route Dist.    : 192.0.2.1:201
MPLS Label     : VNI 201
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
Add Paths Send : Default
Last Modified  : 00h04m30s

Modified Attributes

Network        : n/a
Nexthop        : 192.0.2.1
Path Id        : None
From           : 192.0.2.1
Res. Nexthop   : 192.168.12.1
Local Pref.    : 100                    Interface Name : int-PE-2-PE-1
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : 10
Connector      : None
Community      : target:64500:201 bgp-tunnel-encap:VXLAN
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
EVPN type      : IP-PREFIX
ESI            : n/a
Tag            : 0
Gateway Address: 172.16.0.1
Prefix         : 172.16.1.0/24
Route Dist.    : 192.0.2.1:201
MPLS Label     : VNI 201
Route Tag      : 2
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                     Dest Class     : 0
Add Paths Send : Default
Last Modified  : 00h04m30s

-------------------------------------------------------------------------------
---snip---
```

Conclusion

SR OS supports not only the EVPN control plane for VXLAN tunnels in Layer 2 applications but also the simultaneous use of EVPN and VXLAN for VPN customers (tenants) with intra and inter-subnet connectivity requirements. R-VPLS services can be configured to provide default gateway connectivity to hosts, IRB backhaul connectivity to VPRN services, and EVPN tunnel connectivity to VPRN services.

When configured to do so, EVPN can advertise IP prefixes and interact with the VPRN RTM to propagate IP prefix connectivity between EVPN and other routing protocols in the VPRN, including IP-VPN. This example has shown how to configure R-VPLS services for all these functions, as well as how to configure routing policies for EVPN-based IP prefixes.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

EVPN Interconnect Ethernet Segments

This chapter provides information about EVPN Interconnect Ethernet Segments.

Topics in this chapter include:

Applicability

This chapter was initially written based on SR OS Release 15.0.R4, but the MD-CLI in the current edition corresponds to SR OS Release 21.2.R2.

Chapters EVPN for MPLS Tunnels, EVPN for VXLAN Tunnels (Layer 2) and EVPN-MPLS Interconnect for EVPN-VXLAN VPLS Services are prerequisite reading.

Overview

SR OS supports Interconnect Ethernet Segments (I-ESs) for VXLAN as per the IETF *draft-ietf-bess-dci-evpn-overlay*. An I-ES is a virtual Ethernet Segment (vES) that allows Data Center Gateways (DCGWs) with two BGP instances (one for EVPN-MPLS and one for EVPN-VXLAN) to handle redundancy in VXLAN access networks. I-ESs support the RFC 7432 multi-homing functions, including single-active and all-active, ESI-label based split-horizon filtering, Designated Forwarder (DF) election, aliasing, and backup functions on remote EVPN-MPLS PEs.

The chapter EVPN-MPLS Interconnect for EVPN-VXLAN VPLS Services describes how VPLS services with two BGP instances are configured and describes a redundant mechanism referred to as Multi-homed anycast configuration for dual BGP-instance VPLS services. The use of I-ESs is recommended over this anycast configuration.

In addition to the EVPN multi-homing features, the main advantages of the I-ES solution compared to the redundant solution (described in Anycast Redundant Solution for Dual BGP-instance Services) are as follows:

- The use of I-ES for redundancy in dual BGP-instance services allows local SAPs on the DCGWs. This is not supported in the anycast solution.
- P2MP mLDP can be provisioned to transport Broadcast, Unknown unicast, and Multicast (BUM) traffic between DCs that use I-ES, without any risk of packet duplication. As described in The use of provider tunnels on multi-homed anycast solutions, packet duplication may occur in the anycast DCGW solution when mLDP is used in the WAN.

When EVPN-MPLS networks are interconnected to EVPN-VXLAN networks, the I-ES concept and procedures apply only to the access VXLAN network; the EVPN-MPLS network does not modify its existing behavior compared to any other ES.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Configuration

Figure 115: EVPN-MPLS interconnect for EVPN-VXLAN - BGP topology shows the topology and infrastructure configuration, which are the same as in chapter EVPN-MPLS Interconnect for EVPN-VXLAN VPLS Services. Read that chapter to see how the PEs are configured at port, IS-IS, and base BGP level.

Figure 115: EVPN-MPLS interconnect for EVPN-VXLAN - BGP topology

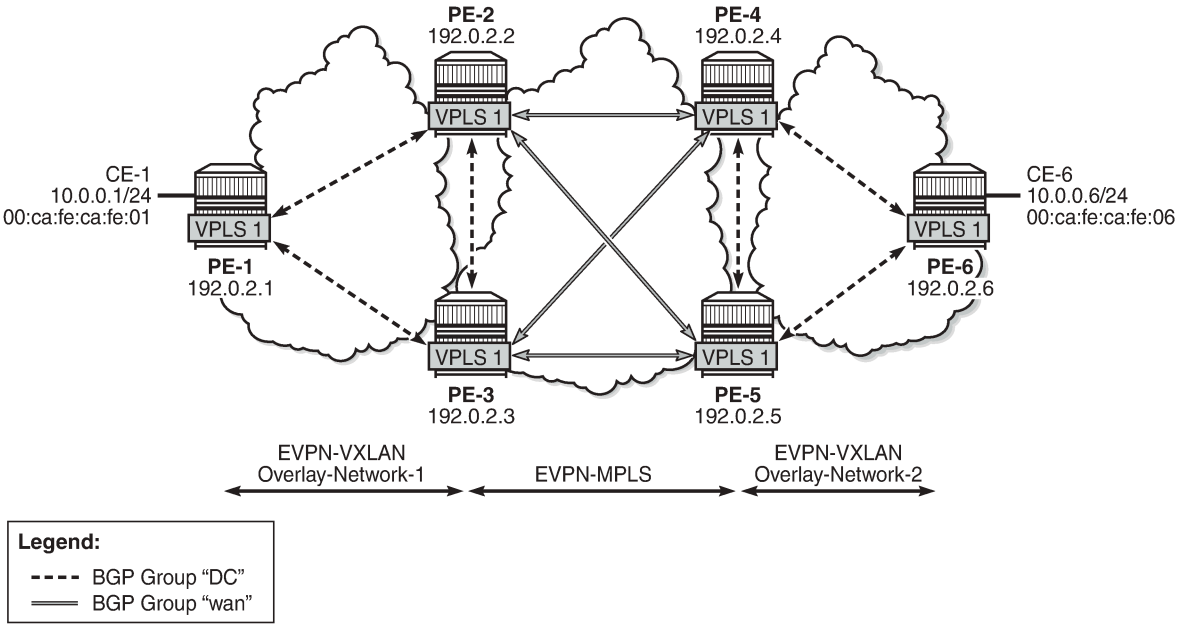

PE-1, PE-2, and PE-3 simulate a data center (DC), shown as Overlay-Network-1, where PE-2 and PE-3 are DCGWs. In the same way, PE-4, PE-5, and PE-6 simulate a remote DC, Overlay-Network-2. Inside each DC, EVPN-VXLAN is used and the two DCGW pairs are connected by EVPN-MPLS. CE-1 and CE-6 are end-to-end connected by EVPN without any VLAN or Pseudowire (PW) hand-off, maintaining all the EVPN advantages across the DC Interconnect (DCI) network.

Interconnect Ethernet Segment (I-ES) configuration

After the base infrastructure is configured (interfaces, IGP, LDP in the core, and BGP EVPN peering sessions, as per Figure 115: EVPN-MPLS interconnect for EVPN-VXLAN - BGP topology), two I-ESs configured on the DCGWs show the use of the Interconnect Ethernet Segments.

The I-ES "I-ES231" is configured on PE-2 and PE-3 as follows:

```
# on PE-2:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "I-ES231" {
                        admin-state enable
                        type virtual
                        esi 00:23:23:23:23:23:23:00:00:01
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        multi-homing-mode all-active
                        df-election {
                            service-carving-mode manual
                            manual {
                                evi 101 {
                                    end 200
                                }
                                preference {
                                    mode non-revertive
                                    value 150
                                }
                            }
                        }
                        association {
                            network-interconnect-vxlan 1 {
                                virtual-ranges {
                                    service-id 1 {
                                        end 100
                                    }
                                    service-id 101 {
                                        end 200
                                    }
                                }
                            }
                        }
                    }
```

```
# on PE-3:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "I-ES231" {
                        admin-state enable
                        type virtual
                        esi 00:23:23:23:23:23:23:00:00:01
                        multi-homing-mode all-active
                        df-election {
                            service-carving-mode manual
                            manual {
                                evi 101 {
                                    end 200
                                }
                                preference {
                                    mode non-revertive
                                    value 50
                                }
                            }
                        }
                        association {
                            network-interconnect-vxlan 1 {
                                virtual-ranges {
                                    service-id 1 {
                                        end 100
                                    }
                                    service-id 101 {
                                        end 200
                                    }
                                }
                            }
                        }
                    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

On PE-1 and PE-2, the preceding configuration associates I-ES "I-ES231" with the VXLAN instance 1 in services contained in the range VPLS 1 to 100 and 101 to 200. The I-ES is modeled as a virtual ES, where:

- Two commands are needed within the ethernet-segment context: **network-interconnect-vxlan** and **virtual-ranges service-id <svc-id> end <svc-id>**.
  - The **network-interconnect-vxlan** command identifies the VXLAN instance associated with the virtual ES. Only value 1 is supported in SR OS Release 21.2.R2.

```
*[ex:/configure service system bgp evpn ethernet-segment "vES-23" association]
A:admin@PE-1# network-interconnect-vxlan ?

 [network-interconnect-vxlan-id] <number>
 <number>  - <1>

    Vxlan instance id multi-homed with this Ethernet segment entry.
```

    The **network-interconnect-vxlan** command is rejected in non-virtual ESs:

```
*[ex:/configure service system bgp evpn ethernet-segment "ES-23" association]
A:admin@PE-1# network-interconnect-vxlan 1
MINOR: MGMT_CORE #2203: configure service system bgp evpn ethernet-segment "ES-23"
 association network-interconnect-vxlan 1 - Invalid element - network-interconnect-vxlan
 allowed only on virtual ethernet-segments
```

  - The **service-id** command associates the specific service range with the ES.
  - The other ES association options (port, lag, sdp, vc-id-range, dot1q, and qinq) cannot be combined with a **network-interconnect-vxlan** instance in an ES.
  - The rest of the ES configuration options are supported. The **source-bmac-lsb** is blocked because the I-ES cannot be associated with I-VPLS or PBB-Epipe services.
  - All the services with two BGP instances associate the VXLAN destinations and ingress VXLAN instance with the ES.
- Multiple services (for example, 1 to 200 in the CLI above) can be associated with the same ES.
  - Up to eight service ranges per VXLAN instance can be configured. Ranges may overlap within the same ES (and not between different ESs). In this example, two non-overlapping ranges are configured to show the service range configuration, although a single range containing all the services could have been configured.
  - The service range may be configured before the service is, and it can be changed on the fly.
- When the **network-interconnect-vxlan** I-ES is configured, the ES operational state depends exclusively on the ES admin state.
  - Because the I-ES is not associated with a physical port or SDP, when testing the non-revertive service-carving manual mode, an ethernet-segment admin-state disable/enable will result in the node sending its own administrative preference and "Do not preempt" (pref/DP) values, and taking over if pref/DP is higher than the current DF. This is because when the ES is enabled, the peer ES routes are not present at the EVPN application layer, so the PE will send its own admin pref/DP values. Therefore, for I-ESs, the non-revertive mode will only work for node failures. See the chapter for more information about the preference-based and non-revertive DF election modes.
- There are no restrictions in the service-carving mode supported by I-ESs. In this example, preference-based service-carving is configured, but modes auto and (non-preference-based) manual are also supported.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- As described in the Preference-based and Non-revertive EVPN DF Election chapter, the service-carving context is configured with an EVI range that will pick up the lowest preference value when electing a DF for the service, whereas the non-configured EVI services will pick up the highest value when electing a DF. In this example, this means that, of the services allowed in the I-ES, that is, 1 to 200, services 1 to 100 will elect the highest Preference PE as DF, whereas services 101 to 200 will elect the lowest Preference PE.

PE-4 and PE-5 are configured with I-ES "I-ES451". The configuration of I-ES451 is similar to that of I-ES231; only single-active mode is configured, instead of all-active mode.

```
# on PE-4:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "I-ES451" {
                        admin-state enable
                        type virtual
                        esi 00:45:45:45:45:45:45:00:00:01
                        multi-homing-mode single-active
                        df-election {
                            service-carving-mode manual
                            manual {
                                evi 101 {
                                    end 200
                                }
                                preference {
                                    mode non-revertive
                                    value 150
                                }
                            }
                        }
                        association {
                            network-interconnect-vxlan 1 {
                                virtual-ranges {
                                    service-id 1 {
                                        end 100
                                    }
                                    service-id 101 {
                                        end 200
                                    }
                                }
                            }
                        }
                    }
                }
```

```
# on PE-5:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "I-ES451" {
                        admin-state enable
                        type virtual
                        esi 00:45:45:45:45:45:45:00:00:01
                        multi-homing-mode single-active
                        df-election {
                            service-carving-mode manual
                            manual {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        evi 101 {
                            end 200
                        }
                        preference {
                            mode non-revertive
                            value 50
                        }
                    }
                }
                association {
                    network-interconnect-vxlan 1 {
                        virtual-ranges {
                            service-id 1 {
                                end 100
                            }
                            service-id 101 {
                                end 200
                            }
                        }
                    }
                }
            }
        }
```

In this example, VPLS 1 will be configured and associated with the preceding I-ESs. Figure 116: VPLS service and association with I-ESs shows an example of VPLS 1 and how it is associated with the I-ESs.

Figure 116: VPLS service and association with I-ESs

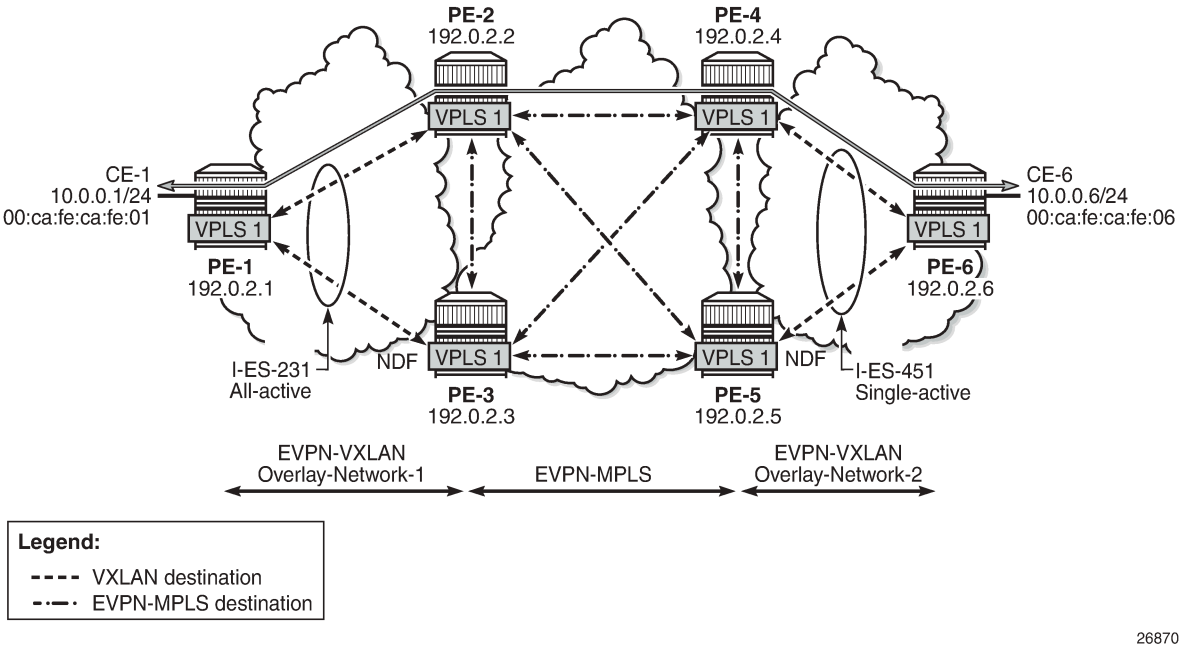

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The configuration of VPLS 1 for PE-1, PE-2, and PE-3 is as follows. VPLS 101 is also configured in all the PEs in a similar way as VPLS 1, but not shown here. Also, the VPLS 1 configuration on the rest of the PEs is equivalent to the one in PE-1, PE-2, and PE-3, as follows:

```
# on PE-1:
configure {
    service {
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            vxlan {
                instance 1 {
                    vni 1
                }
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 1
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
            sap 1/2/1:1 {
            }
        }
```

```
# on PE-2:
configure {
    service {
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            vxlan {
                instance 1 {
                    vni 1
                }
            }
            bgp 1 {
                route-distinguisher "192.0.2.2:1"
            }
            bgp 2 {
                route-distinguisher "192.0.2.2:2"
            }
            bgp-evpn {
                evi 1
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
                mpls 2 {
                    admin-state enable
                    ingress-replication-bum-label true
                    ecmp 2
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
```

```
# on PE-3:
configure {
    service {
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            vxlan {
                instance 1 {
                    vni 1
                }
            }
            bgp 1 {
                route-distinguisher "192.0.2.3:1"
            }
            bgp 2 {
                route-distinguisher "192.0.2.3:2"
            }
            bgp-evpn {
                evi 1
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
                mpls 2 {
                    admin-state enable
                    ingress-replication-bum-label true
                    ecmp 2
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
        }
```

As in the case of any other ESs, the association of instance and service is based on the ES configuration and there is no extra configuration required at the service level to make that association. The existing **show** commands that are used to check the status of the ES can be used to check the I-ESs. For example, on I-ES231:

```
[/]
A:admin@PE-2# show service system bgp-evpn ethernet-segment name "I-ES231" all

===============================================================================
Service Ethernet Segment
===============================================================================
Name                    : I-ES231
Eth Seg Type            : Virtual
Admin State             : Enabled            Oper State         : Up
ESI                     : 00:23:23:23:23:23:23:00:00:01
Multi-homing            : allActive          Oper Multi-homing  : allActive
ES SHG Label            : 524278
Source BMAC LSB         : <none>
VXLAN Instance Id       : 1
ES Activation Timer     : 3 secs (default)
Oper Group              : (Not Specified)
Svc Carving             : manual             Oper Svc Carving   : manual
Cfg Range Type          : lowest-pref

-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
DF Pref Election Information
-------------------------------------------------------------------------------
Preference     Preference     Last Admin Change        Oper Pref      Do No
Mode           Value                                   Value          Preempt
-------------------------------------------------------------------------------
non-revertive  150            05/03/2021 13:01:53      150            Enabled
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
EVI Ranges
-------------------------------------------------------------------------------
From                                  To
-------------------------------------------------------------------------------
101                                   200
-------------------------------------------------------------------------------
ISID Ranges: <none>
===============================================================================

===============================================================================
EVI Information
===============================================================================
EVI                 SvcId               Actv Timer Rem      DF
-------------------------------------------------------------------------------
1                   1                   0                   yes
101                 101                 0                   no
-------------------------------------------------------------------------------
Number of entries: 2
===============================================================================

-------------------------------------------------------------------------------
DF Candidate list
-------------------------------------------------------------------------------
EVI                                   DF Address
-------------------------------------------------------------------------------
1                                     192.0.2.2
1                                     192.0.2.3
101                                   192.0.2.2
101                                   192.0.2.3
-------------------------------------------------------------------------------
Number of entries: 4
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
---snip---

===============================================================================
Vxlan Instance Service Ranges
===============================================================================
Svc Range Start          Svc Range End            Last Changed
-------------------------------------------------------------------------------
1                        100                      05/03/2021 13:01:53
101                      200                      05/03/2021 13:01:53
-------------------------------------------------------------------------------
Number of Entries: 2
===============================================================================
```

The **show service id 1 vxlan instance 1 oper-flags** command shows the status of a VXLAN instance in the service. A service VXLAN instance will raise the oper-flag **MhStandby** (multi-homing standby) due to any of the following reasons:

- The PE is (single-active) non-Designated Forwarder (NDF) for that I-ES.
- The VXLAN service is added to the I-ES and either the ES is disabled or **bgp-evpn>mpls** is disabled in all the services included in the ES.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

For example, because PE-5 is an NDF in I-ES451, the MhStandby flag will show "true":

```
[/]
A:admin@PE-5# show service id 1 vxlan instance 1 oper-flags

===============================================================================
VPLS VXLAN oper flags
===============================================================================
MhStandby                          : true
===============================================================================
```

EVPN route handling in dual BGP-instance VPLSs with I-ES

The configuration of I-ESs on DCGWs with two BGP instances has the following impact on the advertisement and process of the BGP-EVPN routes:

- EVPN MAC/IP routes:
  - MAC/IP routes received on the EVPN-MPLS BGP instance will be re-advertised to the EVPN-VXLAN BGP instance with the ESI set to zero in SR OS Release 21.2.R2.
  - EVPN-VXLAN PE/NVEs (Network Virtual Edge devices) in the DC will receive the same MAC address from two (or more) different MAC/IP routes from the DCGWs. The EVPN-VXLAN PE/NVEs will perform regular EVPN MAC/IP route selection.
  - MAC/IP routes received on the EVPN-VXLAN BGP instance will be re-advertised to the EVPN-MPLS BGP instance with the configured non-zero I-ESI value, assuming the VXLAN instance is not in the MhStandby operational state. MAC/IP routes received on the EVPN-VXLAN BGP instance will be dropped if the VXLAN instance is in the MhStandby state.
  - EVPN-MPLS PEs in the WAN will receive the same MAC address from two (or more) DCGWs, set with the same ESI. EVPN-MPLS PEs will perform regular aliasing and backup functions.
- ES routes are exchanged for the I-ES. They should be sent only to the MPLS network and not to the VXLAN side. This can be achieved by using router policies. In any case, because ES routes use an ES-import route-target extended community, they should not be imported by VXLAN PEs.
- Auto-discover per ES (AD per-ES) and AD per-EVI routes are also advertised for the I-ES. They should be sent only to the MPLS network and not to the VXLAN network. As for ES routes, router policies can be used to prevent AD routes being sent to VXLAN peers.

Required BGP policies to avoid control plane loops

Usually, the use of router policies is required when I-ESs are used for redundancy, to avoid control plane loops with MAC/IP routes. The control plane loops to be avoided are as follows:

1. Loops created by remote MAC addresses (learned on remote PE SAPs):
   a. Remote EVPN-MPLS MAC/IP routes are re-advertised into EVPN-VXLAN with a Site of Origin (SOO) extended community (added by a BGP peer or vsi-export policy) identifying the DCGW pair. The other DCGW in the pair will drop EVPN-VXLAN MAC routes tagged with the self SOO. Router policies to add SOO and drop routes received with self SOO are needed.
   b. Also, when remote EVPN-VXLAN MAC/IP routes are re-advertised into EVPN-MPLS, the DCGWs will automatically drop EVPN-MPLS MAC/IP routes received with their own non-zero I-ESI. No router policies are needed for this.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

2. Loops created by local SAP MAC addresses:
   a. Local SAP MACs are learned and MAC/IP routes are advertised into both BGP instances. The MAC/IP routes advertised in the EVPN-VXLAN instance will be dropped by the peer based on the SOO router policies, as described in (1a) above, and DCGW local MACs will always be learned over the EVPN-MPLS destinations between the DCGWs.
   b. Because only EVPN-MPLS destinations exist between the DCGWs, EVPN-VXLAN MAC/IP and IMET routes exchanged between the DCGWs will be discarded and EVPN-VXLAN destinations will not be created between them.

As an example, the following BGP peer policies on PE-2 and PE-3 achieve the goals described above (similar policies would be configured on PE-4 and PE-5) and summarized as follows:

- Avoid sending service VXLAN routes to MPLS peers, and service MPLS routes to VXLAN peers.
- Avoid sending AD and ES routes to VXLAN peers.
- Add SOO to VXLAN routes to be sent to the ES peer.
- Drop VXLAN routes received from the ES peer.

```
# on PE-2, PE-3:
configure {
    policy-options {
        community "SOO-DCGW-23" {
            member "origin:64500:23" { }
        }
        community "mpls" {
            member "bgp-tunnel-encap:MPLS" { }
        }
        community "vxlan" {
            member "bgp-tunnel-encap:VXLAN" { }
        }
```

The following policy prevents the router from sending service VXLAN routes to MPLS peers:

```
        policy-statement "allow only mpls" {
            entry 10 {
                from {
                    family [evpn]
                    community {
                        name "vxlan"
                    }
                }
                action {
                    action-type reject
                }
            }
        }
```

The following policy makes sure the router exports only routes that include the VXLAN encapsulation:

```
        policy-statement "allow only vxlan" {
            entry 10 {
                from {
                    family [evpn]
                    community {
                        name "vxlan"
                    }
                }
                action {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    action-type accept
                }
            }
            default-action {
                action-type reject
            }
        }
```

The following import policy avoids importing routes with self SOO:

```
        policy-statement "drop SOO-DCGW-23" {
            entry 10 {
                from {
                    family [evpn]
                    community {
                        name "SOO-DCGW-23"
                    }
                }
                action {
                    action-type reject
                }
            }
        }
```

The following export policy adds SOO but only to VXLAN routes. This allows the peer to drop routes based on the SOO, without affecting the MPLS routes.

```
        policy-statement "add SOO to vxlan routes" {
            entry 10 {
                from {
                    family [evpn]
                    community {
                        name "vxlan"
                    }
                }
                action {
                    action-type accept
                    community {
                        add ["SOO-DCGW-23"]
                    }
                }
            }
            default-action {
                action-type accept
            }
        }
```

The BGP configuration for PE-2 and PE-3 is as follows:

```
# on PE-2:
configure {
    router "Base" {
        autonomous-system 64500
        router-id 192.0.2.2
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            family {
                ipv4 false
                evpn true
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            rapid-update {
                evpn true
            }
            group "dc" {
                type internal
                export {
                    policy ["allow only vxlan"]
                }
            }
            group "wan" {
                type internal
                export {
                    policy ["allow only mpls"]
                }
            }
            neighbor "192.0.2.1" {
                group "dc"
            }
            neighbor "192.0.2.3" {
                group "dc"
                import {
                    policy ["drop SOO-DCGW-23"]
                }
                export {
                    policy ["add SOO to vxlan routes"]
                }
            }
            neighbor "192.0.2.4" {
                group "wan"
            }
            neighbor "192.0.2.5" {
                group "wan"
            }
```

```
# on PE-3:
configure {
    router "Base" {
        autonomous-system 64500
        router-id 192.0.2.3
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            family {
                ipv4 false
                evpn true
            }
            rapid-update {
                evpn true
            }
            group "dc" {
                type internal
                export {
                    policy ["allow only vxlan"]
                }
            }
            group "wan" {
                type internal
                export {
                    policy ["allow only mpls"]
                }
            }
            neighbor "192.0.2.1" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            group "dc"
        }
        neighbor "192.0.2.2" {
            group "dc"
            import {
                policy ["drop SOO-DCGW-23"]
            }
            export {
                policy ["add SOO to vxlan routes"]
            }
        }
        neighbor "192.0.2.4" {
            group "wan"
        }
        neighbor "192.0.2.5" {
            group "wan"
        }
```

Single-active multi-homing operation

When the I-ES is configured as **single-active** and **admin-state enabled** (assuming at least one service is associated), the DCGWs will send ES and AD routes as usual for any ES, and run DF election based on the ES routes, with the candidate list being pruned by the AD routes.

In Figure 116: VPLS service and association with I-ESs, PE-4 and PE-5 are configured with I-ES451, which is a single-active ES. The NDF for a service (PE-5 for VPLS 1 in the example) will perform the following tasks:

- The VXLAN instance on the NDF will enter the MhStandby state and will block ingress and egress traffic on the VXLAN destinations associated with the I-ES.

```
[/]
A:admin@PE-5# show service id 1 vxlan instance 1 oper-flags

===============================================================================
VPLS VXLAN oper flags
===============================================================================
MhStandby                           : true
===============================================================================
```

- MAC/IP routes and FDB process:
  - Advertised MAC/IP routes that are associated with the VXLAN instance are withdrawn.
  - Advertised MAC/IP routes corresponding to local SAP MAC addresses or EVPN-MPLS binding MAC addresses are withdrawn if they were advertised to the EVPN-VXLAN instance.
  - Received MAC/IP routes associated with the VXLAN instance are not installed in FDB. The MAC routes will show as "used" in the **show router bgp routes evpn mac** commands; however, only the MAC addresses received from MPLS (in particular from the ES peer) will be programmed. As an example, the following CLI output shows how MAC address 00:ca:fe:ca:fe:06 is learned on PE-4 (DF) and associated with the VXLAN destination to PE-6, whereas the MAC address is installed associated with an MPLS destination (remote ES) on PE-5 (NDF).

```
[/]
A:admin@PE-4# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
1          00:ca:fe:ca:fe:01 eES:                    Evpn     05/03/21 13:10:58
                             00:23:23:23:23:23:23:00:00:01
1          00:ca:fe:ca:fe:06 vxlan-1:                Evpn     05/03/21 13:10:58
                             192.0.2.6:1
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

```
[/]
A:admin@PE-5# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
1          00:ca:fe:ca:fe:01 eES:                    Evpn     05/03/21 13:10:58
                             00:23:23:23:23:23:23:00:00:01
1          00:ca:fe:ca:fe:06 eES:                    Evpn     05/03/21 13:10:58
                             00:45:45:45:45:45:45:00:00:01
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

- Inclusive Multicast Ethernet Tag (IMET) routes process:
  - IMET-Assisted Replication with replicator role (IMET-AR-R) routes are withdrawn if the VXLAN instance enters the MhStandby state. Only the DF will advertise the IMET-AR-R routes. For more information on AR, see chapter "Layer 2 Multicast Optimization for EVPN-VXLAN - Assisted Replication" in the Layer 2 Services and EVPN section of the *7450 ESS, 7750 SR, and 7950 XRS Advanced Configuration Guide - Part II*.
  - IMET-Ingress Replication advertisements (IMET-IR) routes, in case of NDF (or the MhStandby state), are controlled by the **config>service>vpls>bgp-evpn>vxlan# send-incl-mcast-ir-on-ndf** command.
    - By default, the command is enabled and the router will advertise IMET-IR routes even if the PE is NDF (MhStandby). This will attract BUM traffic (even if the NDF ends up dropping it); however, attracting BUM traffic will also speed up convergence in case of DF switchover. The command works for single-active and all-active.
    - If disabled, the router will withdraw the IMET-IR routes when the PE is NDF and will not attract BUM traffic.

In spite of not sending BUM or unicast traffic, the NDF for a service still creates the VXLAN bindings; however, they are not associated with any MAC addresses and they are flagged as non-multicast capable, or "-" in the Mcast column of the following command:

```
[/]
A:admin@PE-5# show service id 1 vxlan destinations

===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Egress VTEP, VNI
==============================================================================
Instance    VTEP Address                          Egress VNI  EvpnStatic Num
 Mcast       Oper State                            L2 PBR     SupBcasDom MACs
------------------------------------------------------------------------------
1           192.0.2.6                             1           evpn       0
 -           Up                                    No          No
------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 1
------------------------------------------------------------------------------
==============================================================================

==============================================================================
BGP EVPN-VXLAN Ethernet Segment Dest
==============================================================================
Instance  Eth SegId                      Num. Macs     Last Change
------------------------------------------------------------------------------
No Matching Entries
==============================================================================
```

The I-ES DF PE for the service (PE-4) will continue advertising IMET and MAC/IP routes for the associated VXLAN instance. Forwarding will also happen as usual on the DF VXLAN bindings. When the DF PE receives BUM traffic from VXLAN, it will send it, adding the egress ESI label if needed.

All-active multi-homing operation

The same considerations as in single-active for ES and AD routes and DF election apply to all-active multi-homing. In Figure 116: VPLS service and association with I-ESs, PE-2 and PE-3 are configured with I-ES231, which is an all-active ES. The NDF PE for a service (PE-3 for VPLS 1, in the example) will show the following behavior:

- The VXLAN instance on the NDF will not enter the MhStandby state because it will still forward unicast traffic:

```
[/]
A:admin@PE-3# show service id 1 vxlan instance 1 oper-flags

===============================================================================
VPLS VXLAN oper flags
===============================================================================
MhStandby                          : false
===============================================================================
```

- MAC/IP routes and FDB process: MAC/IP routes are received, installed, and advertised as in the DF router.
- IMET routes process:
  - As in the single-active case, IMET-AR-R routes are withdrawn on the NDF. Only the DF will advertise the IMET-AR-R routes.
  - Also, as in the single-active case, IMET-IR advertisement from the NDF will be controlled by the **config>service>vpls>bgp-evpn>vxlan# send-incl-mcast-ir-on-ndf** command. Advertising the IMET-IR route from the NDF will attract BUM traffic from the VXLAN PEs to the NDF, even though the unknown unicast traffic will be forwarded only when it is safe to do so. See section All-active multi-homing and unknown unicast forwarding on the NDF for more information about unknown unicast forwarding.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Contrary to the behavior in single-active multi-homing, in all-active, the NDF will forward unknown unicast to the VXLAN PEs as usual, but block broadcast and multicast in the upstream and downstream direction. In our example, the NDF for VPLS 1 (PE-3) will show the VXLAN destinations created as "U" (Unknown unicast) in the Mcast column of the **show service id 1 vxlan** command, as follows:

```
[/]
A:admin@PE-3# show service id 1 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
Instance    VTEP Address                          Egress VNI  EvpnStatic Num
 Mcast       Oper State                            L2 PBR     SupBcasDom MACs
-------------------------------------------------------------------------------
1           192.0.2.1                             1           evpn       1
 U           Up                                    No          No
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 1
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN-VXLAN Ethernet Segment Dest
===============================================================================
Instance  Eth SegId                      Num. Macs     Last Change
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

All-active multi-homing and unknown unicast forwarding on the NDF

The unknown unicast traffic will be transmitted on the (all-active multi-homing) NDF in the upstream and downstream directions only in those cases where there is no risk of packet duplication. The router considers there is no risk when transmitting an unknown unicast packet on the NDF if:

- Unknown unicast packet arrives without an ESI label.
- Unknown unicast packet arrives without a BUM label (label advertised by an IMET route as opposed to a MAC/IP route).
- Unknown unicast packet passes a MAC Source Address (MAC SA) suppression (MAC SA lookup does not yield an entry associated with the I-ES).

The following examples show how unknown unicast traffic is handled in all-active I-ESs.

Figure 117: All-active multi-homing and unknown unicast example 1 shows an example with two DCGWs where (all-active) I-ES-1 is defined.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 117: All-active multi-homing and unknown unicast example 1

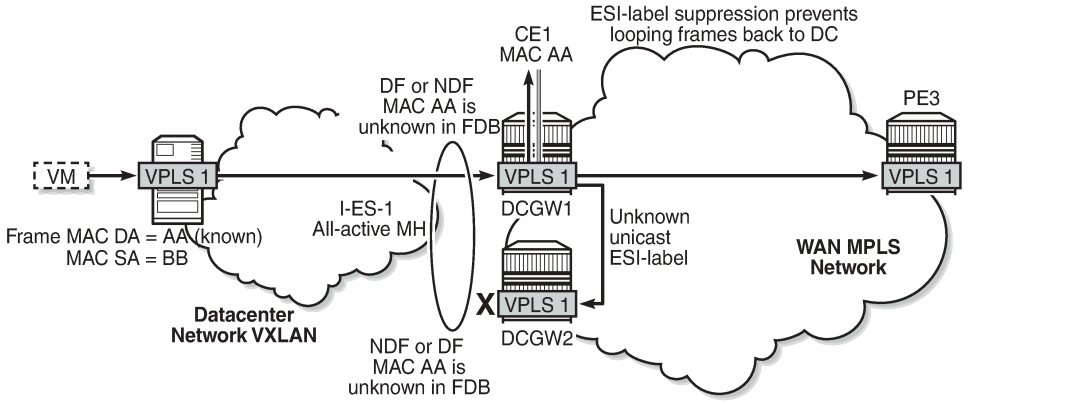

The VXLAN PE/NVE transmits known unicast traffic, whereas DCGW1 has not learned the MAC address yet. Regardless of the DCGW1 being DF or NDF, it will accept unknown unicast and will flood to local SAPs and EVPN destinations. When sending to DCGW2, the router will send the ESI label identifying the I-ES. DCGW2 will not send unknown traffic back to the DC due to the ESI-label suppression on the I-ES.

Figure 118: All-active multi-homing and unknown unicast example 2 shows a similar example where the VXLAN node sends known unicast with MAC Destination Address (MAC DA) "AA" to DCGW2.

Figure 118: All-active multi-homing and unknown unicast example 2

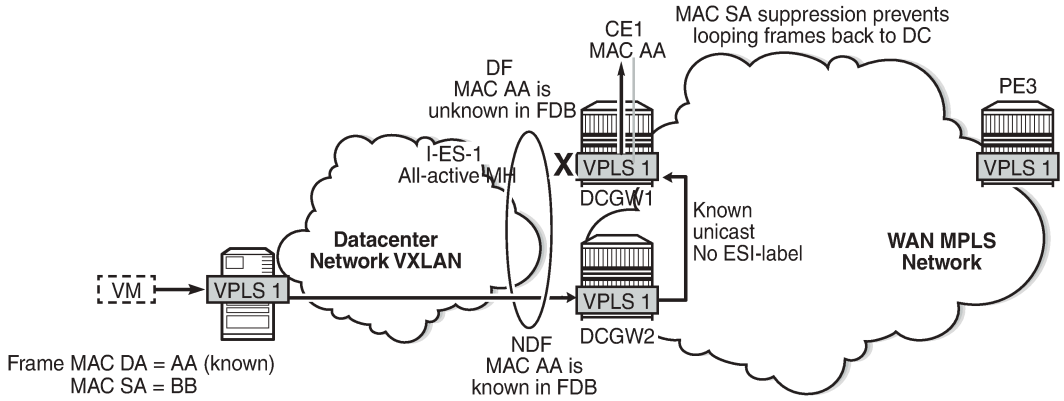

DCGW2 does a MAC lookup and sends the frame as known unicast to DCGW1 via the EVPN-MPLS destination. However, MAC AA is unknown in DCGW1 for some reason (such as FDB limit exceeded, SAP failure, and so on). In this case, DCGW1 will flood the frame to CE1 and not to the VXLAN network. Even though the frame is not coming with an ESI label, the DCGW1 router does a MAC SA suppression and will not send unknown unicast frames to the I-ES. MAC SA suppression means that the router will do a MAC SA lookup on the FDB and will suppress the flooding to the I-ES if the MAC SA is learned on the I-ES (as in Figure 118: All-active multi-homing and unknown unicast example 2).

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 119: All-active multi-homing and unknown unicast example 3 shows an example in which the NDF forwards "no-risk" unknown unicast traffic to avoid black-holes.

Figure 119: All-active multi-homing and unknown unicast example 3

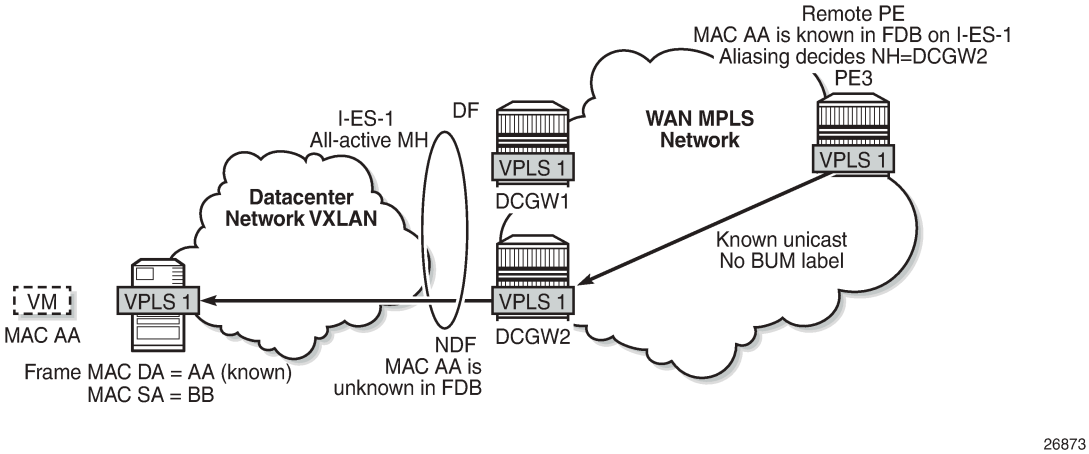

PE3 receives unicast traffic with MAC DA = AA. The MAC address is known in the FDB and associated with I-ES-1; therefore, because PE3 is configured to do aliasing to DCGW1 and DCGW2 (bgp-evpn>mpls# ecmp 2), a packet hash determines that it has to be sent to DCGW2 (NDF). The packet arrives at DCGW2 with a unicast label. DCGW2 does a lookup and MAC AA is unknown for some reason (such as FDB limit exceeded, MAC not learned yet, and so on). In this case, DCGW2 will forward the packet to the I-ES VXLAN bindings, even if it is NDF. This behavior avoids black-hole periods in the network for unicast traffic.

Finally, in some cases, the unknown unicast forwarding behavior on the NDF may cause some transient packet duplication that can be avoided by configuring the **send-incl-mcast-ir-on-ndf** command. The following example shows the use of this command to avoid transient packet duplication. Figure 120: All-active multi-homing and send-incl-mcast-ir-on-ndf true shows how transient packet duplication may occur with the default setting (**send-incl-mcast-ir-on-ndf true**).

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 120: All-active multi-homing and send-incl-mcast-ir-on-ndf true

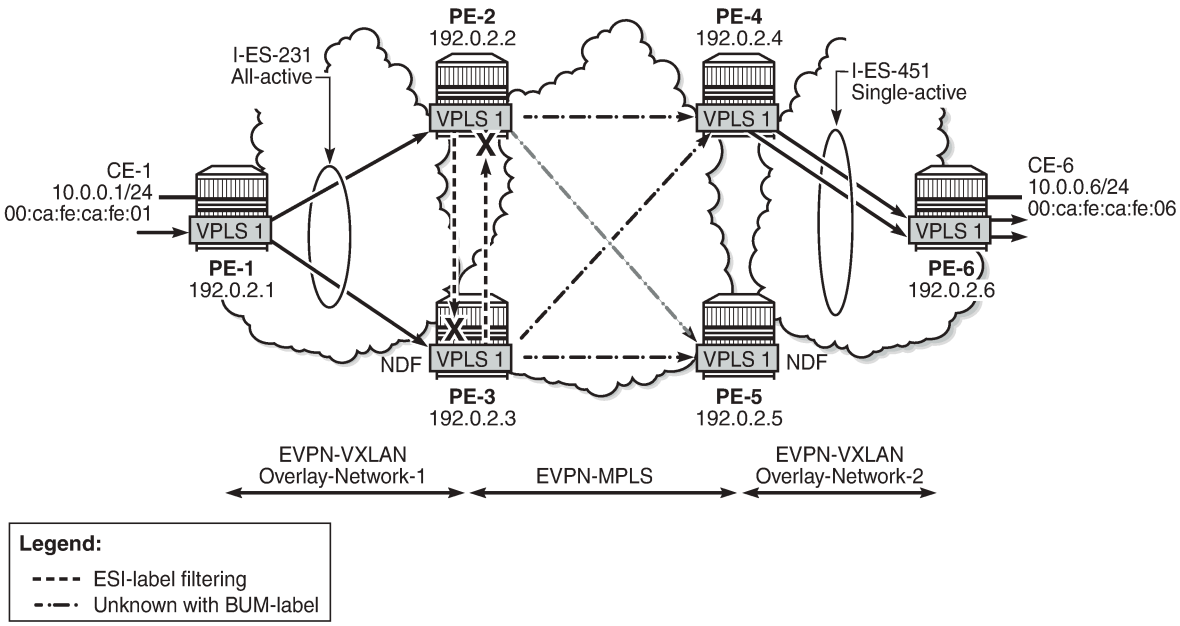

Transient packet duplication may occur when sending unknown unicast from CE-1 to CE-6, if **send-incl-mcast-ir-on-ndf true** is configured in PE-3 and PE-2. To show this, we clear the FDBs in all the PEs in the example as well as the ARP caches on the CEs.

The following command is executed in all the PEs and CEs:

```
[/]
A:admin@PE-1# clear service id 1 fdb all

[/]
A:admin@PE-1# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

The following command clears the ARP table of the VPRN instance (defined in PE-1 using a loop) simulating CE-1:

```
[/]
A:admin@PE-1# clear router 300 arp all

[/]
A:admin@PE-1# show router 300 arp

===============================================================================
ARP Table (Service: 300)
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
IP Address      MAC Address       Expiry    Type   Interface
-------------------------------------------------------------------------------
10.0.0.1        00:ca:fe:ca:fe:01 00h00m00s Oth[I] local
-------------------------------------------------------------------------------
No. of ARP Entries: 1
===============================================================================
```

When ICMP traffic is sent from CE-1 to CE-6, a duplicate entry occurs on CE-1:

```
[/]
A:admin@PE-1# ping 10.0.0.6 router-instance "VPRN 300"
PING 10.0.0.6 56 data bytes
64 bytes from 10.0.0.6: icmp_seq=1 ttl=64 time=13.2ms.
64 bytes from 10.0.0.6: icmp_seq=1 ttl=64, duplicate.
64 bytes from 10.0.0.6: icmp_seq=2 ttl=64 time=5.27ms.
64 bytes from 10.0.0.6: icmp_seq=3 ttl=64 time=5.25ms.
64 bytes from 10.0.0.6: icmp_seq=4 ttl=64 time=4.73ms.
64 bytes from 10.0.0.6: icmp_seq=5 ttl=64 time=4.80ms.

---- 10.0.0.6 PING Statistics ----
5 packets transmitted, 5 packets received, 1 duplicate
round-trip min = 4.73ms, avg = 6.66ms, max = 13.2ms, stddev = 3.29ms
```

This duplicate entry occurs because the packet gets to CE-6 twice and CE-6 sends two unicast ICMP reply messages back. From the CE-1 packet walkthrough:

- PE-1 floods the packet to PE-2 and PE-3 because the CE-6 MAC DA is unknown and it has VXLAN multicast destinations to them.
- PE-2 floods the unknown unicast packet to all the remote PEs because it is DF for I-ES231. PE-2 will add an ESI label when sending to PE-3, and a BUM label when sending to all of them.
- PE-3 is NDF for I-ES231, but it floods the packet because the I-ES is all-active and the unknown unicast packet is considered low risk. The packet arrives with no ESI label, no BUM label (in VXLAN, VNIs are the same for unicast and BUM), and the MAC SA suppression passes because the packet is coming from the I-ES and not from MPLS. PE-3 uses a BUM label when flooding the packet and an ESI label when sending to PE-2.
- PE-4 receives two unknown unicast packets and forwards both to PE-6.
- PE-5 does not forward because it is NDF. This is true regardless of the I-ES being single-active or all-active (if all-active, the packet will not be forwarded because it arrives with a BUM label).

This packet duplication situation is transient and it will stop as soon as the two MAC addresses are learned on the PEs. However, if needed, this situation can be avoided by configuring **send-incl-mcast-ir-on-ndf false**:

```
# on PE-2, PE-3:
configure {
    service {
        vpls "VPLS 1" {
            bgp-evpn {
                vxlan 1 {
                    send-incl-mcast-ir-on-ndf false
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

This command will make the NDF (PE-3) withdraw the IMET-IR route; therefore, PE-1 will only flood unknown unicast packets to the DF (PE-2). The following IMET-IR routes are received on PE-1: one route sent by DF PE-2 for VPLS 1 and two routes for VPLS 101.

```
[/]
A:admin@PE-1# show router bgp routes evpn incl-mcast
===============================================================================
 BGP Router ID:192.0.2.1        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Inclusive-Mcast Routes
===============================================================================
Flag  Route Dist.          OrigAddr
      Tag                  NextHop
-------------------------------------------------------------------------------
u*>i  192.0.2.2:1          192.0.2.2
      0                    192.0.2.2

u*>i  192.0.2.2:101        192.0.2.2
      0                    192.0.2.2

u*>i  192.0.2.3:101        192.0.2.3
      0                    192.0.2.3

-------------------------------------------------------------------------------
Routes : 3
===============================================================================
```

If a DF switchover occurs in the I-ES, the new DF would advertise the IMET-IR route and the new NDF would withdraw it.

After clearing FDBs and ARP caches again, the test is repeated with no packet duplication. Figure 121: All-active multi-homing and send-incl-mcast-ir-on-ndf false shows how PE-1 does not send unknown unicast to PE-3 (NDF) anymore and, therefore, there is no duplication.

```
[/]
A:admin@PE-1# ping 10.0.0.6 router-instance "VPRN 300"
PING 10.0.0.6 56 data bytes
64 bytes from 10.0.0.6: icmp_seq=1 ttl=64 time=15.3ms.
64 bytes from 10.0.0.6: icmp_seq=2 ttl=64 time=5.32ms.
64 bytes from 10.0.0.6: icmp_seq=3 ttl=64 time=5.33ms.
64 bytes from 10.0.0.6: icmp_seq=4 ttl=64 time=5.44ms.
64 bytes from 10.0.0.6: icmp_seq=5 ttl=64 time=4.98ms.

---- 10.0.0.6 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 4.98ms, avg = 7.26ms, max = 15.3ms, stddev = 4.00ms
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 121: All-active multi-homing and send-incl-mcast-ir-on-ndf false

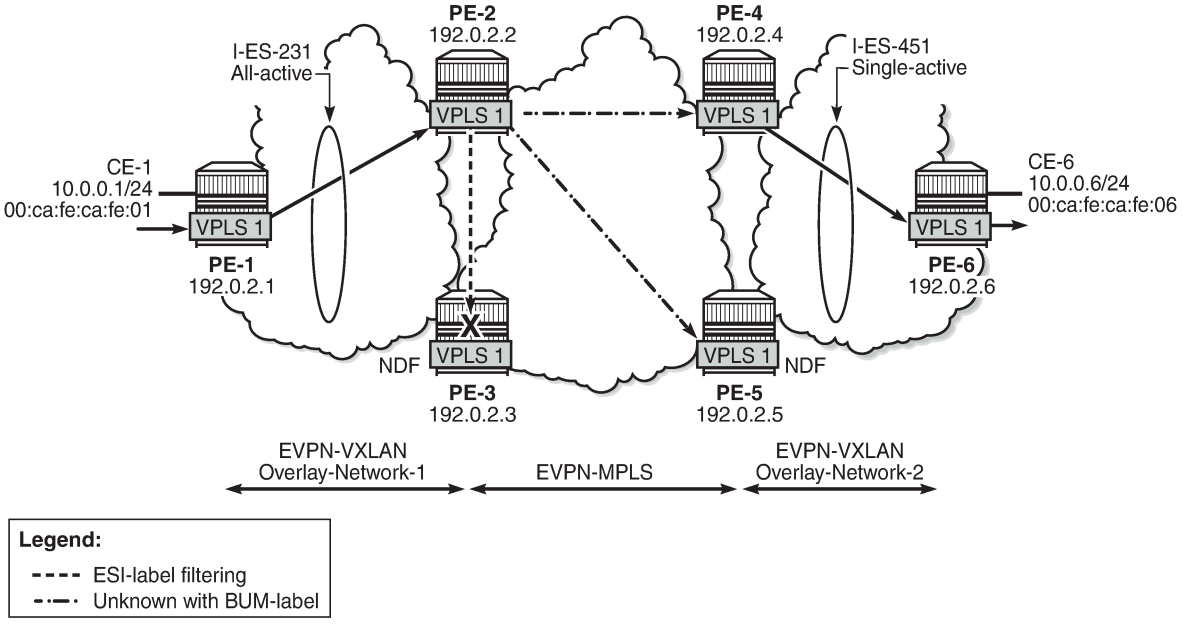

Local SAPs and provider tunnels along with I-ES

As described in the Overview section, the main advantages of the I-ES solution over the anycast redundant solution for dual BGP-instance services are the support of local SAPs and P2MP mLDP trees without packet duplication. This section shows the configuration of local SAPs and provider tunnels along with I-ES in VPLS services. The local SAPs can, at the same time, belong to an ES or a vES.

As an example, VPLS 1 on PE-2 is reconfigured as follows (similar configuration on PE-3, with provider tunnel also configured on PE-4 and PE-5):

```
# on PE-2:
configure {
    service {
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            vxlan {
                instance 1 {
                    vni 1
                }
            }
            bgp 1 {
                route-distinguisher "192.0.2.2:1"
            }
            bgp 2 {
                route-distinguisher "192.0.2.2:2"
            }
            bgp-evpn {
                evi 1
                vxlan 1 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    admin-state enable
                    vxlan-instance 1
                }
                mpls 2 {
                    admin-state enable
                    ingress-replication-bum-label true
                    ecmp 2
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            sap lag-1:1 {
            }
            provider-tunnel {
                inclusive {
                    admin-state enable
                    owner bgp-evpn-mpls
                    root-and-leaf true
                    mldp
                }
            }
```

To have EVPN multi-homing from a CE locally connected to PE-2 and PE-3, an additional ES is configured on PE-2 and PE-3 that will include the local SAPs in VPLS 1, as follows:

```
# on PE-2, PE-3:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "vES232" {
                        admin-state enable
                        type virtual
                        esi 00:23:23:23:23:23:23:00:00:02
                        multi-homing-mode all-active
                        association {
                            lag "lag-1" {
                                virtual-ranges {
                                    dot1q {
                                        q-tag 1 {
                                            end 1
                                        }
                                    }
                                }
                            }
                        }
```

Troubleshooting and debugging

Common troubleshooting commands to operate dual BGP-instance VPLS services are in the corresponding section of EVPN-MPLS Interconnect for EVPN-VXLAN VPLS Services. Also, ES and virtual ES can be troubleshot by using the commands described in chapter EVPN for MPLS Tunnels.

As well, the following **show** commands are specific to the use of I-ES in the router:

```
[/]
A:admin@PE-2# show service id 1 vxlan instance 1 oper-flags
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
VPLS VXLAN oper flags
===============================================================================
MhStandby                           : false
===============================================================================
```

```
[/]
A:admin@PE-2# show service vxlan-instance-using ethernet-segment

===============================================================================
VXLAN Ethernet-Segment Information
===============================================================================
SvcId       VXLAN Instance      ES Name                         Status
-------------------------------------------------------------------------------
1           1                   I-ES231                         DF
101         1                   I-ES231                         NDF
===============================================================================
```

```
[/]
A:admin@PE-2# show service vxlan-instance-using ethernet-segment name "I-ES231"

===============================================================================
VXLAN Ethernet-Segment Information
===============================================================================
SvcId                            VXLAN Instance                 Status
-------------------------------------------------------------------------------
1                                1                              DF
101                              1                              NDF
===============================================================================
```

Conclusion

Based on *draft-ietf-bess-dci-evpn-overlay*, SR OS supports the connectivity of Layer 2 EVPN-VXLAN services to an EVPN-MPLS network. This chapter complements the chapter EVPN-MPLS Interconnect for EVPN-VXLAN VPLS Services by describing how redundancy can be improved with the use of I-ES multi-homing, a concept standardized in *draft-ietf-bess-dci-evpn-overlay*.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

EVPN Interconnect Ethernet Segments in Dual EVPN-VXLAN Instance VPLS Services

This chapter provides information about EVPN Interconnect Ethernet Segments in Dual EVPN-VXLAN Instance VPLS Services.

Topics in this chapter include:

Applicability

The information and configuration in this chapter are based on SR OS Release 21.7.R1. EVPN multi-homing on dual VXLAN instance VPLS services is supported on SR OS Release 19.10.R1, and later.

Overview

Some service providers are deploying large Data Centers (DCs) where SR OS routers are used as leaf switches in a VXLAN fabric. In those cases, all-active multi-homing can provide redundancy and maximize the bandwidth utilization.

SR OS supports Interconnect Ethernet Segments (I-ESs) for VXLAN as per RFC 9014. Chapter EVPN Interconnect Ethernet Segments (I-ESs) describes how I-ESs allow Data Center Gateways (DCGWs) with two BGP instances (one for EVPN-MPLS and one for EVPN-VXLAN) to handle redundancy in VXLAN access networks, as supported in SR OS 15.0.R4, and later.

This chapter describes similar scenarios with EVPN-VXLAN in the core network instead of EVPN-MPLS. The following scenarios are supported with I-ES in VXLAN instance 1:

- dual instance VPLS with two EVPN-VXLAN instances
- dual instance VPLS with one EVPN-VXLAN instance and one static VXLAN instance
- dual instance VPLS with one EVPN-VXLAN instance and one EVPN-MPLS instance

The first two of these scenarios are described in this chapter.

CLI

Figure 122: Sample topology shows VPLS 1 with different EVPN-VXLAN instances: VXLAN instance 1 in DC 1 (and DC2) and VXLAN instance 2 in the WAN.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 122: Sample topology

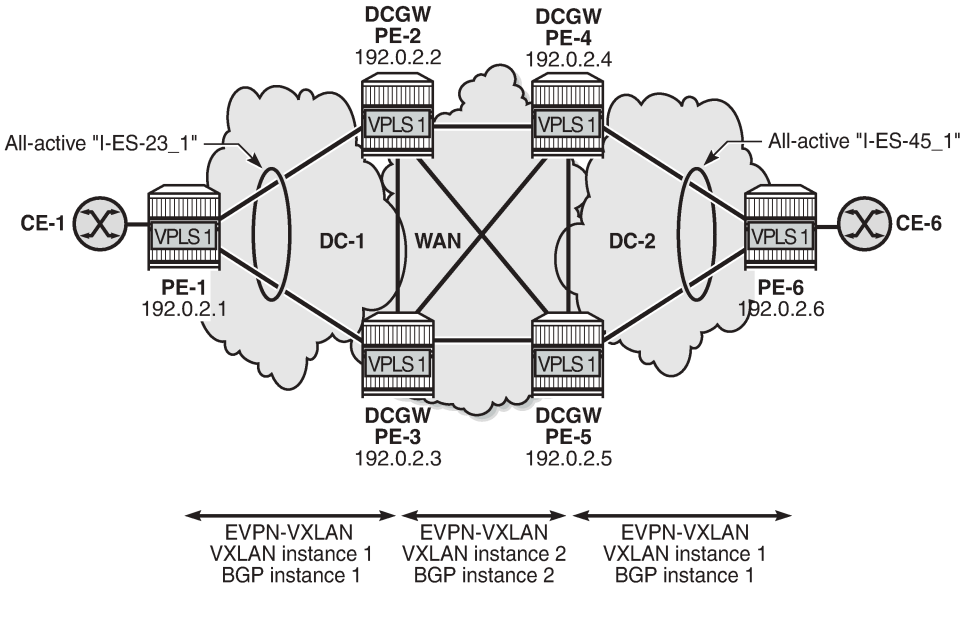

On DCGW PE-2, the following all-active I-ES is configured for VXLAN instance 1 and service id 1:

```
# on DCGW PE-2:
configure {
    service {
        system {
            bgp-auto-rd-range {
                ip-address 192.0.2.2
                community-value {
                    start 1
                    end 1000
                }
            }
            bgp {
                evpn {
                    ethernet-segment "I-ES-23_1" {
                        admin-state enable
                        type virtual
                        esi 00:23:23:23:23:23:23:00:00:01
                        multi-homing-mode all-active
                        df-election {
                            service-carving-mode manual
                            manual {
                                evi 1 {
                                    end 1
                                }
                                preference {
                                    value 100
                                }
                            }
                        }
                        association {
                            network-interconnect-vxlan 1 {
                                virtual-ranges {
                                    service-id 1 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                                        end 1
                                    }
                                }
                            }
                        }
                    }
                }
```

The following command configures VPLS 1 with dual EVPN-VXLAN instance. VXLAN instance 1 is a member of the I-ES and VXLAN instance 2 is configured with **mh-mode network** and **routes>auto-disc>advertise true**:

```
# on DCGW PE-2:
configure {
    service {
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            vxlan {
                instance 1 {
                    vni 11
                    rx-discard-on-ndf bum
                }
                instance 2 {
                    vni 12
                }
            }
            bgp 1 {
                route-distinguisher auto-rd
                route-target {
                    export "target:64500:11"
                    import "target:64500:11"
                }
            }
            bgp 2 {
                route-distinguisher auto-rd
                route-target {
                    export "target:64500:12"
                    import "target:64500:12"
                }
            }
            bgp-evpn {
                evi 1
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                    default-route-tag 0xb
                    ecmp 2
                    routes {
                        auto-disc {
                            advertise true
                        }
                    }
                }
                vxlan 2 {
                    admin-state enable
                    vxlan-instance 2
                    default-route-tag 0xc
                    ecmp 2
                    mh-mode network
                    routes {
                        auto-disc {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                            advertise true
                        }
                    }
                }
            }
            sap 1/2/1:1 {
            }
        }
```

By default, the multi-homing mode for EVPN-VXLAN is access, but for VXLAN instance 2, it is modified to **mh-mode network**. The following error is raised when attempting to configure VXLAN instance 1—as a member of an I-ES—with **mh-mode network**:

```
[ex:/configure service vpls "VPLS 1" bgp-evpn vxlan 1]
A:admin@PE-2# mh-mode network

*[ex:/configure service vpls "VPLS 1" bgp-evpn vxlan 1]
A:admin@PE-2# commit
MINOR: MGMT_CORE #4001: configure service vpls "VPLS 1" bgp-evpn vxlan 1 mh-mode - mh-mode
 network not supported when vxlan instance is member of ethernet-segment - configure service
 system bgp evpn ethernet-segment "I-ES-23_1" association network-interconnect-vxlan 1 virtual-
ranges service-id 1
```

With **mh-mode network** configured, it is mandatory to configure **routes>auto-disc>advertise true**; for **mh-mode access**, it is optional. When **routes>auto-disc>advertise** is enabled in an access instance associated to an I-ES, AD per-ES/EVI routes and MAC/IP routes are advertised for the I-ES.

The following AD per-EVI route is sent by DCGW PE-2:

```
10 2021/09/29 17:53:31.575 CEST MINOR: DEBUG #2001 Base Peer 1: 192.0.2.1
"Peer 1: 192.0.2.1: UPDATE
Peer 1: 192.0.2.1 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 81
    Flag: 0x90 Type: 14 Len: 36 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
        Type: EVPN-AD Len: 25 RD: 192.0.2.2:1 ESI: 00:23:23:23:23:23:23:00:00:01,
                      tag: 0 Label: 11 (Raw Label: 0xb) PathId:
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        origin:64500:23
        target:64500:11
        bgp-tunnel-encap:VXLAN
"
```

For MAC routes and their ESI value for an access VXLAN instance, the following redistribution considerations apply.

- With **mh-mode access** and **routes>auto-disc>advertise true** configured, MAC routes are redistributed from the instance network to the instance access with the I-ESI if present, regardless of the original ESI.
- With **mh-mode access** and **routes>auto-disc>advertise false**, MAC routes are redistributed with zero ESI, regardless of the original ESI.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The following EVPN-MAC route is sent by DCGW PE-2 with I-ESI 00:23:23:23:23:23:23:00:00:01 of "I-ES-23_1":

```
20 2021/09/29 17:53:31.577 CEST MINOR: DEBUG #2001 Base Peer 1: 192.0.2.1
"Peer 1: 192.0.2.1: UPDATE
Peer 1: 192.0.2.1 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 89
    Flag: 0x90 Type: 14 Len: 44 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
        Type: EVPN-MAC Len: 33 RD: 192.0.2.2:1 ESI: 00:23:23:23:23:23:23:00:00:01,
         tag: 0, mac len: 48 mac: 00:ca:fe:ca:fe:06, IP len: 0, IP: NULL, label1: 11
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        origin:64500:23
        target:64500:11
        bgp-tunnel-encap:VXLAN
"
```

The following ES route is sent by DCGW PE-2:

```
29 2021/09/29 17:53:31.661 CEST MINOR: DEBUG #2001 Base Peer 1: 192.0.2.5
"Peer 1: 192.0.2.5: UPDATE
Peer 1: 192.0.2.5 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 71
    Flag: 0x90 Type: 14 Len: 34 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
        Type: EVPN-ETH-SEG Len: 23 RD: 192.0.2.2:0
              ESI: 00:23:23:23:23:23:23:00:00:01, IP-Len: 4 Orig-IP-Addr: 192.0.2.2
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        df-election::DF-Type:Preference/DP:0/DF-Preference:100/AC:1
        target:23:23:23:23:23:23
"
```

The following commands are not supported when **mh-mode network** is configured:

- proxy-arp/nd
- assisted replication
- source-vtep-security

Attempting to enable these unsupported commands while a BGP-EVPN VXLAN instance has **mh-mode network** triggers error messages, as follows:

- proxy-arp

```
*[ex:/configure service vpls "VPLS 1" proxy-arp]
A:admin@PE-2# commit
MINOR: SVCMGR #12: configure service vpls "VPLS 1" bgp-evpn vxlan 2 mh-mode - Inconsistent
 Value error - mh-mode network not supported with proxy-arp - configure service vpls "VPLS
 1" proxy-arp
MINOR: MGMT_CORE #4001: configure service vpls "VPLS 1" - multiple bgp-evpn instances not
 supported with proxy-arp
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- proxy-nd

```
*[ex:/configure service vpls "VPLS 1" proxy-nd]
A:admin@PE-2# commit
MINOR: SVCMGR #12: configure service vpls "VPLS 1" bgp-evpn vxlan 2 mh-mode - Inconsistent
 Value error - mh-mode network not supported with proxy-nd - configure service vpls "VPLS 1"
 proxy-nd
MINOR: MGMT_CORE #4001: configure service vpls "VPLS 1" - multiple bgp-evpn instances not
 supported with proxy-nd
```

- assisted replication

```
[ex:/configure service vpls "VPLS 1" vxlan instance 2 assisted-replication]
A:admin@PE-2# replicator

*[ex:/configure service vpls "VPLS 1" vxlan instance 2 assisted-replication]
A:admin@PE-2# commit
MINOR: MGMT_CORE #4001: configure service vpls "VPLS 1" bgp-evpn vxlan 2 vxlan-instance -
 replicator role on vxlan instance not supported when it is in use by bgp-evpn with mh-mode
 network - configure service vpls "VPLS 1" vxlan instance 2 assisted-replication replicator
```

- source-vtep-security

```
[ex:/configure service vpls "VPLS 1" vxlan instance 2]
A:admin@PE-2# source-vtep-security

*[ex:/configure service vpls "VPLS 1" vxlan instance 2]
A:admin@PE-2# commit
MINOR: MGMT_CORE #4001: configure service vpls "VPLS 1" bgp-evpn vxlan 2 vxlan-instance -
 source-vtep-security on vxlan instance not supported when it is in use by bgp-evpn with mh-
mode network - configure service vpls "VPLS 1" vxlan instance 2 source-vtep-security
```

Local bias

When EVPN-VXLAN is used in the instance network of a dual-instance VPLS service, local bias—as described in RFC 8365—is used for split horizon in all-active I-ESs. In VXLAN, there is no multicast label or multicast BMAC, so BUM traffic is identified by the MAC destination address. The modified forwarding rules for the I-ES-sourced BUM traffic for ingress PE and egress PE are as follows:

- ingress PE
  - The Non-Designated Forwarder (NDF) must discard BUM traffic, so one of the following two commands must be configured in VXLAN instance 1.
    - **send-incl-mcast-ir-on-ndf false**
    - **rx-discard-on-ndf bum**
  - BUM frames received on any SAP or I-ES VXLAN binding are flooded to:
    - local non-ES and single-active DF ES SAPs
    - local all-active ES SAPs (DF and NDF)
    - EVPN-VXLAN destinations (BUM frames received on an I-ES VXLAN binding follow split-horizon rules, so they can only be forwarded to EVPN-VXLAN destinations belonging to the other VXLAN instance.)
- egress PE

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

– Look up source VTEP for BUM frames received on EVPN-VXLAN.
  - If the source VTEP matches a PE with which the local PE shares an ES and a VXLAN service, then the local PE does not forward to the shared local ESs (this includes port, lag, and network-interconnect-VXLAN ESs).
  - The local PE forwards to local ESs that are not shared but only when in DF state.

Figure 123: EVPN-VXLAN network interconnect VXLAN multi-homing and local bias shows the BUM forwarding with local bias procedures in multi-instance VPLS services.

Figure 123: EVPN-VXLAN network interconnect VXLAN multi-homing and local bias

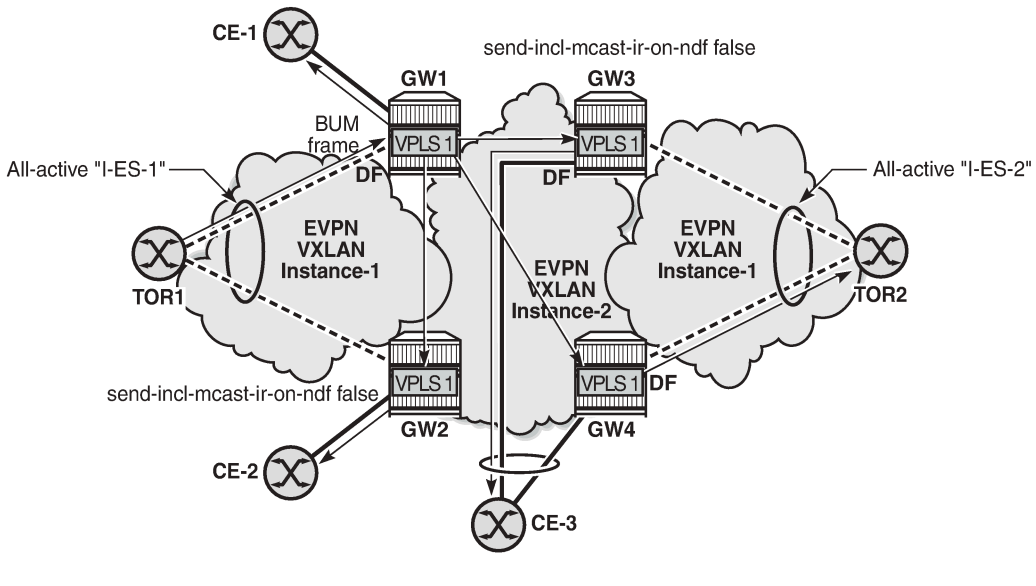

In the example, GW1 and GW2 are configured with **send-incl-mcast-ir-on-ndf false**. TOR1 generates BUM traffic that will only reach DF GW1 and is forwarded as follows.

- Ingress PE GW1 forwards to CE-1 and EVPN-VXLAN destinations GW2, GW3, and GW4.
- Egress PE GW2 identifies the source VTEP as a PE with which I-ES-1 is shared, so it does not forward the BUM frames to the local I-ES. PE GW2 forwards only to the non-shared ES and local SAPs, in this case, to CE-2.
- Egress PE GW3 receives the BUM traffic with a source VTEP that does not match any PE with which GW3 shares an ES, so it forwards to all ESs that are DF, in this case, to CE-3.
- Egress PE GW4 receives the BUM traffic with a source VTEP that does not match any PE with which GW4 shares an ES, so it forwards to all ESs that are DF, in this case, to TOR2 through I-ES-2.

Local bias with static VXLAN on I-ES

When a static VXLAN instance coexists with an EVPN-VXLAN instance in the same VPLS service, traffic blackholes may occur when the static VXLAN instance is associated to an all-active I-ES. This is because, when multi-homing is used with an EVPN-VXLAN network instance, the NDF PE always discards unknown unicast traffic to the static VXLAN instance (this is not the case with EVPN-MPLS if the unknown traffic has a BUM label).

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 124: All-active I-ES NDF PE-5 drops unknown unicast traffic shows the packet blackhole for unknown unicast traffic at all-active I-ES NDF PE-5.

Figure 124: All-active I-ES NDF PE-5 drops unknown unicast traffic

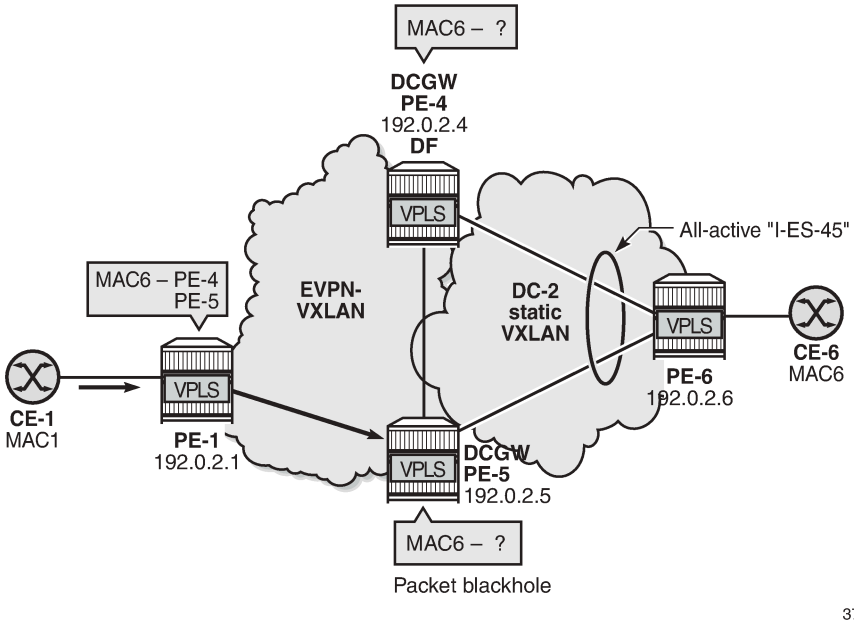

In the event that the remote PE-1 has learned the destination MAC address MAC6 via I-ES-45 EVPN destination, but the DCGWs PE-4 and PE-5 do not know MAC6, regular aliasing procedures allow that PE-1 sends unicast traffic with destination MAC6 to the NDF PE-5, which does not know MAC6 and drops all unknown unicast traffic, creating a blackhole for the flow.

When a static VXLAN instance coexists with an EVPN-VXLAN instance in the same VPLS service, Nokia recommends using a single-active I-ES or an anycast solution without I-ES instead of an all-active I-ES.

Configuration

Figure 125: Sample topology shows the sample topology with six SR OS nodes:

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 125: Sample topology

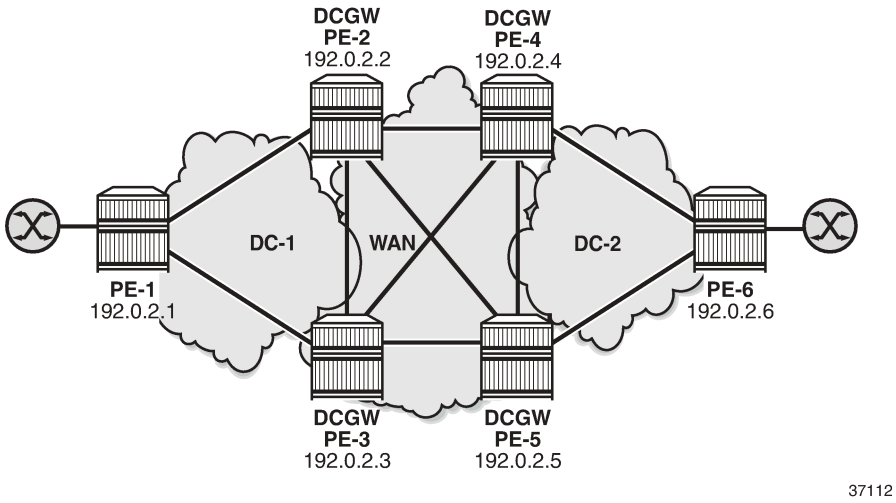

The initial configuration includes:

- Cards, MDAs, and ports
- Router interfaces
- IS-IS on all interfaces (level 1 in the DCs; level 2 in the WAN)

BGP is configured for the EVPN address family. PE-1 acts as Route Reflector (RR) in DC 1 and PE-6 as RR in DC 2; no RR is used in the WAN.

The BGP configuration on RR PE-1 in DC 1 is as follows. The BGP configuration on RR PE-6 in DC2 is similar.

```
# on PE-1:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            family {
                ipv4 false
                evpn true
            }
            cluster {
                cluster-id 192.0.2.1
            }
            rapid-update {
                evpn true
            }
            group "DC" {
                type internal
            }
            neighbor "192.0.2.2" {
                group "DC"
            }
            neighbor "192.0.2.3" {
                group "DC"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
```

On DCGWs PE-2 and PE-3, BGP is configured as follows. The policies are explained in the next section.

```
# on PE-2, PE-3:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            family {
                ipv4 false
                evpn true
            }
            rapid-update {
                evpn true
            }
            group "DC" {
                type internal
                import {
                    policy ["drop SOO-DCGW-23"]
                }
                export {
                    policy ["export DC routes and add SOO"]
                }
            }
            group "WAN" {
                type internal
                export {
                    policy ["export WAN routes only"]
                }
            }
            neighbor "192.0.2.1" {
                group "DC"
            }
            neighbor "192.0.2.4" {
                group "WAN"
            }
            neighbor "192.0.2.5" {
                group "WAN"
            }
        }
```

On DCGWs PE-4 and PE-5, BGP is configured as follows. The policies are explained in the next section.

```
# on PE-4, PE-5:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            family {
                ipv4 false
                evpn true
            }
            rapid-update {
                evpn true
            }
            group "DC" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                type internal
                import {
                    policy ["drop SOO-DCGW-45"]
                }
                export {
                    policy ["export DC routes and add SOO"]
                }
            }
            group "WAN" {
                type internal
                export {
                    policy ["export WAN routes only"]
                }
            }
            neighbor "192.0.2.6" {
                group "DC"
            }
            neighbor "192.0.2.2" {
                group "WAN"
            }
            neighbor "192.0.2.3" {
                group "WAN"
            }
        }
```

The following examples are configured:

- All-active multi-homing I-ESs in dual EVPN-VXLAN instance VPLS
- Single-active multi-homing I-ES when static VXLAN coexists with EVPN-VXLAN in the same VPLS

All-active multi-homing I-ESs in dual EVPN-VXLAN instance VPLS

Figure 126: All-active multi-homing for I-ESs shows the example topology with the service VPLS 1 on all nodes and two all-active I-ESs:

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 126: All-active multi-homing for I-ESs

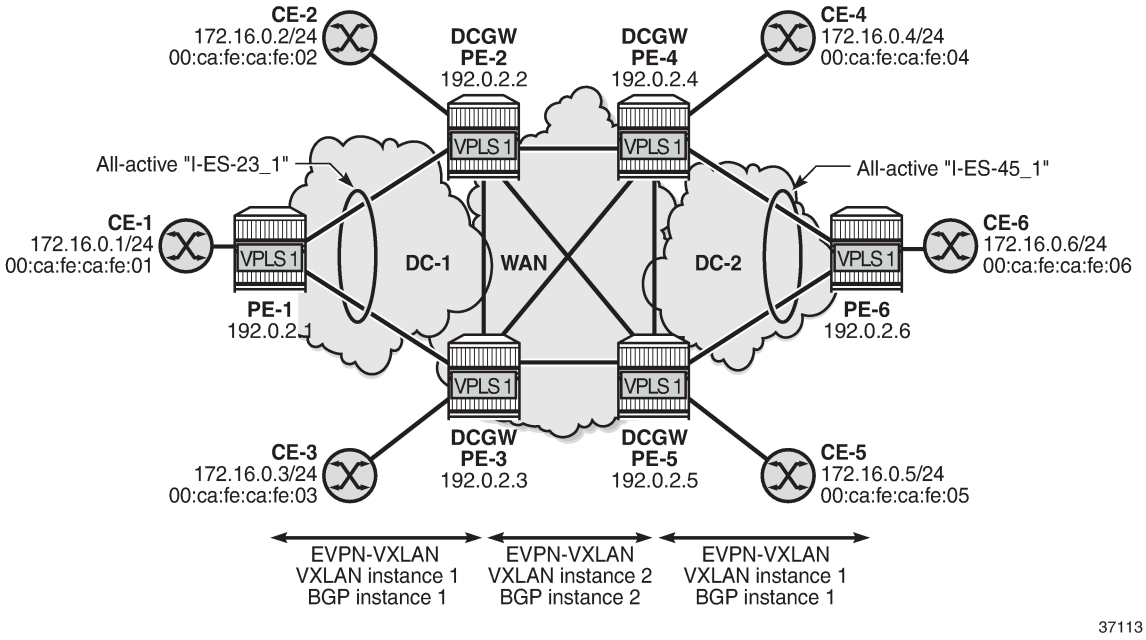

On PE-1, VPLS 1 is configured as follows. The configuration on PE-6 is similar.

```
# on PE-1:
configure {
    service {
        system {
            bgp-auto-rd-range {
                ip-address 192.0.2.1
                community-value {
                    start 1
                    end 1000
                }
            }
        }
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            vxlan {
                instance 1 {
                    vni 11
                }
            }
            bgp 1 {
                route-distinguisher auto-rd
                route-target {
                    export "target:64500:11"
                    import "target:64500:11"
                }
            }
            bgp-evpn {
                evi 1
                vxlan 1 {
                    admin-state enable
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    vxlan-instance 1
                    ecmp 2
                }
            }
            sap 1/2/1:1 {
            }
        }
```

On DCGW PE-2, the following all-active multi-homing I-ES is configured for VXLAN instance 1 and service id 1. The configuration on DCGW PE-3 is similar, but the preference value is 150 instead of 100.

```
# on PE-2:
configure {
    service {
        system {
            bgp-auto-rd-range {
                ip-address 192.0.2.2
                community-value {
                    start 1
                    end 1000
                }
            }
            bgp {
                evpn {
                    ethernet-segment "I-ES-23_1" {
                        admin-state enable
                        type virtual
                        esi 00:23:23:23:23:23:23:00:00:01
                        multi-homing-mode all-active
                        df-election {
                            service-carving-mode manual
                            manual {
                                evi 1 {
                                    end 1
                                }
                                preference {
                                    value 100       # on PE-3: preference value 150
                                }
                            }
                        }
                        association {
                            network-interconnect-vxlan 1 {
                                virtual-ranges {
                                    service-id 1 {
                                        end 1
                                    }
                                }
                            }
                        }
                    }
                }
            }
```

On DCGWs PE-4 and PE-5, the following I-ES is configured:

```
# on PE-4, PE-5:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "I-ES-45_1" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        admin-state enable
                        type virtual
                        esi 00:45:45:45:45:45:45:00:00:01
                        multi-homing-mode all-active
                        association {
                            network-interconnect-vxlan 1 {
                                virtual-ranges {
                                    service-id 1 {
                                        end 1
                                    }
                                }
                            }
                        }
                    }
                }
            }
```

On DCGWs PE-2, PE-3, PE-4, and PE-5, VPLS 1 is configured as follows. The **rx-discard-on-ndf bum** command makes the NDF drop any BUM traffic in VXLAN instance 1. VXLAN instance 2 is configured with **mh-mode network** and **routes>auto-disc>advertise true**.

```
# on PE-2, PE-3, PE-4, PE-5:
configure {
    service {
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            vxlan {
                instance 1 {
                    vni 11
                    rx-discard-on-ndf bum
                }
                instance 2 {
                    vni 12
                }
            }
            bgp 1 {
                route-distinguisher auto-rd
                route-target {
                    export "target:64500:11"
                    import "target:64500:11"
                }
            }
            bgp 2 {
                route-distinguisher auto-rd
                route-target {
                    export "target:64500:12"
                    import "target:64500:12"
                }
            }
            bgp-evpn {
                evi 1
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                    default-route-tag 0xb
                    ecmp 2
                    routes {
                        auto-disc {
                            advertise true
                        }
                    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
                vxlan 2 {
                    admin-state enable
                    vxlan-instance 2
                    default-route-tag 0xc
                    ecmp 2
                    mh-mode network
                    routes {
                        auto-disc {
                            advertise true
                        }
                    }
                }
            }
            sap 1/2/1:1 {
            }
        }
```

On PE-2 and PE-3, the following policies are configured.

- The import policy "drop SOO-DCGW-23" in group "DC" is used to drop all VXLAN instance 1 routes between PE-2 and PE-3.
- The export policy "export WAN routes only" in group "WAN" is applied to avoid sending VXLAN instance 1 routes to the WAN PEs.
- The export policy "export DC routes and add SOO" in group "DC" is used to tag VXLAN instance 1 routes with community "SOO-23".

```
# on PE-2, PE-3:
configure {
    policy-options {
        community "SOO-23" {
            member "origin:64500:23" { }
        }
        policy-statement "drop SOO-DCGW-23" {               # import in group "DC"
            entry 10 {
                from {
                    family [evpn]
                    community {
                        name "SOO-23"
                    }
                }
                action {
                    action-type reject
                }
            }
            default-action {
                action-type accept
            }
        }
        policy-statement "export WAN routes only" {         # export in group "WAN"
            entry 10 {
                from {
                    family [evpn]
                    tag 11
                }
                action {
                    action-type reject
                }
            }
            default-action {
                action-type accept
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
        }
        policy-statement "export DC routes and add SOO" {     # export in group "DC"
            entry 10 {
                from {
                    family [evpn]
                    tag 11
                }
                action {
                    action-type accept
                    community {
                        add ["SOO-23"]
                    }
                }
            }
            default-action {
                action-type accept
            }
        }
```

On PE-4 and PE-5, the following policies are configured:

```
# on PE-4, PE-5:
configure {
    policy-options {
        community "SOO-45" {
            member "origin:64500:45" { }
        }
        policy-statement "drop SOO-DCGW-45" {               # import in group "DC"
            entry 10 {
                from {
                    family [evpn]
                    community {
                        name "SOO-45"
                    }
                }
                action {
                    action-type reject
                }
            }
            default-action {
                action-type accept
            }
        }
        policy-statement "export WAN routes only" {         # export in group "WAN"
            entry 10 {
                from {
                    family [evpn]
                    tag 11
                }
                action {
                    action-type reject
                }
            }
            default-action {
                action-type accept
            }
        }
        policy-statement "export DC routes and add SOO" {   # export in group "DC"
            entry 10 {
                from {
                    family [evpn]
                    tag 11
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
                action {
                    action-type accept
                    community {
                        add ["SOO-45"]
                    }
                }
            }
            default-action {
                action-type accept
            }
        }
```

For VPLS 1, PE-2 is DF and PE-3 is NDF in the I-ES "I-ES-23_1":

```
[/]
A:admin@PE-2# show service id 1 ethernet-segment
No sap entries
No sdp entries

===============================================================================
VXLAN Ethernet-Segment Information
===============================================================================
VXLAN Instance        Eth-Seg                       Status
-------------------------------------------------------------------------------
1                     I-ES-23_1                     DF
===============================================================================
```

```
[/]
A:admin@PE-3# show service id 1 ethernet-segment
No sap entries
No sdp entries

===============================================================================
VXLAN Ethernet-Segment Information
===============================================================================
VXLAN Instance        Eth-Seg                       Status
-------------------------------------------------------------------------------
1                     I-ES-23_1                     NDF
===============================================================================
```

PE-4 is NDF and PE-5 is DF in the I-ES "I-ES-45_1":

```
[/]
A:admin@PE-4# show service vxlan-instance-using ethernet-segment

===============================================================================
VXLAN Ethernet-Segment Information
===============================================================================
SvcId       VXLAN Instance      ES Name                       Status
-------------------------------------------------------------------------------
1           1                   I-ES-45_1                     NDF
===============================================================================
```

```
[/]
A:admin@PE-5# show service vxlan-instance-using ethernet-segment

===============================================================================
VXLAN Ethernet-Segment Information
===============================================================================
SvcId       VXLAN Instance      ES Name                       Status
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
1           1                   I-ES-45_1                       DF
===============================================================================
```

On leaf PE-1, the VXLAN destinations in VXLAN instance 1 are the following:

```
[/]
A:admin@PE-1# show service id 1 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
Instance    VTEP Address                            Egress VNI  EvpnStatic Num
 Mcast       Oper State                              L2 PBR     SupBcasDom MACs
-------------------------------------------------------------------------------
1           192.0.2.2                               11          evpn       0
 BUM         Up                                      No          No
1           192.0.2.3                               11          evpn       0
 BUM         Up                                      No          No
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 2
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN-VXLAN Ethernet Segment Dest
===============================================================================
Instance  Eth SegId                       Num. Macs     Last Change
-------------------------------------------------------------------------------
1         00:23:23:23:23:23:23:00:00:01   5             09/29/2021 17:54:10
-------------------------------------------------------------------------------
Number of entries: 1
-------------------------------------------------------------------------------
===============================================================================
```

On DCGW PE-2, the VXLAN destinations in VXLAN instances 1 and 2 are the following:

```
[/]
A:admin@PE-2# show service id 1 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
Instance    VTEP Address                            Egress VNI  EvpnStatic Num
 Mcast       Oper State                              L2 PBR     SupBcasDom MACs
-------------------------------------------------------------------------------
1           192.0.2.1                               11          evpn       1
 BUM         Up                                      No          No
2           192.0.2.3                               12          evpn       1
 BUM         Up                                      No          No
2           192.0.2.4                               12          evpn       1
 BUM         Up                                      No          No
2           192.0.2.5                               12          evpn       1
 BUM         Up                                      No          No
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 4
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN-VXLAN Ethernet Segment Dest
===============================================================================
Instance  Eth SegId                       Num. Macs     Last Change
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
2         00:45:45:45:45:45:00:00:01   1             09/29/2021 17:54:35
-------------------------------------------------------------------------------
Number of entries: 1
-------------------------------------------------------------------------------
===============================================================================
```

ECMP 2 is configured, so aliasing is used. PE-1 can reach the I-ES "I-ES-23_1" in VXLAN instance 1 via PE-2 and PE-3:

```
[/]
A:admin@PE-1# show service id 1 vxlan esi 00:23:23:23:23:23:00:00:01

===============================================================================
BGP EVPN-VXLAN Ethernet Segment Dest
===============================================================================
Instance  Eth SegId                       Num. Macs     Last Change
-------------------------------------------------------------------------------
1         00:23:23:23:23:23:00:00:01   5             09/29/2021 17:54:10
-------------------------------------------------------------------------------
Number of entries: 1
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN-VXLAN Dest TEP Info
===============================================================================
Instance  TEP Address                   Egr VNI             Last Change
-------------------------------------------------------------------------------
1         192.0.2.2                     11                  09/29/2021 17:53:32
1         192.0.2.3                     11                  09/29/2021 17:54:10
-------------------------------------------------------------------------------
Number of entries : 2
-------------------------------------------------------------------------------
===============================================================================
```

In a similar way, PE-4 can reach the I-ES "I-ES-23_1" via PE-2 and PE-3 in VXLAN instance 2:

```
[/]
A:admin@PE-4# show service id 1 vxlan esi 00:23:23:23:23:23:00:00:01

===============================================================================
BGP EVPN-VXLAN Ethernet Segment Dest
===============================================================================
Instance  Eth SegId                       Num. Macs     Last Change
-------------------------------------------------------------------------------
2         00:23:23:23:23:23:00:00:01   1             09/29/2021 17:54:21
-------------------------------------------------------------------------------
Number of entries: 1
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN-VXLAN Dest TEP Info
===============================================================================
Instance  TEP Address                   Egr VNI             Last Change
-------------------------------------------------------------------------------
2         192.0.2.2                     12                  09/29/2021 17:54:21
2         192.0.2.3                     12                  09/29/2021 17:54:21
-------------------------------------------------------------------------------
Number of entries : 2
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
```

The following command on PE-2 shows the ES information for "I-ES-23_1": DF status, DF candidate list, VXLAN instance service range, and so on:

```
[/]
A:admin@PE-2# show service system bgp-evpn ethernet-segment name "I-ES-23_1" all

===============================================================================
Service Ethernet Segment
===============================================================================
Name                    : I-ES-23_1
Eth Seg Type            : Virtual
Admin State             : Enabled           Oper State         : Up
ESI                     : 00:23:23:23:23:23:23:00:00:01
Oper ESI                : 00:23:23:23:23:23:23:00:00:01
Auto-ESI Type           : None
AC DF Capability        : Include
Multi-homing            : allActive         Oper Multi-homing  : allActive
ES SHG Label            : 524287
Source BMAC LSB         : None
VXLAN Instance Id       : 1
ES Activation Timer     : 3 secs (default)
Oper Group              : (Not Specified)
Svc Carving             : manual            Oper Svc Carving   : manual
Cfg Range Type          : lowest-pref

-------------------------------------------------------------------------------
DF Pref Election Information
-------------------------------------------------------------------------------
Preference     Preference     Last Admin Change        Oper Pref      Do No
Mode           Value                                   Value          Preempt
-------------------------------------------------------------------------------
revertive      100            09/29/2021 17:53:32      100            Disabled
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
EVI Ranges
-------------------------------------------------------------------------------
From                                To
-------------------------------------------------------------------------------
1                                   1
-------------------------------------------------------------------------------
ISID Ranges: <none>
===============================================================================

===============================================================================
EVI Information
===============================================================================
EVI                 SvcId               Actv Timer Rem      DF
-------------------------------------------------------------------------------
1                   1                   0                   yes
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================

-------------------------------------------------------------------------------
DF Candidate list
-------------------------------------------------------------------------------
EVI                                 DF Address
-------------------------------------------------------------------------------
1                                   192.0.2.2
1                                   192.0.2.3
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
Number of entries: 2
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------

===============================================================================

---snip---

===============================================================================
Vxlan Instance Service Ranges
===============================================================================
Svc Range Start          Svc Range End            Last Changed
-------------------------------------------------------------------------------
1                        1                        09/29/2021 17:53:32
-------------------------------------------------------------------------------
Number of Entries: 1
===============================================================================
```

When traffic is sent between CE-1 and CE-6, the EVPN-MAC routes are sent with I-ESI. The FDB for VPLS 1 on PE-1 shows I-ESI 00:23:23:23:23:23:23:00:00:01 of "I-ES-23_1" as source identifier for MAC address 00:ca:fe:ca:fe:06 of CE-6:

```
[/]
A:admin@PE-1# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
1          00:ca:fe:ca:fe:01 sap:1/2/1:1             L/0      09/29/21 18:04:32
1          00:ca:fe:ca:fe:06 eES:                    Evpn     09/29/21 18:04:32
                             00:23:23:23:23:23:23:00:00:01
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

On PE-2, the FDB for VPLS 1 shows I-ESI 00:45:45:45:45:45:45:00:00:01 of "I-ES-45_1" as source identifier for MAC address 00:ca:fe:ca:fe:06 of CE-6:

```
[/]
A:admin@PE-2# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
1          00:ca:fe:ca:fe:01 vxlan-1:                Evpn     09/29/21 18:04:32
                             192.0.2.1:11
1          00:ca:fe:ca:fe:06 eES:                    Evpn     09/29/21 18:04:32
                             00:45:45:45:45:45:45:00:00:01
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
```

Single-active multi-homing I-ES when static VXLAN coexists with EVPN-VXLAN in the same VPLS

Figure 127: I-ES with EVPN-VXLAN in DC 1 and static VXLAN in DC2 shows the sample topology for VPLS 2 with static VXLAN in DC 2 and the single-active "I-ES-45_2" on PE-4 and PE-5.

Figure 127: I-ES with EVPN-VXLAN in DC 1 and static VXLAN in DC2

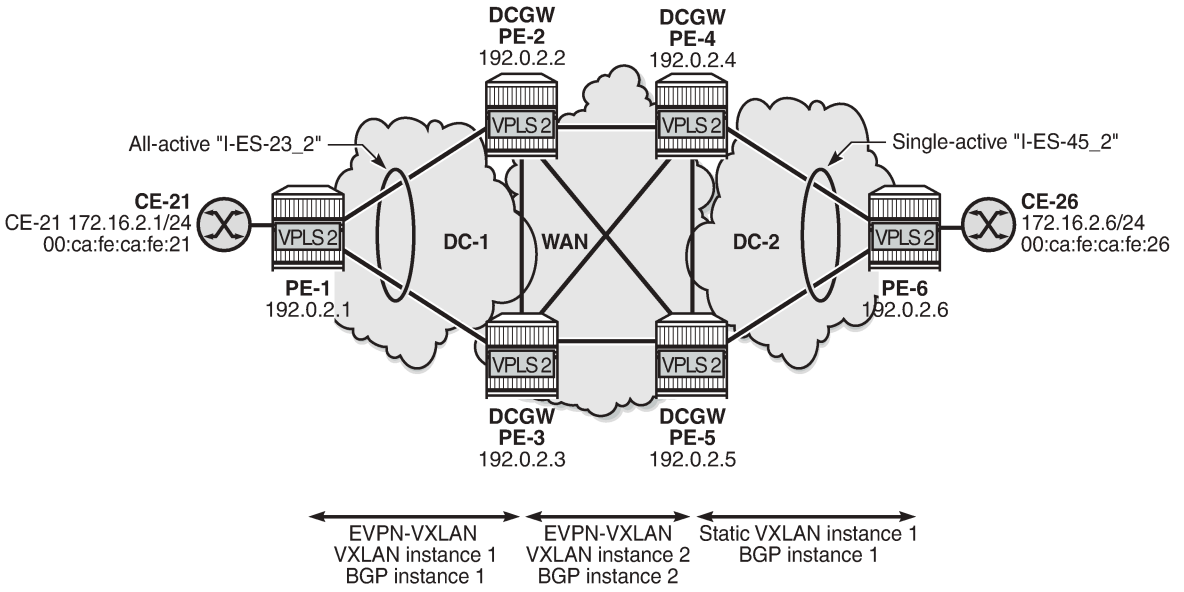

The configuration for VPLS 2 on PE-1, PE-2, and PE-3 is similar to the configuration for VPLS 1, so only the configuration on PE-4, PE-5, and PE-6 is shown.

On PE-6, VPLS 2 is configured with static VXLAN using non-anycast VTEP addresses:

```
# on PE-6:
configure {
    service {
        vpls "VPLS 2" {
            admin-state enable
            service-id 2
            customer "1"
            vxlan {
                instance 1 {
                    vni 21
                    egress-vtep 192.0.2.4 { }
                    egress-vtep 192.0.2.5 { }
                }
            }
            sap 1/2/1:2 {
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

To avoid blackholes, the I-ES between DCGWs PE-4 and PE-5 must not be all-active.

On PE-4 and PE-5, the single-active I-ES "I-ES-45_2" is configured as follows:

```
# on PE-4, PE-5:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "I-ES-45_2" {
                        admin-state enable
                        type virtual
                        esi 00:45:45:45:45:45:45:00:00:02
                        multi-homing-mode single-active
                        association {
                            network-interconnect-vxlan 1 {
                                virtual-ranges {
                                    service-id 2 {
                                        end 2
                                    }
                                }
                            }
                        }
                    }
                }
            }
```

On PE-4 and PE-5, VPLS 2 is configured as follows:

```
# on PE-4, PE-5:
configure {
    service {
       vpls "VPLS 2" {
            admin-state enable
            service-id 2
            customer "1"
            vxlan {
                instance 1 {
                    vni 21
                    egress-vtep 192.0.2.6 { }
                }
                instance 2 {
                    vni 22
                }
            }
            bgp 2 {
                route-distinguisher auto-rd
                route-target {
                    export "target:64500:22"
                    import "target:64500:22"
                }
            }
            bgp-evpn {
                evi 2
                vxlan 2 {
                    admin-state enable
                    vxlan-instance 2
                    default-route-tag 0x16      # default route tag 22
                    ecmp 2
                    mh-mode network
                    routes {
                        auto-disc {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                            advertise true
                        }
                    }
                }
            }
            sap 1/2/1:2 {
            }
        }
```

The policies on all DCGWs must be modified with tag 21 for VXLAN instance 1 in VPLS 2, as follows:

```
# on PE-2, PE-3:
configure {
    policy-options {
        policy-statement "export WAN routes only" {
            entry 20 {
                from {
                    family [evpn]
                    tag 21
                }
                action {
                    action-type reject
                }
            }
            default-action {
                action-type accept
            }
        }
        policy-statement "export DC routes and add SOO" {
            entry 20 {
                from {
                    family [evpn]
                    tag 21
                }
                action {
                    action-type accept
                    community {
                        add ["SOO-23"]
                    }
                }
            }
            default-action {
                action-type accept
            }
        }
```

DCGW PE-5 is NDF for "I-ES-45_2":

```
[/]
A:admin@PE-5# show service id 2 ethernet-segment
No sap entries
No sdp entries

===============================================================================
VXLAN Ethernet-Segment Information
===============================================================================
VXLAN Instance        Eth-Seg                         Status
-------------------------------------------------------------------------------
1                     I-ES-45_2                       NDF
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

On PE-5, the status of VXLAN instance 1 in VPLS 2 is mhStandby, as follows:

```
[/]
A:admin@PE-5# show service id 2 vxlan
===============================================================================
VPLS VXLAN
===============================================================================
Vxlan Src Vtep IP: N/A

===============================================================================
Vxlan Instance
===============================================================================
VXLAN Instance                VNI         AR          Oper-flags   VTEP
                                                                   security
-------------------------------------------------------------------------------
1                             21          none        mhStandby    disabled
2                             22          none        none         disabled
-------------------------------------------------------------------------------
Number of Entries : 2
-------------------------------------------------------------------------------
===============================================================================
```

The VXLAN destinations in VPLS 2 on PE-5 are the following:

```
[/]
A:admin@PE-5# show service id 2 vxlan destinations
===============================================================================
Egress VTEP, VNI
===============================================================================
Instance    VTEP Address                          Egress VNI  EvpnStatic Num
 Mcast       Oper State                            L2 PBR     SupBcasDom MACs
-------------------------------------------------------------------------------
1           192.0.2.6                             21          static     0
 -           Up                                    No          No
2           192.0.2.2                             22          evpn       0
 BUM         Up                                    No          No
2           192.0.2.3                             22          evpn       0
 BUM         Up                                    No          No
2           192.0.2.4                             22          evpn       0
 BUM         Up                                    No          No
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 4
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN-VXLAN Ethernet Segment Dest
===============================================================================
Instance  Eth SegId                     Num. Macs     Last Change
-------------------------------------------------------------------------------
2         00:23:23:23:23:23:23:00:00:02   1           09/29/2021 18:14:49
-------------------------------------------------------------------------------
Number of entries: 1
-------------------------------------------------------------------------------
===============================================================================
```

Note:
An anycast solution without I-ES can also be configured when an EVPN-VXLAN coexists with a static VXLAN.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Conclusion

Service providers can use I-ESs for better bandwidth utilization and redundancy in large DCs. EVPN all-active multi-homing I-ESs can be used in dual EVPN-VXLAN instance VPLS services. However, when a static VXLAN instance coexists with EVPN-VXLAN in the same VPLS, a single-active multi-homing I-ES (or an anycast solution without I-ES) is required to avoid blackholes.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

EVPN Multi-Homing for VXLAN VPLS Services

This chapter provides information about EVPN Multi-Homing for VXLAN VPLS Services.

Topics in this chapter include:

Applicability

The information and configuration in this chapter are based on SR OS Release 21.7.R1.

EVPN multi-homing has been supported in SR OS for EVPN-MPLS and PBB-EVPN in SR OS Release 13.0.R4 and later. SR OS Release 16.0 introduced EVPN multi-homing for EVPN-VXLAN on Epipe services. EVPN-VXLAN multi-homing in a single VXLAN instance VPLS or R-VPLS service—as specified in RFC 8365—is supported in SR OS Release 19.5.R1, and later.

Before you read this chapter, ensure you are familiar with the concepts in the EVPN for VXLAN Tunnels (Layer 2) chapter.

Overview

Some Service Providers are deploying large Telco cloud Data Centers (DCs) where SR OS nodes are used as leaf switches in a VXLAN fabric. In those cases, all-active multi-homing can provide redundancy and maximize the bandwidth use.

The multi-homing procedures consist of three components:

- Designated Forwarder (DF) election
  - The PEs attached to the same Ethernet Segment (ES) elect a single PE as DF to:
    - forward all traffic, in case of single-active mode
    - forward all Broadcast, Unknown unicast, Multicast (BUM) traffic, in case of all-active mode
- split-horizon
  - BUM traffic received from a peer ES PE is filtered so that it is not looped back to the CE that first transmitted the frame.
  - in EVPN-VXLAN services, split-horizon is only used with all-active mode and makes use of the local bias procedure described in RFC 8365.
- aliasing
  - PEs that are not attached to the ES can process non-zero Ethernet Segment Identifier (ESI) MAC/IP routes and AD routes and create ES destinations to which per-flow Equal Cost Multi-Path (ECMP) can be applied.
  - Aliasing only applies to all-active mode.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Split-horizon using local bias

In EVPN-MPLS services, split-horizon filtering uses ESI labels. VXLAN does not support ESI labels or MPLS labels. In EVPN-VXLAN services, the split-horizon filtering is based on the tunnel source IP address. In RFC 8365, this forwarding is referred to as local bias. Local bias works as follows:

- Every PE knows the IP addresses associated with the other PEs with which it has shared multi-homed ESs.
- The ingress PE replicates locally to all directly attached ESs, regardless of the DF state, for all flooded traffic coming from the access interfaces. BUM frames received on any SAP are flooded to:
  - local non-ES SAPs and non-ES SDP bindings
  - local all-active ES SAPs (DF and NDF)
  - local single-active ES SDP bindings and SAPs (DF only)
  - EVPN-VXLAN destinations
- When an egress PE receives a BUM frame from a VXLAN binding, it looks up the source IP address in the tunnel header and filters out the frame on all local interfaces connected to ESs that are shared with the ingress PE. The following rules apply to egress PE forwarding for EVPN-VXLAN services.
  1. The source VTEP is looked up for BUM frames received on EVPN-VXLAN.
  2. The router checks if the source VTEP matches one of the PEs with which the egress PE shared both an ES and a VXLAN service.
     - If there is a match, the egress PE is not forwarding to the shared ES local SAPs.
     - If there is no match, the egress PE forwards to ES SAPs in DF state (as usual).

Figure 128: Split-horizon filtering based on tunnel source IP address shows an example of local bias forwarding for BUM frames.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 128: Split-horizon filtering based on tunnel source IP address

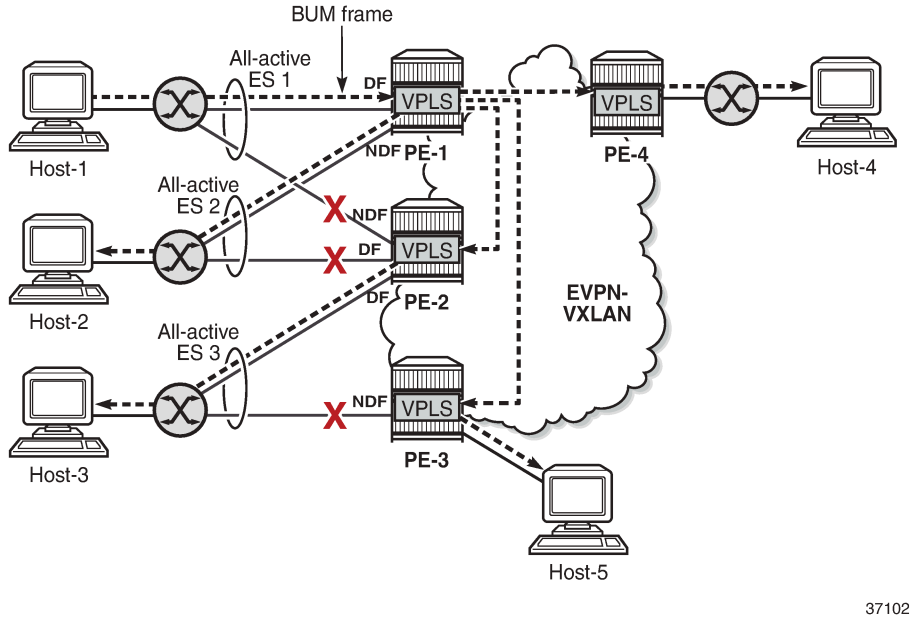

In this example, BUM frames sent by Host-1 are treated as follows.

- Ingress node PE-1 receives BUM frames from Host-1 and forwards them to the other PEs (EVPN-VXLAN destinations) and the local all-active ES SAP toward Host-2, even though the SAP is in NDF state.
- Egress node PE-2 receives BUM frames on VXLAN. PE-2 identifies the source VTEP as a PE with which two all-active ESs are shared, so it does not forward the BUM frames to the two shared ESs. PE-2 forwards the BUM frames to the non-shared ES toward Host-3 because it is in DF state.
- Egress node PE-3 receives BUM traffic from PE-1, with which it does not share any ESs, so it forwards the BUM frames based on normal rules: it does not forward them toward Host-3, because the ES SAP is in NDF state. PE-3 only forwards toward Host-5.
- PE-4 does not share any ESs with PE-1, so the normal rules apply. PE-4 forwards the BUM frames toward Host-4.

Known limitations for local bias

In VXLAN, there are no BUM labels or any tunnel indication that can identify BUM traffic. The egress PE must solely rely on the Customer MAC (CMAC) destination address and this may create transient issues.

- Duplicate unicast traffic may occur when the CMAC destination address MAC1 is unknown on the ingress PE-3, while known on the egress PEs (PE-1 and PE-2). Figure 129: Duplicate unicast packets when MAC1 is unknown on PE-3 only shows that a packet with destination MAC1 arrives at PE-3, where it is flooded via ingress replication to PE-1 and PE-2, where MAC1 is known. PE-1 and PE-2 both forward the packets with CMAC destination MAC1 to CE-1, so multiple copies are sent to CE-1.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 129: Duplicate unicast packets when MAC1 is unknown on PE-3 only

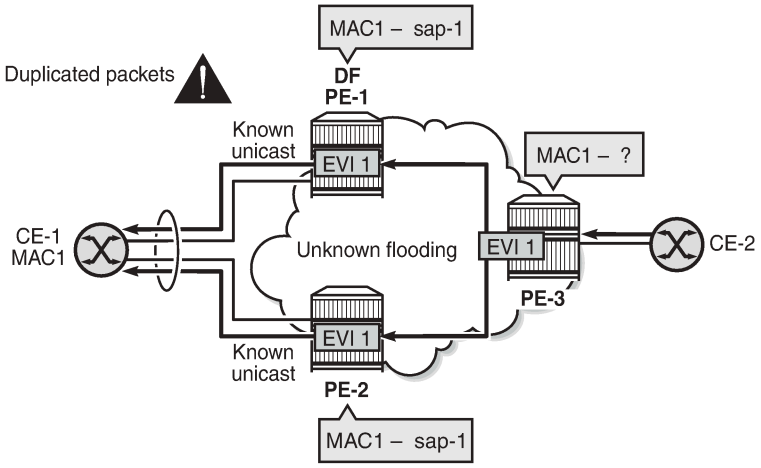

- A blackhole may occur when the CMAC destination address MAC1 is known on PE-3, but unknown on PE-1 and PE-2 and the aliasing hashing on PE-3 picks up the path to the NDF, where unknown unicast traffic is dropped, as shown in Figure 130: Packet blackhole for traffic on NDF PE-2 when MAC1 is known on PE-3 only. When the path to the DF is picked, no problem occurs, because the DF forwards BUM traffic.

  Figure 130: Packet blackhole for traffic on NDF PE-2 when MAC1 is known on PE-3 only

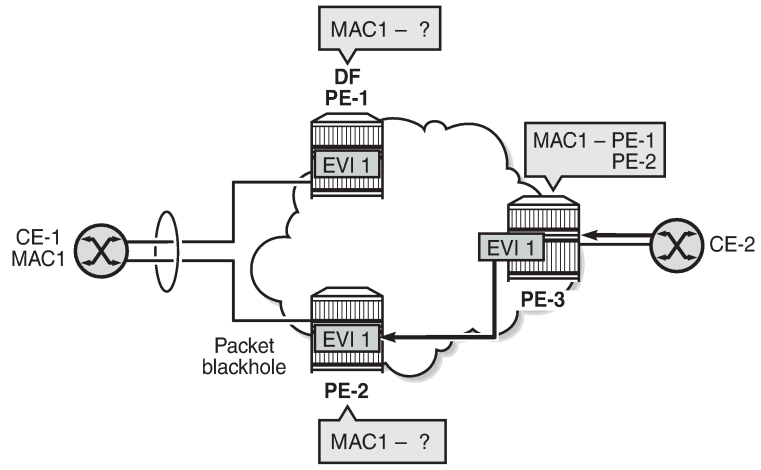

- A blackhole can be created when a remote SAP is disabled, as shown in Figure 131: Blackhole created when a remote SAP is disabled.

  Under normal circumstances, when CE-3 sends BUM traffic to ingress node PE-3, the local bias mechanism on PE-3 forwards the BUM packets to SAP3, even though it is NDF for the ES. The BUM traffic is also flooded to PE-2, where it is forwarded to CE-2, but not to SAP2, because the ES is shared with PE-3.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

When SAP3 is manually disabled (**admin-state disable**), PE-3 withdraws the AD per-EVI route corresponding to SAP3. This does not change the local bias filtering for SAP2 on PE-2, so when CE-3 sends BUM traffic, it can neither be forwarded to CE-23 via SAP3 nor by PE-2.

Figure 131: Blackhole created when a remote SAP is disabled

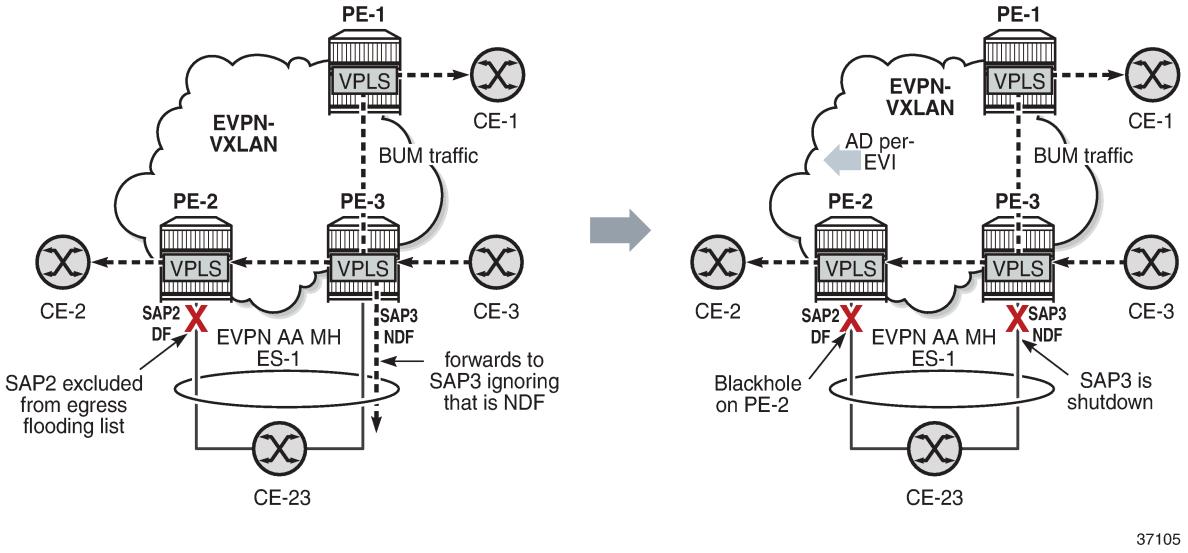

CLI

The multi-homing capabilities are enabled in all the PEs attached to the VPLS service by configuring the options **routes auto-disc advertise** and **mh-mode network** in the **vpls bgp-evpn vxlan** context.

The **routes auto-disc advertise** option is by default disabled, but it can be enabled as follows:

```
*[ex:/configure service vpls "VPLS 2" bgp-evpn vxlan 1 routes auto-disc]
A:admin@PE-2# advertise true
```

This **routes auto-disc advertise** command is only configurable for EVPN-VXLAN VPLS services and is implicitly enabled on all instances where it is not configurable. **routes auto-disc advertise** is required in nodes with local ESs and remote ESs to process and enable the creation of ES destinations.

When **routes auto-disc advertise** is enabled, BGP-EVPN:

- processes Auto-Discovery per EVPN instance (AD per-EVI) routes and AD per-ES routes
- processes MAC/IP routes with non-zero Ethernet Segment Identifier (ESI) — without resetting the ESI to zero
- creates ES destinations upon receiving MAC/IP routes and AD per-ES/EVI routes with non-zero ESI.

The **mh-mode** option can be configured with the values **access** or **network**. For EVPN-VXLAN services, the default value is **access**. The following command configures **mh-mode network**:

```
*[ex:/configure service vpls "VPLS 2" bgp-evpn vxlan 1]
A:admin@PE-2# mh-mode network
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

When **mh-mode network** is configured, BGP-EVPN:

- activates multi-homing for the local ES SAPs or SDP-bindings and creates ES associations and related processes, such as:
  - the local bias mode allowing the system to add all-active SAPs to the flooding list regardless of the DF state
  - the source VTEP lookup mode
- runs DF election for the ESs associated with the service
- triggers the advertisement of AD per-ES routes, AD per-EVI routes, and non-zero MAC/IP routes for the ESs in the service.

Configuration

The following examples are configured:

- EVPN-VXLAN multi-homing with system IPv4 VTEP addresses
- EVPN-VXLAN multi-homing with non-system IPv4 VTEP addresses
- EVPN-VXLAN multi-homing with non-system IPv6 VTEP addresses

EVPN-VXLAN multi-homing with system IPv4 VTEP addresses

Figure 132: Example topology shows the topology with three all-active multi-homing ESs and one single-active multi-homing ES. This example shows the configuration for virtual Ethernet Segments, as described in the Virtual Ethernet Segments chapter, but non-virtual ES can also be used.

Figure 132: Example topology

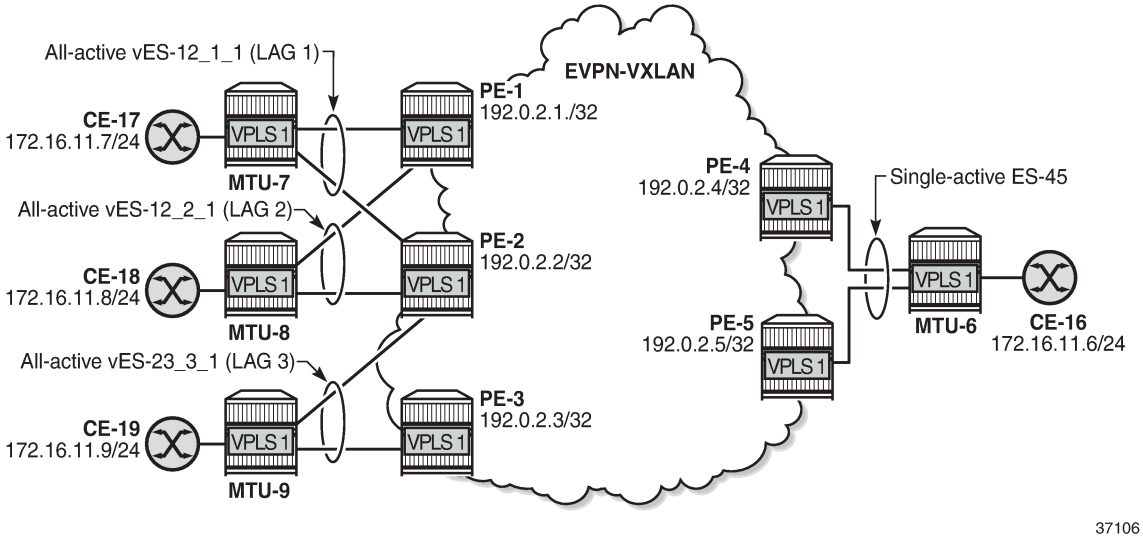

The initial configuration on the PEs includes:

- cards, MDAs, ports

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- LAG 1 on MTU-7, PE-1, PE-2

  LAG 2 on MTU-8, PE-1, PE-2

  LAG 3 on MTU-9, PE-2, PE-3
- router interfaces
- IS-IS between the PEs
- SR-ISIS between PE-4 and MTU-6 and between PE-5 and MTU-6 (and TLDP for SDP signaling)

BGP is configured between the PEs for the EVPN address family. PE-1 acts as route reflector, as follows:

```
# on RR PE-1:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            rapid-update {
                evpn true
            }
            group "internal" {
                peer-as 64500
                family {
                    evpn true
                }
                cluster {
                    cluster-id 192.0.2.1
                }
            }
            neighbor "192.0.2.2" {
                group "internal"
            }
            neighbor "192.0.2.3" {
                group "internal"
            }
            neighbor "192.0.2.4" {
                group "internal"
            }
            neighbor "192.0.2.5" {
                group "internal"
            }
        }
```

ES configuration

The all-active ESs "vES-12_1_1" and "vES-12_2_1" are configured on PE-1 and PE-2. The configuration on PE-1 is as follows. The configuration on PE-2 is similar, but with different preference values.

```
# on PE-1:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "vES-12_1_1" {
                        admin-state enable
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        type virtual
                        esi 00:12:12:12:12:12:12:00:01:01
                        multi-homing-mode all-active
                        df-election {
                            service-carving-mode manual
                            manual {
                                evi 1 {
                                    end 1
                                }
                                preference {
                                    value 100       # on PE-2: preference value 150
                                }
                            }
                        }
                        association {
                            lag "lag-1" {
                                virtual-ranges {
                                    dot1q {
                                        q-tag 1 {
                                            end 1
                                        }
                                    }
                                }
                            }
                        }
                    }
                    ethernet-segment "vES-12_2_1" {
                        admin-state enable
                        type virtual
                        esi 00:12:12:12:12:12:12:00:02:01
                        multi-homing-mode all-active
                        df-election {
                            service-carving-mode manual
                            manual {
                                evi 1 {
                                    end 1
                                }
                                preference {
                                    value 150       # on PE-2: preference value 100
                                }
                            }
                        }
                        association {
                            lag "lag-2" {
                                virtual-ranges {
                                    dot1q {
                                        q-tag 1 {
                                            end 1
                                        }
                                    }
                                }
                            }
                        }
                    }
                }
```

On PE-2 and PE-3, the all-active ES "vES-23_3_1" is configured in a similar way:

```
# on PE-2:
configure {
    service {
        system {
            bgp {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                evpn {
                    ethernet-segment "vES-23_3_1" {
                        admin-state enable
                        type virtual
                        esi 00:23:23:23:23:23:23:00:03:01
                        multi-homing-mode all-active
                        df-election {
                            service-carving-mode manual
                            manual {
                                evi 1 {
                                    end 1
                                }
                                preference {
                                    value 100       # on PE-3: preference value 150
                                }
                            }
                        }
                        association {
                            lag "lag-3" {
                                virtual-ranges {
                                    dot1q {
                                        q-tag 1 {
                                            end 1
                                        }
                                    }
                                }
                            }
                        }
                    }
                }
```

On PE-4 and PE-5, the single-active ES "ES-45" is configured, as follows:

```
# on PE-4:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "ES-45" {
                        admin-state enable
                        esi 00:45:45:45:45:45:45:00:00:01
                        multi-homing-mode single-active
                        df-election {
                            service-carving-mode manual
                            manual {
                                evi 1 {
                                    end 1
                                }
                                preference {
                                    value 100       # on PE-5: preference value 150
                                }
                            }
                        }
                        association {
                            sdp 46 {                # on PE-5: sdp 56
                            }
                        }
                    }
                }
            }
        }
        sdp 46 {                                    # on PE-5: sdp 56
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            admin-state enable
            delivery-type mpls
            sr-isis true
            far-end {
                ip-address 192.0.2.6
            }
        }
```

VPLS configuration

VPLS 1 is configured on PE-2 as follows. The configuration is similar on PE-1 and PE-3.

```
# on PE-2:
configure {
    service {
        system {
            bgp-auto-rd-range {
                ip-address 192.0.2.2           # different values on different PEs
                community-value {
                    start 1
                    end 1000
                }
            }
        }
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            vxlan {
                instance 1 {
                    vni 1
                }
            }
            bgp 1 {
                route-distinguisher auto-rd
                route-target {
                    export "target:64500:1"
                    import "target:64500:1"
                }
            }
            bgp-evpn {
                evi 1
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                    ecmp 2
                    mh-mode network
                    routes {
                        auto-disc {
                            advertise true
                        }
                    }
                }
            }
            sap lag-1:1 {                   # LAG 1 also on PE-1, not on PE-3
            }
            sap lag-2:1 {                   # LAG 2 also on PE-1, not on PE-3
            }
            sap lag-3:1 {                   # LAG 3 also on PE-3, not on PE-1
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The EVPN-VXLAN multi-homing capabilities are enabled in the PEs attached to VPLS 1 by the commands **routes auto-disc advertise** and **mh-mode network**. The **routes auto-disc advertise** command enables the advertisement and processing of multi-homing routes, and the **mh-mode network** command activates the DF election procedures.

ECMP is required for per-flow load balancing for VXLAN ES destinations with two or more next hops. In this example, ECMP is configured with a value of 2.

On PE-4, VPLS 1 is configured as follows. The configuration on PE-5 is similar.

```
# on PE-4:
configure {
    service {
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            vxlan {
                instance 1 {
                    vni 1
                }
            }
            bgp 1 {
                route-distinguisher auto-rd
                route-target {
                    export "target:64500:1"
                    import "target:64500:1"
                }
            }
            bgp-evpn {
                evi 1
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                    ecmp 2
                    mh-mode network
                    routes {
                        auto-disc {
                            advertise true
                        }
                    }
                }
            }
            spoke-sdp 46:1 {                        # on PE-5: spoke-sdp 56:1
            }
        }
```

Show commands

The following command shows that the commands **mh-mode network** and **routes auto-disc advertise** are enabled:

```
[/]
A:admin@PE-2# show service id 1 bgp-evpn

===============================================================================
BGP EVPN Table
===============================================================================
MAC Advertisement  : Enabled            Unknown MAC Route  : Disabled
CFM MAC Advertise  : Disabled
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Creation Origin    : manual
MAC Dup Detn Moves : 5                  MAC Dup Detn Window: 3
MAC Dup Detn Retry : 9                  Number of Dup MACs : 0
MAC Dup Detn BH    : Disabled
IP Route Advert    : Disabled
Sel Mcast Advert   : Disabled

EVI                : 1
Ing Rep Inc McastAd: Enabled
Accept IVPLS Flush : Disabled

-------------------------------------------------------------------------------
Detected Duplicate MAC Addresses            Time Detected
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN VXLAN Information
===============================================================================
Admin Status       : Enabled            Bgp Instance       : 1
Vxlan Instance     : 1
Max Ecmp Routes    : 2
Default Route Tag  : none
Send EVPN Encap    : Enabled
Imet-Ir routes     : Enabled
MH Mode            : network
Auto Disc Route Adv: Enabled
Oper Group         :
===============================================================================
```

The following command shows that PE-1 is DF for the all-active ES vES-12_1_1 and NDF for the all-active ES vES-12_2_1:

```
[/]
A:admin@PE-1# show service id 1 ethernet-segment

===============================================================================
SAP Ethernet-Segment Information
===============================================================================
SAP                   Eth-Seg                         Status
-------------------------------------------------------------------------------
lag-1:1               vES-12_1_1                      DF
lag-2:1               vES-12_2_1                      NDF
===============================================================================
No sdp entries
No vxlan instance entries
```

The following command shows that PE-2 is NDF for the all-active ES vES-12_1_1 and DF for the other two all-active ESs:

```
[/]
A:admin@PE-2# show service id 1 ethernet-segment

===============================================================================
SAP Ethernet-Segment Information
===============================================================================
SAP                   Eth-Seg                         Status
-------------------------------------------------------------------------------
lag-1:1               vES-12_1_1                      NDF
lag-2:1               vES-12_2_1                      DF
lag-3:1               vES-23_3_1                      DF
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
No sdp entries
No vxlan instance entries
```

PE-3 is NDF for the all-active multi-homing ES vES-23_3_1:

```
[/]
A:admin@PE-3# show service id 1 ethernet-segment

===============================================================================
SAP Ethernet-Segment Information
===============================================================================
SAP                   Eth-Seg                         Status
-------------------------------------------------------------------------------
lag-3:1               vES-23_3_1                      NDF
===============================================================================
No sdp entries
No vxlan instance entries
```

PE-4 is DF for the single-active multi-homing ES ES-45:

```
[/]
A:admin@PE-4# show service id 1 ethernet-segment
No sap entries

===============================================================================
SDP Ethernet-Segment Information
===============================================================================
SDP                   Eth-Seg                         Status
-------------------------------------------------------------------------------
46:1                  ES-45                           DF
===============================================================================
No vxlan instance entries
```

PE-5 is NDF for the single-active multi-homing ES ES-45:

```
[/]
A:admin@PE-5# show service id 1 ethernet-segment
No sap entries

===============================================================================
SDP Ethernet-Segment Information
===============================================================================
SDP                   Eth-Seg                         Status
-------------------------------------------------------------------------------
56:1                  ES-45                           NDF
===============================================================================
No vxlan instance entries
```

The following command shows the VXLAN destinations for VPLS 1 on PE-3; the system addresses of the other PEs act as destination VTEP addresses.

```
[/]
A:admin@PE-3# show service id 1 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
Instance    VTEP Address                          Egress VNI  EvpnStatic Num
 Mcast       Oper State                            L2 PBR     SupBcasDom MACs
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
1           192.0.2.1                                 1           evpn       0
 BUM         Up                                        No          No
1           192.0.2.2                                 1           evpn       0
 BUM         Up                                        No          No
1           192.0.2.4                                 1           evpn       0
 BUM         Up                                        No          No
1           192.0.2.5                                 1           evpn       0
 BUM         Up                                        No          No
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 4
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN-VXLAN Ethernet Segment Dest
===============================================================================
Instance  Eth SegId                       Num. Macs     Last Change
-------------------------------------------------------------------------------
1         00:12:12:12:12:12:12:00:01:01   1             09/27/2021 16:42:17
1         00:12:12:12:12:12:12:00:02:01   1             09/27/2021 16:42:17
1         00:45:45:45:45:45:45:00:00:01   1             09/27/2021 16:42:17
-------------------------------------------------------------------------------
Number of entries: 3
-------------------------------------------------------------------------------
===============================================================================
```

The following command on PE-3 shows the EVPN-VXLAN destination next hops (192.0.2.1 and 192.0.2.2) for alias ESI 00:12:12:12:12:12:12:00:01:01. The VTEP addresses 192.0.2.1 and 192.0.2.2 are the system addresses of PE-1 and PE-2.

```
[/]
A:admin@PE-3# show service id 1 vxlan esi 00:12:12:12:12:12:12:00:01:01

===============================================================================
BGP EVPN-VXLAN Ethernet Segment Dest
===============================================================================
Instance  Eth SegId                       Num. Macs     Last Change
-------------------------------------------------------------------------------
1         00:12:12:12:12:12:12:00:01:01   1             09/27/2021 16:42:17
-------------------------------------------------------------------------------
Number of entries: 1
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN-VXLAN Dest TEP Info
===============================================================================
Instance  TEP Address                   Egr VNI             Last Change
-------------------------------------------------------------------------------
1         192.0.2.1                     1                   09/27/2021 16:42:17
1         192.0.2.2                     1                   09/27/2021 16:42:17
-------------------------------------------------------------------------------
Number of entries : 2
-------------------------------------------------------------------------------
===============================================================================
```

Tools command to check local bias

The following **tools** command on PE-2 checks whether local bias is enabled for the peers in ES "vES-12_1_1". The output lists the PEs that are in the candidate DF election list for the ES and whether

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

local bias procedures are enabled on them. In this case, only peer 192.0.2.1 is in the list and local bias is enabled. The output is similar for ES "vES-12_2_1".

```
[/]
A:admin@PE-2# tools dump service system bgp-evpn ethernet-segment "vES-12_1_1" local-bias
-------------------------------------------------------------------------------
[09/27/2021 16:45:44] Vxlan Local Bias Information
-------------------------------------------------------------------+-----------
Peer                                                               | Enabled
-------------------------------------------------------------------+-----------
192.0.2.1                                                          | Yes
-------------------------------------------------------------------------------
```

The PE can only enable local bias procedures on a maximum of three PEs that are attached to the same ES and use multi-homed VXLAN services. If more than three PEs exist, the PEs are ordered by preference or IP address and only the top three PEs are considered for local bias. The order is as follows:

- lowest IP address (automatic service-carving)
- lowest preference (manual service-carving with configured EVI)
- highest preference (manual service-carving without configured EVI)

The following **tools** command on PE-2 shows that local bias is enabled for peer 192.0.2.3 in ES "vES-23_3_1":

```
[/]
A:admin@PE-2# tools dump service system bgp-evpn ethernet-segment "vES-23_3_1" local-bias
-------------------------------------------------------------------------------
[09/27/2021 16:45:44] Vxlan Local Bias Information
-------------------------------------------------------------------+-----------
Peer                                                               | Enabled
-------------------------------------------------------------------+-----------
192.0.2.3                                                          | Yes
-------------------------------------------------------------------------------
```

Verify local bias for BUM traffic in all-active multi-homing ESs

Unknown unicast traffic is generated on MTU-7. This traffic is received in ingress queue 11 for SAP lag-1:1 on ingress node PE-1. The following **monitor** command — in classic CLI — monitors SAP lag-1:1 in VPLS 1 on PE-1:

```
*A:PE-1# monitor service id 1 sap lag-1:1

===============================================================================
Monitor statistics for Service 1 SAP lag-1:1
===============================================================================
---snip---
-------------------------------------------------------------------------------
Sap per Queue Stats
-------------------------------------------------------------------------------
                        Packets                 Octets

Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio           : 0                       0
Off. LowPrio          : 0                       0
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
For. OutProf          : 0                       0

Ingress Queue 11 (Multipoint) (Priority)
Off. Combined         : 6                       408
Off. Managed          : 0                       0
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 6                       408

Egress Queue 1
For. In/InplusProf    : 0                       0
For. Out/ExcProf      : 0                       0
Dro. In/InplusProf    : 0                       0
Dro. Out/ExcProf      : 0                       0

===============================================================================
```

On the ingress node PE-1, the local bias mechanism forwards this BUM traffic toward EVPN-VXLAN destinations, and also to the local SAPs of all-active ESs, regardless of the DF state. In this case, the local bias mechanism forwards the BUM traffic to lag-2:1 toward MTU-8, even though PE-1 is NDF in ES "vES-12_2_1".

```
*A:PE-1# monitor service id 1 sap lag-2:1

===============================================================================
Monitor statistics for Service 1 SAP lag-2:1
===============================================================================
-------------------------------------------------------------------------------
---snip---
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

                        Packets                 Octets
CPM Ingress           : 0                       0
Forwarding Engine Stats
Dropped               : 0                       0
Received Valid        : 0                       0
Off. HiPrio           : 0                       0
Off. LowPrio          : 0                       0
Off. Uncolor          : 0                       0
Off. Managed          : 0                       0

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Queueing Stats(Egress QoS Policy 1)
Dro. In/InplusProf    : 0                       0
Dro. Out/ExcProf      : 0                       0
For. In/InplusProf    : 0                       0
For. Out/ExcProf      : 6                       408
-------------------------------------------------------------------------------
```

The egress PEs PE-2 and PE-3 receive the BUM traffic on the EVPN-VXLAN terminations. On egress PEs, the local bias mechanism filters BUM traffic based on the source IP address 192.0.2.1 of PE-1. PE-2 does not forward the traffic to the local SAPs lag-1:1 and lag-2:1, because PE-2 shares the all-active ESs

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

"vES-12_1_1" and "vES-12_2_1" with PE-1. However, PE-2 forwards the BUM traffic to the non-shared ES "vES-23_3_1" because it is DF.

The following **monitor** commands show that PE-2 does not send any traffic toward SAP lag-1:1 or SAP lag-2:1.

```
*A:PE-2# monitor service id 1 sap lag-1:1
---snip---

Queueing Stats(Egress QoS Policy 1)
Dro. In/InplusProf    : 0                       0
Dro. Out/ExcProf      : 0                       0
For. In/InplusProf    : 0                       0
For. Out/ExcProf      : 0                       0
---snip---
```

```
*A:PE-2# monitor service id 1 sap lag-2:1
---snip---

Queueing Stats(Egress QoS Policy 1)
Dro. In/InplusProf    : 0                       0
Dro. Out/ExcProf      : 0                       0
For. In/InplusProf    : 0                       0
For. Out/ExcProf      : 0                       0
---snip---
```

The following **monitor** command shows that PE-2 forwards the traffic to SAP lag-3:1 toward MTU-9:

```
*A:PE-2# monitor service id 1 sap lag-3:1
---snip---

Queueing Stats(Egress QoS Policy 1)
Dro. In/InplusProf    : 0                       0
Dro. Out/ExcProf      : 0                       0
For. In/InplusProf    : 0                       0
For. Out/ExcProf      : 6                       408
---snip---
```

Egress node PE-3 receives BUM traffic on VXLAN and filters on IP address 192.0.2.1, but there are no shared ESs with PE-1. PE-3 is NDF for the non-shared ES vES-23_3_1, so it does not forward the traffic to SAP lag-3:1, as follows:

```
*A:PE-3# monitor service id 1 sap lag-3:1
---snip---

Queueing Stats(Egress QoS Policy 1)
Dro. In/InplusProf    : 0                       0
Dro. Out/ExcProf      : 0                       0
For. In/InplusProf    : 0                       0
For. Out/ExcProf      : 0                       0
---snip---
```

EVPN-VXLAN multi-homing with non-system IPv4 VTEP addresses

Figure 133: Non-system IPv4 VTEP multi-homing for VXLAN VPLS 2 shows the non-system IPv4 addresses to be used as VTEP addresses.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 133: Non-system IPv4 VTEP multi-homing for VXLAN VPLS 2

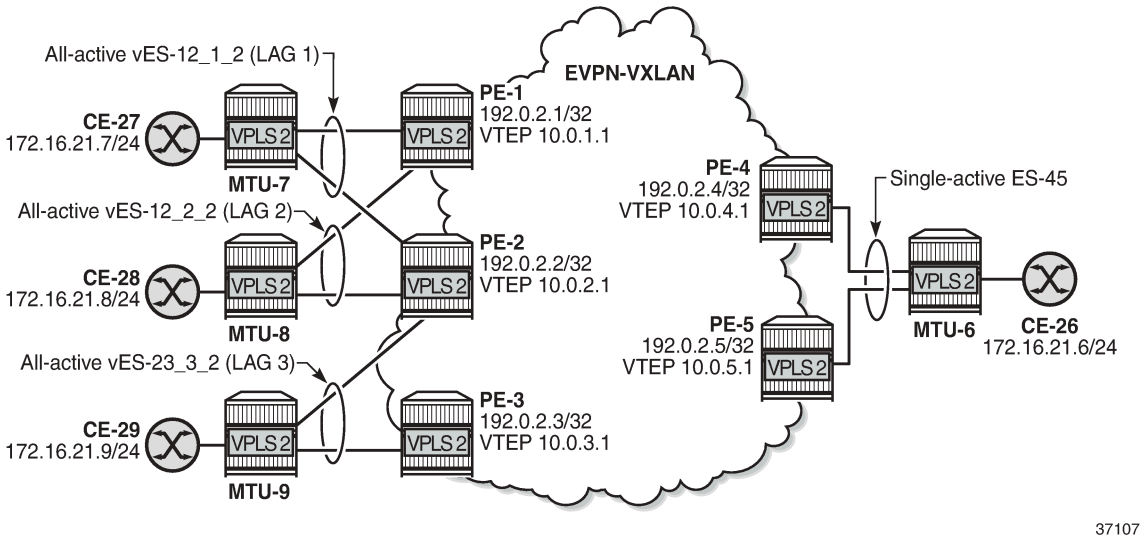

Forwarding Path Extension (FPE), as described in the VXLAN Forwarding Path Extension chapter, is configured on all PEs. The configuration on PE-1 is as follows:

```
# on PE-1:
configure {
    fwd-path-ext {
        sdp-id-range {
            start 10000
            end 10127
        }
        fpe 1 {
            path {
                pxc 1
            }
            application {
                vxlan-termination {
                }
            }
        }
    }
    port 1/2/6 {
        admin-state enable
        ethernet {
            mode hybrid
            dot1x {
                tunneling true
            }
        }
    }
    port pxc-1.a {
        admin-state enable
    }
    port pxc-1.b {
        admin-state enable
    }
    port-xc {
        pxc 1 {
            admin-state enable
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            port-id 1/2/6
        }
    }
    router "Base" {
        interface "loopback1" {
            loopback
            ipv4 {
                primary {
                    address 10.0.1.0
                    prefix-length 31
                }
            }
            ipv6 {
                address 2001:db8::10:0 {
                    prefix-length 127
                }
            }
        }
        isis 0 {
            interface "loopback1"
                passive true
            }
        }
    }
    service {
        system {
            vxlan {
                tunnel-termination 10.0.1.1 {
                    fpe-id 1
                }
                tunnel-termination 2001:db8::10:1 {
                    fpe-id 1
                }
            }
```

The configuration on the other PEs is similar but with different IP addresses, for example, 10.0.2.1 on PE-2, 10.0.3.1 on PE-3, and so on.

The non-system IP address in each of the PEs in the ES must match in the following three commands for the local PE to be considered suitable for DF election:

- orig-ip 10.0.x.1 (ES)

  The **orig-ip** command modifies the originating IP address in the ES routes advertised for the ES and makes the system use this IP address when adding the local PE as DF candidate.

- **route-next-hop** 10.0.x.1 (ES)

  The **route-next-hop** command changes the next hop of the ES routes and AD per-ES routes to the configured address.

- **vxlan source-vtep** 10.0.x.1 (VPLS)

  The **vxlan source-vtep** command makes the router use the configured IP address as the VXLAN tunnel source IP address (source VTEP) for originating VXLAN-encapsulated frames for the service. The source VTEP is also used to set the BGP NLRI next hop in EVPN route advertisements for the services.

The following all-active multi-homing ESs are configured on PE-2 with non-system IPv4 address 10.0.2.1:

```
# on PE-2:
configure {
    service {
        system {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            bgp {
                evpn
                    ethernet-segment "vES-12_1_2" {
                        admin-state enable
                        type virtual
                        esi 00:12:12:12:12:12:12:00:01:02
                        orig-ip 10.0.2.1
                        route-next-hop 10.0.2.1
                        multi-homing-mode all-active
                        df-election {
                            service-carving-mode manual
                            manual {
                                preference {
                                    value 150
                                }
                            }
                        }
                        association {
                            lag "lag-1" {
                                virtual-ranges {
                                    dot1q {
                                        q-tag 2 {
                                            end 2
                                        }
                                    }
                                }
                            }
                        }
                    }
                    ethernet-segment "vES-12_2_2" {
                        admin-state enable
                        type virtual
                        esi 00:12:12:12:12:12:12:00:02:02
                        orig-ip 10.0.2.1
                        route-next-hop 10.0.2.1
                        multi-homing-mode all-active
                        df-election {
                            service-carving-mode manual
                            manual {
                                preference {
                                    value 100
                                }
                            }
                        }
                        association {
                            lag "lag-2" {
                                virtual-ranges {
                                    dot1q {
                                        q-tag 2 {
                                            end 2
                                        }
                                    }
                                }
                            }
                        }
                    }
                    ethernet-segment "vES-23_3_2" {
                        admin-state enable
                        type virtual
                        esi 00:23:23:23:23:23:23:00:03:02
                        orig-ip 10.0.2.1
                        route-next-hop 10.0.2.1
                        multi-homing-mode all-active
                        df-election {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                            service-carving-mode manual
                            manual {
                                preference {
                                    value 100
                                }
                            }
                        }
                        association {
                            lag "lag-3" {
                                virtual-ranges {
                                    dot1q {
                                        q-tag 2 {
                                            end 2
                                        }
                                    }
                                }
                            }
                        }
                    }
                }
            }
```

The ES configuration on the other PEs is similar, but with different IP addresses and preference values.

VPLS 2 is configured with source VTEP 10.0.2.1 on PE-2:

```
# on PE-2:
configure {
    service {
        vpls "VPLS 2" {
            admin-state enable
            service-id 2
            customer "1"
            vxlan {
                source-vtep 10.0.2.1        # different IP address on different PEs
                instance 1 {
                    vni 2
                }
            }
            bgp 1 {
                route-distinguisher auto-rd
                route-target {
                    export "target:64500:2"
                    import "target:64500:2"
                }
            }
            bgp-evpn {
                evi 2
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                    ecmp 2
                    mh-mode network
                    routes {
                        auto-disc {
                            advertise true
                        }
                    }
                }
            }
            sap lag-1:2 {         # lag-1 is shared with PE-1
            }
            sap lag-2:2 {         # lag-2 is shared with PE-1
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            sap lag-3:2 {         # lag-3 is shared with PE-3
            }
        }
```

The configuration on the other PEs is similar.

Verification

The following command shows the DF status for the different ESs in VPLS 2 on PE-1:

```
[/]
A:admin@PE-1# show service id 2 ethernet-segment

===============================================================================
SAP Ethernet-Segment Information
===============================================================================
SAP                   Eth-Seg                         Status
-------------------------------------------------------------------------------
lag-1:2               vES-12_1_2                      NDF
lag-2:2               vES-12_2_2                      DF
===============================================================================
No sdp entries
No vxlan instance entries
```

The following command on PE-1 shows that the source VTEP for VPLS 2 is 10.0.1.1:

```
[/]
A:admin@PE-1# show service id 2 vxlan
===============================================================================
VPLS VXLAN
===============================================================================
Vxlan Src Vtep IP: 10.0.1.1

===============================================================================
Vxlan Instance
===============================================================================
VXLAN Instance                VNI         AR          Oper-flags   VTEP
                                                                   security
-------------------------------------------------------------------------------
1                             2           none        none         disabled
-------------------------------------------------------------------------------
Number of Entries : 1
-------------------------------------------------------------------------------
===============================================================================
```

The following command on PE-1 shows the (non-system) VXLAN destinations for VPLS 2:

```
[/]
A:admin@PE-1# show service id 2 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
Instance    VTEP Address                          Egress VNI  EvpnStatic Num
 Mcast       Oper State                            L2 PBR     SupBcasDom MACs
-------------------------------------------------------------------------------
1           10.0.2.1                              2           evpn       0
 BUM         Up                                    No          No
1           10.0.3.1                              2           evpn       0
 BUM         Up                                    No          No
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
1           10.0.4.1                                  2           evpn       0
 BUM         Up                                        No          No
1           10.0.5.1                                  2           evpn       0
 BUM         Up                                        No          No
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 4
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN-VXLAN Ethernet Segment Dest
===============================================================================
Instance  Eth SegId                         Num. Macs     Last Change
-------------------------------------------------------------------------------
1         00:23:23:23:23:23:23:00:03:02   1             09/27/2021 16:59:29
1         00:45:45:45:45:45:00:00:02   1             09/27/2021 17:00:28
-------------------------------------------------------------------------------
Number of entries: 2
-------------------------------------------------------------------------------
===============================================================================
```

The non-system VTEP addresses in the all-active multi-homing ES with ESI 00:23:23:23:23:23:23:00:03:02 are 10.0.2.1 and 10.0.3.1, as follows:

```
[/]
A:admin@PE-1# show service id 2 vxlan esi 00:23:23:23:23:23:23:00:03:02

===============================================================================
BGP EVPN-VXLAN Ethernet Segment Dest
===============================================================================
Instance  Eth SegId                         Num. Macs     Last Change
-------------------------------------------------------------------------------
1         00:23:23:23:23:23:23:00:03:02   1             09/27/2021 16:59:29
-------------------------------------------------------------------------------
Number of entries: 1
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN-VXLAN Dest TEP Info
===============================================================================
Instance  TEP Address                   Egr VNI             Last Change
-------------------------------------------------------------------------------
1         10.0.2.1                      2                   09/27/2021 16:59:29
1         10.0.3.1                      2                   09/27/2021 16:59:29
-------------------------------------------------------------------------------
Number of entries : 2
-------------------------------------------------------------------------------
===============================================================================
```

EVPN-VXLAN multi-homing with non-system IPv6 VTEP addresses

Figure 134: Non-system IPv6 VTEP multi-homing for VXLAN VPLS 2 shows the non-system IPv6 addresses to be used as VTEP addresses.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 134: Non-system IPv6 VTEP multi-homing for VXLAN VPLS 2

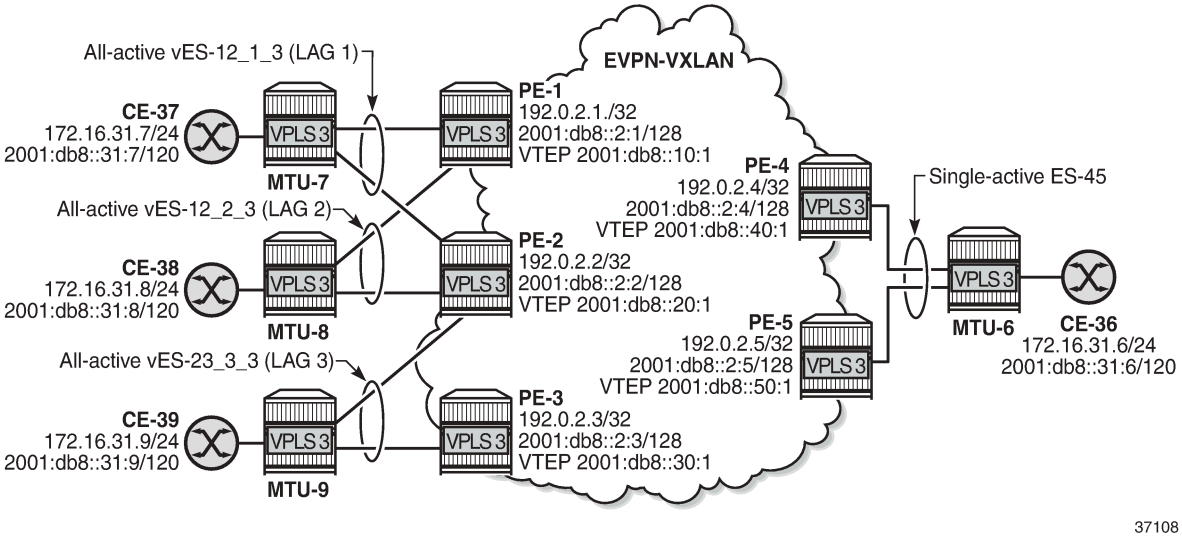

Between the PEs, the router interfaces have IPv6 addresses as well as IPv4 addresses, and **ipv6-routing native** is configured in IS-IS on the PEs. FPE is configured with VXLAN termination 2001:db8::x0:1 on PE-x.

The following all-active multi-homing ESs with non-system IPv6 addresses are configured on PE-2:

```
# on PE-2:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "vES-12_1_3" {
                        admin-state enable
                        type virtual
                        esi 00:12:12:12:12:12:12:00:01:03
                        orig-ip 2001:db8::20:1
                        route-next-hop 2001:db8::20:1
                        multi-homing-mode all-active
                        association {
                            lag "lag-1" {
                                virtual-ranges {
                                    dot1q {
                                        q-tag 3 {
                                            end 3
                                        }
                                    }
                                }
                            }
                        }
                    }
                    ethernet-segment "vES-12_2_3" {
                        admin-state enable
                        type virtual
                        esi 00:12:12:12:12:12:12:00:02:03
                        orig-ip 2001:db8::20:1
                        route-next-hop 2001:db8::20:1
                        multi-homing-mode all-active
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        association {
                            lag "lag-2" {
                                virtual-ranges {
                                    dot1q {
                                        q-tag 3 {
                                            end 3
                                        }
                                    }
                                }
                            }
                        }
                        ethernet-segment "vES-23_3_3" {
                            admin-state enable
                            type virtual
                            esi 00:23:23:23:23:23:23:00:03:03
                            orig-ip 2001:db8::20:1
                            route-next-hop 2001:db8::20:1
                            multi-homing-mode all-active
                            association {
                                lag "lag-3" {
                                    virtual-ranges {
                                        dot1q {
                                            q-tag 3 {
                                                end 3
                                            }
                                        }
                                    }
                                }
                            }
                        }
                    }
                }
```

"VPLS 3" is configured with non-system source VTEP 2001:db8::x0:1, as follows:

```
# on PE-2:
configure {
    service {
        vpls "VPLS 3" {
            admin-state enable
            service-id 3
            customer "1"
            vxlan {
                source-vtep 2001:db8::20:1
                instance 1 {
                    vni 3
                }
            }
            bgp 1 {
                route-distinguisher auto-rd
                route-target {
                    export "target:64500:3"
                    import "target:64500:3"
                }
            }
            bgp-evpn {
                evi 3
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                    ecmp 2
                    mh-mode network
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    routes {
                        auto-disc {
                            advertise true
                        }
                    }
                }
            }
            sap lag-1:3 {         # lag-1 shared with PE-1
            }
            sap lag-2:3 {         # lag-2 shared with PE-1
            }
            sap lag-3:3 {         # lag-3 shared with PE-3
            }
        }
```

Verification

The following command on PE-1 shows that the source VTEP is 2001:db8::10:1 for VPLS 3:

```
[/]
A:admin@PE-1# show service id 3 vxlan
================================================================================
VPLS VXLAN
================================================================================
Vxlan Src Vtep IP: 2001:db8::10:1

================================================================================
Vxlan Instance
================================================================================
VXLAN Instance                VNI         AR          Oper-flags   VTEP
                                                                   security
--------------------------------------------------------------------------------
1                             3           none        none         disabled
--------------------------------------------------------------------------------
Number of Entries : 1
--------------------------------------------------------------------------------
================================================================================
```

The following command on PE-1 shows the non-system IPv6 destination VTEPs for VPLS 3:

```
[/]
A:admin@PE-1# show service id 3 vxlan destinations

================================================================================
Egress VTEP, VNI
================================================================================
Instance    VTEP Address                          Egress VNI  EvpnStatic Num
 Mcast       Oper State                            L2 PBR     SupBcasDom MACs
--------------------------------------------------------------------------------
1           2001:db8::20:1                        3           evpn       0
 BUM         Up                                    No          No
1           2001:db8::30:1                        3           evpn       0
 BUM         Up                                    No          No
1           2001:db8::40:1                        3           evpn       0
 BUM         Up                                    No          No
1           2001:db8::50:1                        3           evpn       0
 BUM         Up                                    No          No
--------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 4
--------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================

===============================================================================
BGP EVPN-VXLAN Ethernet Segment Dest
===============================================================================
Instance  Eth SegId                       Num. Macs     Last Change
-------------------------------------------------------------------------------
1         00:23:23:23:23:23:23:00:03:03   1             09/27/2021 17:20:28
1         00:45:45:45:45:45:00:00:03   1             09/27/2021 17:06:28
-------------------------------------------------------------------------------
Number of entries: 2
-------------------------------------------------------------------------------
===============================================================================
```

The following command on PE-3 shows that VTEPs 2001:db8::10:1 and 2001:db8::20:1 are destinations in the all-active ES with ESI 00:12:12:12:12:12:12:00:01:03:

```
[/]
A:admin@PE-3# show service id 3 vxlan esi 00:12:12:12:12:12:12:00:01:03

===============================================================================
BGP EVPN-VXLAN Ethernet Segment Dest
===============================================================================
Instance  Eth SegId                       Num. Macs     Last Change
-------------------------------------------------------------------------------
1         00:12:12:12:12:12:12:00:01:03   1             09/27/2021 17:28:29
-------------------------------------------------------------------------------
Number of entries: 1
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN-VXLAN Dest TEP Info
===============================================================================
Instance  TEP Address                   Egr VNI             Last Change
-------------------------------------------------------------------------------
1         2001:db8::10:1                3                   09/27/2021 17:28:29
1         2001:db8::20:1                3                   09/27/2021 17:28:29
-------------------------------------------------------------------------------
Number of entries : 2
-------------------------------------------------------------------------------
===============================================================================
```

Debug

With debugging enabled for BGP updates, the following debug message on PE-3 shows that the NextHop value is changed in the EVPN-AD routes:

```
29 2021/09/27 17:36:42.781 CEST MINOR: DEBUG #2001 Base Peer 1: 192.0.2.1
"Peer 1: 192.0.2.1: UPDATE
Peer 1: 192.0.2.1 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 85
    Flag: 0x90 Type: 14 Len: 48 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 16 Global NextHop 2001:db8::30:1
        Type: EVPN-AD Len: 25 RD: 192.0.2.3:3 ESI: 00:23:23:23:23:23:23:00:03:03,
                      tag: MAX-ET Label: 0 (Raw Label: 0x0) PathId:
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:3
        esi-label:524285/All-Active
"
```

The following EVPN-ETH-SEG message on PE-3 shows that the NextHop value and Orig-IP-Addr is modified to the value 2001:db8::30:1.

```
26 2021/09/27 17:36:42.781 CEST MINOR: DEBUG #2001 Base Peer 1: 192.0.2.1
"Peer 1: 192.0.2.1: UPDATE
Peer 1: 192.0.2.1 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 95
    Flag: 0x90 Type: 14 Len: 58 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 16 Global NextHop 2001:db8::30:1        Type: EVPN-ETH-SEG Len: 35 RD:
 192.0.2.3:0
         ESI: 00:23:23:23:23:23:23:00:03:03, IP-Len: 16 Orig-IP-Addr: 2001:db8::30:1
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        df-election::DF-Type:Auto/DP:0/DF-Preference:0/AC:1
        target:23:23:23:23:23:23
"
```

Conclusion

All-active and single-active multi-homing can be configured for EVPN-VXLAN VPLSs. On all-active ESs, split-horizon for BUM traffic is based on local-bias, as described in RFC 8365.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

EVPN VPWS Services with SRv6 Transport

This chapter provides information about SRv6 support for EVPN-VPWS overlay services.

Topics in this chapter include:

Applicability

The information and configuration in this chapter are based on SR OS Release 22.10.R2. SRv6 support for EVPN-VPWS overlay services is supported on FP-based platforms with FP4-based network ports in SR OS Release 22.7.R1 and later.

Chapter EVPN for MPLS Tunnels is prerequisite reading.

Overview

Service providers prefer an optimized, standardized, and unified control plane for VPNs. EVPN-VPWS is supported in SRv6 networks that may also run other EVPN-based services, such as EVPN-based VPLS services or Layer 3 EVPN IFL (interface-less) services. From a control plane perspective, EVPN-VPWS is a simplified point-to-point version of RFC 7432, because there is no need to advertise MAC/IP advertisement routes in VPWS. EVPN-VPWS is described in RFC 8214, and the signaling aspects to support SRv6 are specified in RFC 9252.

EVPN-VPWS supports all-active multihoming (per-flow load-balancing multihoming) as well as single-active multihoming (per-service load-balancing multihoming), using the same Ethernet segments (ESs) used for EVPN-based VPLS services. EVPN-VPWS uses route type 1 and route type 4; it does not use route types 2, 3, or 5, because MAC/IP routes, inclusive multicast routes, or IP-prefix routes are not required.

EVPN-VPWS uses AD per-EVI routes, and optionally, if multihoming is used, AD per-ES and ES routes are required:

- route type 1 - Auto-discovery per EVPN instance (AD per-EVI). This route type is used in all EVPN-VPWS scenarios, with or without multihoming. For EVPN-VPWS, the Ethernet tag field is encoded with the local attachment circuit (AC) of the advertising PE. This value is configured using the **configure service epipe** *<service-name>* **bgp-evpn local-attachment-circuit** *<name>* **eth-tag** *<eth-tag>* command. The route distinguisher (RD), label, and the Ethernet segment identifier (ESI) are encoded as for EVPN-based VPLS. The label field is used as service label. In case of multihoming, AD per-EVI routes containing the same ESI are used to provide aliasing and a backup path to the PEs part of the ES. The L2 MTU field is encoded with the service MTU configured in the Epipe. The flags used for EVPN-VPWS are:
  - Flag C: this flag is set if a control word is configured in the service; however, this does not apply if the transport is SRv6.
  - Flag P: this flag is set if the advertising PE is a primary PE.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- If no multihoming is used, there is no primary PE (P = 0).
- In all-active multihoming, all PEs in the ES are primary (P = 1).
- In single-active multihoming, only one PE per-EVI in the ES is a primary (P = 1).

– Flag B: this flag is set if the advertising PE is a backup PE.

- Flag B is only set in case of single-active multihoming and only for one PE, even if more than two PEs are present in the same single-active ES. The backup PE is the winner of the second designated forwarder (DF) election (excluding the DF). The remaining non-DF PEs send B = 0.

If there is no multihoming, the ESI, flag P, and flag B are set to zero.

- route type 1 - Auto-discovery per Ethernet segment (AD per-ES). This route type has the same encoding as for EVPN-based VPLS. The AD per-ES route is only used in multihoming scenarios where it is advertised from the PE for each ES. This route type carries the ESI label (used for split-horizon, but only for VPLS services and not for Epipe services) and can affect procedures such as the DF election, as well as the aliasing on remote PEs.
- route type 4 - ES route. This route type has the same encoding as for EVPN-based VPLS. The ES route is only used in multihoming scenarios. This route type advertises a local configured ES. The exchange of this route type can discover remote PEs that are part of the same ES and the DF election algorithm among them.

Configuration

Figure 135: EVPN-VPWS example topology shows the example topology that is used throughout this chapter.

Figure 135: EVPN-VPWS example topology

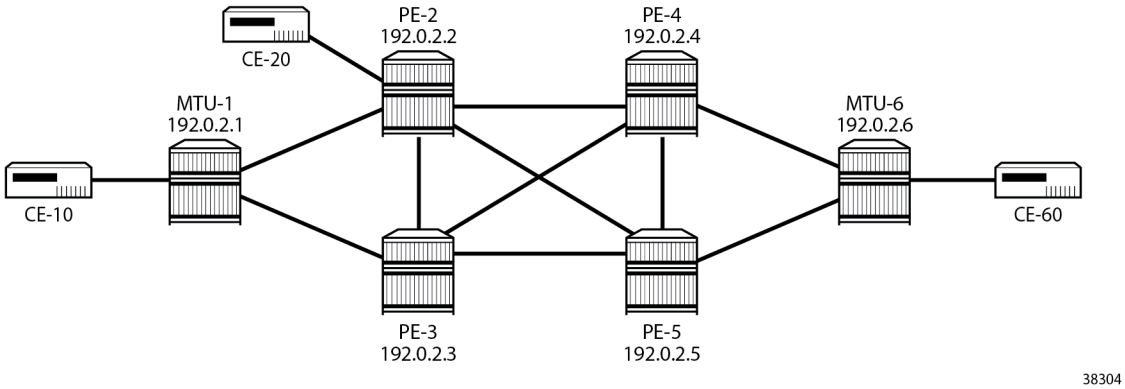

The example topology consists of six SR OS nodes with the following initial configuration:

- Network (or hybrid) ports interconnect the core PEs with configured router interfaces.
- MTU-1 is a pure Ethernet aggregator. The ports toward the core PEs are access ports. Likewise, the ports on PE-2 and PE-3 toward MTU-1 are access ports.
- Core PEs and MTU-6 run IS-IS on all interfaces.
- Link LDP is configured between all PEs, and toward and from MTU-6.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- EVPN uses BGP for exchanging reachability information at the service level. Therefore, BGP peering sessions must be established among the core PEs for the EVPN family. Although a separate router is typically used, in this chapter, PE-2 is used as route reflector with the following BGP configuration:

```
[/]
A:admin@PE-2# configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-import true
            vpn-apply-export true
            peer-ip-tracking true
            rapid-withdrawal true
            split-horizon true
            rapid-update {
                evpn true
            }
            group "gr_v6_internal" {
                family {
                    evpn true
                }
                cluster {
                    cluster-id 1.1.1.1
                }
                peer-as 64500
                extended-nh-encoding {
                    ipv4 true
                    vpn-ipv4 true
                }
                advertise-ipv6-next-hops {
                    evpn true
                }
            }
            neighbor 2001:db8::2:3 {
                group "gr_v6_internal"
            }
            neighbor 2001:db8::2:4 {
                group "gr_v6_internal"
            }
            neighbor 2001:db8::2:5 {
                group "gr_v6_internal"
            }
        }
    }
}
```

The BGP configuration on the other PEs is as follows:

```
[/]
A:admin@PE-3#, A:admin@PE-4#,A:admin@PE-5# configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-import true
            vpn-apply-export true
            peer-ip-tracking true
            rapid-withdrawal true
            split-horizon true
            rapid-update {
                evpn true
            }
            group "gr_v6_internal" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                family {
                    evpn true
                }
                peer-as 64500
                extended-nh-encoding {
                    ipv4 true
                    vpn-ipv4 true
                }
                advertise-ipv6-next-hops {
                    evpn true
                }
            }
            neighbor 2001:db8::2:2 {
                group "gr_v6_internal"
            }
        }
    }
}
```

The following sections describe the EVPN-VPWS scenarios:

- SRv6 tunnels in EVPN-VPWS services without multihoming
- SRv6 tunnels in EVPN-VPWS services with all-active multihoming
- SRv6 tunnels in EVPN-VPWS services with single-active multihoming

SRv6 tunnels in EVPN-VPWS services without multihoming

BGP-EVPN can be enabled in Epipe services with either SAPs or spoke SDPs at the access, as shown in Figure 136: Example topology for EVPN-VPWS without multihoming.

Figure 136: Example topology for EVPN-VPWS without multihoming

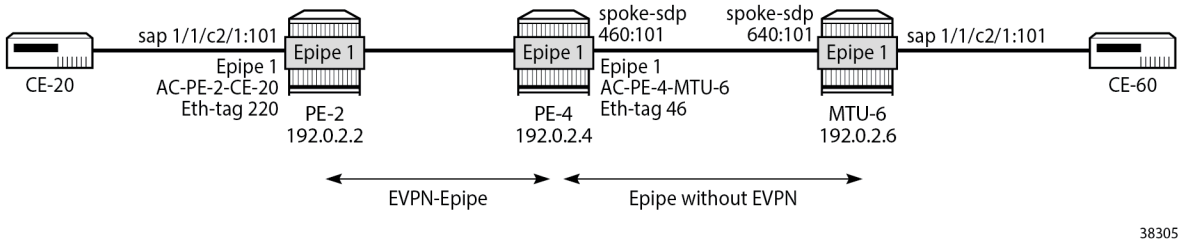

On PE-2, Epipe 1 is configured as follows:

```
[/]
A:admin@PE-2# configure {
    service {
        epipe "Epipe-1" {
            customer "1"
            service-id 1
            segment-routing-v6 1 {
                locator "loc_Epipe-1" {
                    function {
                        end-dx2 {
                        }
                    }
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            bgp 1 {
            }
            bgp-evpn {
                local-attachment-circuit AC-PE-2-CE-20 {
                    eth-tag 220
                }
                remote-attachment-circuit AC-PE-4-MTU-6 {
                    eth-tag 46
                }
                evi 10
                segment-routing-v6 1 {
                    srv6 {
                        instance 1
                        default-locator "loc_Epipe-1"
                    }
                    # source-address 2001:db8::2:2    # defined for SRv6 on router level
                    admin-state enable
                }
            }
            sap 1/1/c2/1:101 {
            }
            admin-state enable
        }
    }
}
```

On PE-4, the service configuration is as follows:

```
[/]
A:admin@PE-4# configure {
    service {
        sdp 460 {
            far-end {
                ip-address 192.0.2.6
            }
            admin-state enable
        }
        epipe "Epipe-1" {
            customer "1"
            service-id 1
            segment-routing-v6 1 {
                locator "loc_Epipe-1" {
                    function {
                        end-dx2 {
                        }
                    }
                }
            }
            bgp 1 {
            }
            bgp-evpn {
                local-attachment-circuit AC-PE-4-MTU-6 {
                    eth-tag 46
                }
                remote-attachment-circuit AC-PE-2-CE-20 {
                    eth-tag 220
                }
                evi 10
                segment-routing-v6 1 {
                    srv6 {
                        instance 1
                        default-locator "loc_Epipe-1"
                    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    # source-address 2001:db8::2:4    # defined for SRv6 on router level
                    admin-state enable
                }
            }
            spoke-sdp 460:101 {
            }
            admin-state enable
        }
    }
}
```

The following commands are relevant for the EVPN-VPWS configuration:

- the **bgp** <*bgp-instance*> command enables the context for the BGP configuration relevant to the service. The **bgp** context configures the common BGP parameters for all BGP families in the service, such as the RD and the route target (RT). Even if the general BGP parameters for the service are auto-derived, the **bgp** context must be enabled.

```
[/]
A:admin@PE-2# configure {
    service {
        epipe "Epipe-1" {
            bgp 1 ?

 bgp

 adv-service-mtu       - Advertised service MTU value
 apply-groups          - Apply a configuration group at this level
 apply-groups-exclude  - Exclude a configuration group at this level
 pw-template-binding   + Enter the pw-template-binding list instance
 route-distinguisher   - High-order 6 bytes that are used as string to compose VSI-ID for
 use in NLRI
 route-target          + Enter the route-target context
 vsi-export            - VSI export policies
 vsi-import            - VSI import policies
```

- The following parameters can be configured in the **bgp-evpn** context:

```
[/]
A:admin@PE-2# configure {
    service {
        epipe "Epipe-1" {
            bgp-evpn ?

 bgp-evpn

 apply-groups          - Apply a configuration group at this level
 apply-groups-exclude  - Exclude a configuration group at this level
 evi                   - EVPN ID
 local-attachment-     + Enter the local-attachment-circuit list instance
  circuit
 mpls                  + Enter the mpls list instance
 remote-attachment-    + Enter the remote-attachment-circuit list instance
  circuit
 segment-routing-v6    + Enter the segment-routing-v6 list instance
 vxlan                 + Enter the vxlan list instance
```

  – The **evi** command configures a 2-byte or 3-byte EVPN identifier (EVI) used for auto-deriving the service RD, service RT, and for the service carving (or DF election) when multihoming is used. For 2-byte EVIs, the auto-derivation of RD and RT is as follows:

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- RD system-ip:evi
- RT autonomous-system:evi

The EVI values must be unique in the system, regardless of the type of service they are assigned to (Epipe or VPLS).

– The **local-attachment-circuit** and **remote-attachment-circuit** commands configure the two attachment circuits connected by the EVPN-VPWS service. The configured Ethernet tag for the local AC is advertised in the Ethernet tag field of the AD per-EVI route for the Epipe, along with the corresponding RD, RT, and label. Both local and remote Ethernet tags are necessary to bring up the Epipe service. If the received Ethernet tag for the Epipe service matches the configured remote AC Ethernet tag, an EVPN-SRv6 destination is created to the next hop.

  The local Ethernet tag cannot be modified without disabling **bgp-evpn segment-routing-v6** in the Epipe, as shown in the following output:

```
[/]
A:admin@PE-2# configure {
    service {
        epipe "Epipe-1" {
            bgp-evpn {
                local-attachment-circuit AC-PE-2-CE-20 {
                    eth-tag 221
                }
            }
        }
    }
}
MINOR: SVCMGR #8036: configure service epipe "Epipe-1" bgp-evpn local-attachment-circuit
 "AC-PE-2-CE-20" - evpn-vpws ac eth-tag not allowed - cannot change while evpn mpls/
vxlan/srv6 is enabled
```

  Unlike local Ethernet tags, remote Ethernet tags can be modified without disabling bgp-evpn.

– The following configuration options are available for Epipes in the **configure service epipe 1 bgp-evpn segment-routing-v6** context:

```
[/]
A:admin@PE-2# configure {
    service {
        epipe "Epipe-1" {
            bgp-evpn {
                segment-routing-v6 1 ?

 segment-routing-v6

 admin-state           - Administrative state of segment routing over IPv6
 apply-groups          - Apply a configuration group at this level
 apply-groups-exclude  - Exclude a configuration group at this level
 default-route-tag     - Default route tag
 ecmp                  - Maximum ECMP value configured on the service
 evi-three-byte-auto-  - Auto-derive the BGP EVPN route target
  rt
 force-vc-forwarding   - Datapath forwarding to force vlan-vc-type
 oper-group            - Operational group
 resolution            - Resolution options for routes
 route-next-hop        + Enter the route-next-hop context
 source-address        - Source IPv6 address
 srv6                  + Enter the srv6 context
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

This output shows a subset of the options for VPLS services; see chapter EVPN for MPLS Tunnels for a longer list of options.

When the local AC (sap 1/1/c2/1:101) is up, PE-2 sends a BGP EVPN AD per-EVI route that contains Ethernet tag 220 for the local AC:

```
# on PE-2:
4 2023/01/10 23:10:54.278 CET MINOR: DEBUG #2001 Base Peer 1: 2001:db8::2:4
"Peer 1: 2001:db8::2:4: UPDATE
Peer 1: 2001:db8::2:4 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 113
    Flag: 0x90 Type: 14 Len: 36 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
        Type: EVPN-AD Len: 25 RD: 192.0.2.2:10 ESI: ESI-0, tag: 220 Label: 8388448 (Raw Label:
 0x7fff60) PathId:
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:10
        l2-attribute:MTU: 1514 C: 0 P: 0 B: 0
    Flag: 0xc0 Type: 40 Len: 37 Prefix-SID-attr:
       SRv6 Services TLV (37 bytes):-
           Type: SRV6 L2 Service TLV (6)
           Length: 34 bytes, Reserved: 0x0
         SRv6 Service Information Sub-TLV (33 bytes)
            Type: 1 Len: 30 Rsvd1: 0x0
            SRv6 SID: 2001:db8:aaaa:102::
            SID Flags: 0x0 Endpoint Behavior: 0x15 Rsvd2: 0x0
            SRv6 SID Sub-Sub-TLV
               Type: 1 Len: 6
               BL:48 NL:16 FL:20 AL:0 TL:20 TO:64
"
```

The auto-derived RD is 192.0.2.2:10 and the RT is 64500:10.

When the remote AC on PE-4 (spoke sdp 460:101) is up, PE-2 receives the following BGP update from PE-4:

```
# on PE-2:
5 2023/01/10 23:11:18.370 CET MINOR: DEBUG #2001 Base Peer 1: 2001:db8::2:4
"Peer 1: 2001:db8::2:4: UPDATE
Peer 1: 2001:db8::2:4 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 113
    Flag: 0x90 Type: 14 Len: 36 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.4
        Type: EVPN-AD Len: 25 RD: 192.0.2.4:10 ESI: ESI-0, tag: 46 Label: 8388448 (Raw Label:
 0x7fff60) PathId:
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:10
        l2-attribute:MTU: 1514 C: 0 P: 0 B: 0
    Flag: 0xc0 Type: 40 Len: 37 Prefix-SID-attr:
       SRv6 Services TLV (37 bytes):-
           Type: SRV6 L2 Service TLV (6)
           Length: 34 bytes, Reserved: 0x0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
         SRv6 Service Information Sub-TLV (33 bytes)
             Type: 1 Len: 30 Rsvd1: 0x0
             SRv6 SID: 2001:db8:aaaa:104::
             SID Flags: 0x0 Endpoint Behavior: 0x15 Rsvd2: 0x0
             SRv6 SID Sub-Sub-TLV
                Type: 1 Len: 6
                BL:48 NL:16 FL:20 AL:0 TL:20 TO:64
"
```

When the received RT matches and the received Ethernet tag matches the configured remote AC Ethernet tag, the EVPN-SRv6 destination, which consists of a termination endpoint (TEP) and a SID, is created on PE-2 and PE-4:

```
[/]
A:admin@PE-2# show service id 1 segment-routing-v6 instance 1 destinations

===============================================================================
TEP, SID
===============================================================================
Instance  TEP Address                        Segment Id
-------------------------------------------------------------------------------
1         192.0.2.4                          2001:db8:aaaa:104:7fff:6000::
-------------------------------------------------------------------------------
Number of TEP, SID: 1
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
Segment Routing v6 Ethernet Segment Dest
===============================================================================
Instance  Eth SegId                      Num. Macs     Last Change
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

Note:
The egress label for the EVPN-SRv6 destination on PE-4 is 524278. The 24-bit label value in the BGP update debug is 16 (2^4) times as high:

524 278 * 16 = 8 388 448

because the debug message is shown before the router can parse the label field and determine if it corresponds to an MPLS label or a transposed function (20 bits), or to a VXLAN VNI (24 bits).

The BGP AD per-EVI routes for Ethernet tag 46 are shown with the following command:

```
[/]
A:admin@PE-2# show router bgp routes evpn auto-disc tag 46
===============================================================================
 BGP Router ID:192.0.2.2        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Auto-Disc Routes
===============================================================================
Flag  Route Dist.         ESI                           NextHop
      Tag                                               Label
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
u*>i  192.0.2.4:10        ESI-0                          192.0.2.4
      46                                                 524278

-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

The following command shows the BGP EVPN information for Epipe 1:

```
[/]
A:admin@PE-2# show service id 1 bgp-evpn

===============================================================================
BGP EVPN Table
===============================================================================
EVI                : 10                Creation Origin    : manual

-------------------------------------------------------------------------------
Local AC Name              Eth Tag  Endpoint                   Ingress Label
-------------------------------------------------------------------------------
AC-PE-2-CE-20              220                                 0
-------------------------------------------------------------------------------
Number of local ACs : 1

-------------------------------------------------------------------------------
Remote AC Name                     Eth Tag  Endpoint
-------------------------------------------------------------------------------
AC-PE-4-MTU-6                      46
-------------------------------------------------------------------------------
Number of Remote ACs : 1
===============================================================================

===============================================================================
Segment Routing v6 Instance 1 Service 1
===============================================================================
Admin State             : Enabled
Srv6 Instance           : 1
Default Locator         : loc_Epipe-1

Oper Group              : (Not Specified)
Default Route Tag       : 0x0
Source Address          : (Not Specified)
ECMP                    : 1
Force Vlan VC Fwd       : disabled
Next Hop Type           : system-ipv4
Evi 3-byte Auto-RT      : disabled
Route Resolution        : route-table
Force QinQ VC Fwd       : none
MH Mode                 : network
===============================================================================
```

Note:
Each PE sends its service MTU into the L2 MTU field in the l2-attribute in the AD per-EVI route for the Epipe service. The received L2 MTU is checked. In case of a mismatch between the received MTU and the configured service MTU, the router does not set up the EVPN destination and, therefore, the service does not come up.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

SRv6 tunnels in EVPN-VPWS services with multihoming

SR OS supports EVPN multihoming as per RFC 8214.

The EVPN multihoming implementation is based on the concept of the ES. An ES is a logical structure that can be defined in one or more PEs and identifies the CE (or access network) multihoming to the EVPN PEs. An ES is associated with a port, LAG, or SDP object, and is shared by all the services defined on those objects. It can also be shared between Epipe and VPLS services.

Each ES has a unique ESI that is 10 bytes and is manually configured. The ESI is advertised in the control plane to all the PEs in an EVPN network; therefore, it is very important to ensure that the 10-byte ESI value is unique throughout the entire network. Single-homing CEs are assumed to be connected to an ES with ESI = 0 (single-homing ESs are not explicitly configured).

The ES is part of the base BGP-EVPN configuration and is not applied to any EVPN-based VPLS service by default. An ES can be shared by multiple services; a specific SAP or spoke SDP is automatically associated with an ES when the SAP is defined in the same LAG or port configured in the ES, or when the spoke SDP is defined in the same SDP configured in the ES.

Regardless of the multihoming mode, the local Ethernet tag values must match on all the PEs that are part of the same ES. The PEs in the ES use the AD per-EVI routes from the peer PEs to validate the PEs as DF election candidates for an EVI. The DF election is only relevant for single-active multihoming ESs. For Epipes defined in an all-active multihoming ES, there is no DF election required, because all PEs are forwarding traffic and all traffic is treated as unicast.

Aliasing is supported when sending traffic to an ES destination. Assuming ECMP is enabled on the ingress PE (and shared queuing or ingress policing are configured), per-flow load-balancing is performed among all the PEs that advertise P = 1. PEs advertising P = 0 are not considered as next hops for an ES destination.

The following sections show the configuration of:

- an all-active multihoming ES with a LAG associated with it
- a single-active multihoming ES linked to an SDP

Figure 137: Example topology EVPN-VPWS with multihoming shows the example topology has an all-active multihoming ES "ESI-23" with a LAG associated with it in PE-2 and PE-3. A single-active multihoming ES "ESI-45" with an SDP associated with it is configured in PE-4 and PE-5.

Figure 137: Example topology EVPN-VPWS with multihoming

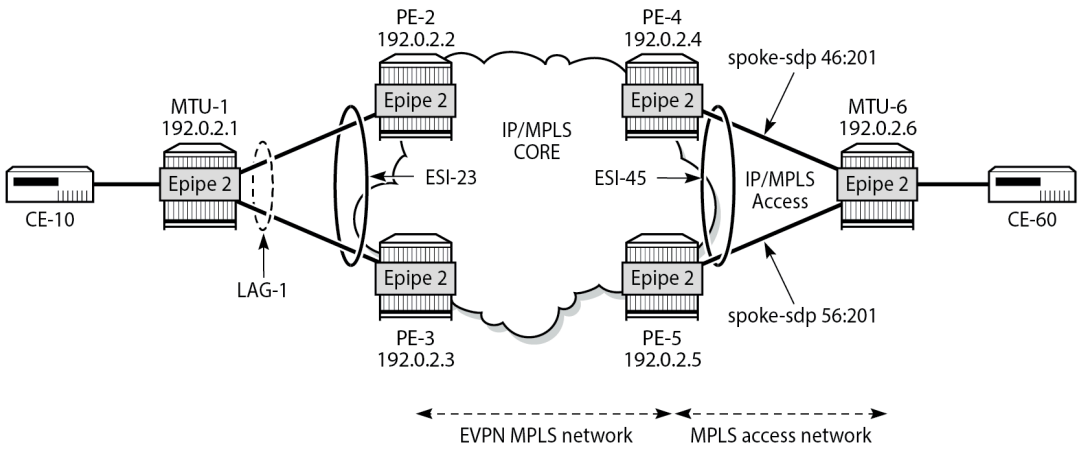

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

SRv6 tunnels in EVPN-VPWS services with all-active multihoming

All-active multihoming allows for per-flow load-balancing. Unlike EVPN-based VPLS services, EVPN-VPWS has no DF election in all-active multihoming. All PEs in the ES are active and the remote PE performs per-flow load-balancing. ESI-23 is configured on PE-2 and PE-3 as all-active multihoming and is associated with LAG 1. This LAG is used as a SAP in Epipe 2 on both PE-2 and PE-3. The configuration of the ES and Epipe 2 is identical on PE-2 and PE-3, including the local AC and remote AC names and Ethernet tags:

```
[/]
A:admin@PE-2#, A:admin@PE-3# configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "ESI-23" {
                        esi 01:00:00:00:00:23:00:00:00:01
                        df-election {
                            es-activation-timer 3
                        }
                        multi-homing-mode all-active
                        association {
                            lag "lag-1" {
                            }
                        }
                        admin-state enable
                    }
                }
            }
        }
        epipe "Epipe-2" {
            customer "1"
            service-id 2
            segment-routing-v6 1 {
                locator "loc_Epipe-2" {
                    function {
                        end-dx2 {
                        }
                    }
                }
            }
            bgp 1 {
            }
            bgp-evpn {
                local-attachment-circuit AC-ESI-23-MTU-1 {
                    eth-tag 231
                }
                remote-attachment-circuit AC-ESI-45-MTU-6 {
                    eth-tag 456
                }
                evi 20
                segment-routing-v6 1 {
                    srv6 {
                        instance 1
                        default-locator "loc_Epipe-2"
                    }
                    ecmp 2
                    # source-address 2001:db8::2:2    # defined for SRv6 on router level
                    admin-state enable
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            sap lag-1:201 {
            }
            admin-state enable
        }
    }
}
```

See chapter EVPN for MPLS Tunnels for a detailed explanation of the configuration parameters of the ES.

In EVPN-VPWS multihoming scenarios, three route types are exchanged: AD per-EVI, AD per-ES, and ES routes. The following ES route (route type 4) for ESI 01:00:00:00:00:23:00:00:00:01, sent by PE-2, is imported at PE-3:

```
# on PE-3:
8 2023/01/10 23:27:55.229 CET MINOR: DEBUG #2001 Base Peer 1: 2001:db8::2:2
"Peer 1: 2001:db8::2:2: UPDATE
Peer 1: 2001:db8::2:2 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 71
    Flag: 0x90 Type: 14 Len: 34 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
        Type: EVPN-ETH-SEG Len: 23 RD: 192.0.2.2:0 ESI: 01:00:00:00:00:23:00:00:00:01, IP-Len:
 4 Orig-IP-Addr: 192.0.2.2
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        df-election::DF-Type:Auto/DP:0/DF-Preference:0/AC:1
        target:00:00:00:00:23:00
"
```

The target 00:00:00:00:23:00 in the extended community is derived from the ESI (bytes 2 to 7) and is only imported by the PEs that are part of the same ES; that is, PE-2 and PE-3 in this example.

At the same time, the following AD per-ES route (route type 1) with maximum Ethernet (MAX-ET) tag (all Fs) and label 0 is sent by RR PE-2 and imported by the rest of the PEs. The following two BGP updates with MAX-ET are received by PE-4:

```
# on PE-4:
15 2023/01/10 23:28:34.279 CET MINOR: DEBUG #2001 Base Peer 1: 2001:db8::2:2
"Peer 1: 2001:db8::2:2: UPDATE
Peer 1: 2001:db8::2:2 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 113
    Flag: 0x90 Type: 14 Len: 36 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
        Type: EVPN-AD Len: 25 RD: 192.0.2.2:20 ESI: 01:00:00:00:00:23:00:00:00:01, tag: MAX-ET
 Label: 0 (Raw Label: 0x0) PathId:
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:20
        esi-label:3/All-Active
    Flag: 0xc0 Type: 40 Len: 37 Prefix-SID-attr:
       SRv6 Services TLV (37 bytes):-
           Type: SRV6 L2 Service TLV (6)
           Length: 34 bytes, Reserved: 0x0
         SRv6 Service Information Sub-TLV (33 bytes)
             Type: 1 Len: 30 Rsvd1: 0x0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                Type: 1 Len: 6
                BL:0 NL:0 FL:0 AL:0 TL:0 TO:0
"
```

```
13 2023/01/10 23:28:34.279 CET MINOR: DEBUG #2001 Base Peer 1: 2001:db8::2:2
"Peer 1: 2001:db8::2:2: UPDATE
Peer 1: 2001:db8::2:2 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 127
    Flag: 0x90 Type: 14 Len: 36 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.3
        Type: EVPN-AD Len: 25 RD: 192.0.2.3:20 ESI: 01:00:00:00:00:23:00:00:00:01, tag: MAX-ET
 Label: 0 (Raw Label: 0x0) PathId:
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0x80 Type: 9 Len: 4 Originator ID: 192.0.2.3
    Flag: 0x80 Type: 10 Len: 4 Cluster ID:
        1.1.1.1
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:20
        esi-label:3/All-Active
    Flag: 0xc0 Type: 40 Len: 37 Prefix-SID-attr:
       SRv6 Services TLV (37 bytes):-
           Type: SRV6 L2 Service TLV (6)
                Type: 1 Len: 6
                BL:0 NL:0 FL:0 AL:0 TL:0 TO:0
"
```

The ESI label is in the extended community, as well as the indication that the multihoming is all-active. Epipe services do not require ESI labels because BUM traffic is not recognized in EVPN-VPWS services. However, because the ES can be shared by Epipe and VPLS services, the AD per-ES route still includes a non-zero ESI label. In this case, the transport is SRv6, so there are no ESI labels. The label field in the ESI-label extended community is an implicit-null value (3) and the included SRv6 Services TLV encodes a SID with value 0.

The following two AD per-EVI routes (route type 1) with Ethernet tag 231 sent by RR PE-2 are received and imported on PE-4:

```
# on PE-4:
14 2023/01/10 23:28:34.279 CET MINOR: DEBUG #2001 Base Peer 1: 2001:db8::2:2
"Peer 1: 2001:db8::2:2: UPDATE
Peer 1: 2001:db8::2:2 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 113
    Flag: 0x90 Type: 14 Len: 36 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
        Type: EVPN-AD Len: 25 RD: 192.0.2.2:20 ESI: 01:00:00:00:00:23:00:00:00:01, tag: 231
 Label: 8388432 (Raw Label: 0x7fff50) PathId:
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:20
        l2-attribute:MTU: 1514 C: 0 P: 1 B: 0
    Flag: 0xc0 Type: 40 Len: 37 Prefix-SID-attr:
       SRv6 Services TLV (37 bytes):-
           Type: SRV6 L2 Service TLV (6)
           Length: 34 bytes, Reserved: 0x0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
         SRv6 Service Information Sub-TLV (33 bytes)
             Type: 1 Len: 30 Rsvd1: 0x0
                Type: 1 Len: 6
                BL:48 NL:16 FL:20 AL:0 TL:20 TO:64
"
```

```
12 2023/01/10 23:28:34.279 CET MINOR: DEBUG #2001 Base Peer 1: 2001:db8::2:2
"Peer 1: 2001:db8::2:2: UPDATE
Peer 1: 2001:db8::2:2 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 127
    Flag: 0x90 Type: 14 Len: 36 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.3
        Type: EVPN-AD Len: 25 RD: 192.0.2.3:20 ESI: 01:00:00:00:00:23:00:00:00:01, tag: 231
 Label: 8388432 (Raw Label: 0x7fff50) PathId:
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0x80 Type: 9 Len: 4 Originator ID: 192.0.2.3
    Flag: 0x80 Type: 10 Len: 4 Cluster ID:
        1.1.1.1
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:20
        l2-attribute:MTU: 1514 C: 0 P: 1 B: 0
    Flag: 0xc0 Type: 40 Len: 37 Prefix-SID-attr:
       SRv6 Services TLV (37 bytes):-
           Type: SRV6 L2 Service TLV (6)
                Type: 1 Len: 6
                BL:48 NL:16 FL:20 AL:0 TL:20 TO:64
"
```

This route type contains the flags for control word (C), primary (P), and backup (B). In all-active multihoming, all nodes are primary (P = 1).

PE-4 learns AD per-EVI and AD per-ES routes for ESI-23 from PE-2 and PE-3, as shown in the following output:

```
[/]
A:admin@PE-4# show router bgp routes evpn auto-disc esi 01:00:00:00:00:23:00:00:00:01
===============================================================================
 BGP Router ID:192.0.2.4        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Auto-Disc Routes
===============================================================================
Flag  Route Dist.         ESI                           NextHop
      Tag                                               Label
-------------------------------------------------------------------------------
u*>i  192.0.2.2:20        01:00:00:00:00:23:00:00:00:01 192.0.2.2
      231                                               524277

u*>i  192.0.2.2:20        01:00:00:00:00:23:00:00:00:01 192.0.2.2
      MAX-ET                                            0

u*>i  192.0.2.3:20        01:00:00:00:00:23:00:00:00:01 192.0.2.3
      231                                               524277
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
u*>i  192.0.2.3:20        01:00:00:00:00:23:00:00:00:01 192.0.2.3
      MAX-ET                                          0

-------------------------------------------------------------------------------
Routes : 4
===============================================================================
```

For Epipe 2 on PE-4, the EVPN VPWS destination is not pointing at a specific TEP, but at ESI-23, as shown in the following output:

```
[/]
A:admin@PE-4# show service id 2 segment-routing-v6 instance 1 destinations

===============================================================================
TEP, SID
===============================================================================
Instance  TEP Address                       Segment Id
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================

===============================================================================
Segment Routing v6 Ethernet Segment Dest
===============================================================================
Instance  Eth SegId                       Num. Macs     Last Change
-------------------------------------------------------------------------------
1         01:00:00:00:00:23:00:00:00:01   0             01/10/2023 23:28:34
-------------------------------------------------------------------------------
Number of entries: 1
-------------------------------------------------------------------------------
===============================================================================
```

When ECMP is greater than 1 on the ingress PE, multiple TEPs can correspond to a specific ESI (aliasing). In this case, ECMP = 2 and PE-4 and PE-5 have two TEP addresses and SIDs for ESI 01:00:00:00:00:23:00:00:00:01, as shown for PE-4:

```
[/]
A:admin@PE-4# show service id 2 segment-routing-v6 esi 01:00:00:00:00:23:00:00:00:01

===============================================================================
Segment Routing v6 Ethernet Segment Dest
===============================================================================
Instance  Eth SegId                       Num. Macs     Last Change
-------------------------------------------------------------------------------
1         01:00:00:00:00:23:00:00:00:01   0             01/10/2023 23:28:34
-------------------------------------------------------------------------------
Number of entries: 1
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
Segment Routing v6 Dest TEP Info
===============================================================================
Instance  TEP Address                     Segment Id          Last Change
-------------------------------------------------------------------------------
1         192.0.2.2                       2001:db8:aaaa:202:* 01/10/2023 23:28:34
1         192.0.2.3                       2001:db8:aaaa:203:* 01/10/2023 23:28:34
-------------------------------------------------------------------------------
Number of entries : 2
-------------------------------------------------------------------------------
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
* indicates that the corresponding row element may have been truncated.
```

Note:
Even if ECMP is configured, the ingress router (where a SAP is configured) does not load-balance the traffic unless shared queuing or ingress policing is configured in the SAP. This is not specific to EVPN, but is generic to the way Epipes forward traffic.

In all-active multihoming for EVPN-VPWS, there is no DF election and all PEs in the ES are active. For ESI-23, both PE-2 and PE-3 are active primary DF, but there are no DF candidates, because there is no DF election:

```
[/]
A:admin@PE-2# show service system bgp-evpn ethernet-segment name "ESI-23" evi evi-1 20

===============================================================================
EVI DF and Candidate List
===============================================================================
EVI           SvcId         Actv Timer Rem      DF  DF Last Change
-------------------------------------------------------------------------------
20            2             0                   yes 01/10/2023 23:27:40
===============================================================================

===============================================================================
DF Candidates                              Time Added           Oper Pref  Do Not
                                                                  Value    Preempt
-------------------------------------------------------------------------------
No entries found
===============================================================================
```

Similarly, on PE-3:

```
[/]
A:admin@PE-3# show service system bgp-evpn ethernet-segment name "ESI-23" evi evi-1 20

===============================================================================
EVI DF and Candidate List
===============================================================================
EVI           SvcId         Actv Timer Rem      DF  DF Last Change
-------------------------------------------------------------------------------
20            2             0                   yes 01/10/2023 23:27:54
===============================================================================

===============================================================================
DF Candidates                              Time Added           Oper Pref  Do Not
                                                                  Value    Preempt
-------------------------------------------------------------------------------
No entries found
===============================================================================
```

To confirm that all-active multihoming is working correctly, the following command shows all information related to a specific ESI; in this case, ESI-23 on PE-2:

```
[/]
A:admin@PE-2# show service system bgp-evpn ethernet-segment name "ESI-23" all

===============================================================================
Service Ethernet Segment
===============================================================================
Name                    : ESI-23
Eth Seg Type            : None
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Admin State             : Enabled            Oper State         : Up
ESI                     : 01:00:00:00:00:23:00:00:00:01
Oper ESI                : 01:00:00:00:00:23:00:00:00:01
Auto-ESI Type           : None
AC DF Capability        : Include
Multi-homing            : allActive          Oper Multi-homing  : allActive
ES SHG Label            : 524277
Source BMAC LSB         : None
Lag                     : lag-1
ES Activation Timer     : 3 secs
Oper Group              : (Not Specified)
Svc Carving             : auto               Oper Svc Carving   : auto
Cfg Range Type          : primary
Vprn NextHop EVI Ranges : <none>
===============================================================================

===============================================================================
EVI Information
===============================================================================
EVI                 SvcId               Actv Timer Rem      DF
-------------------------------------------------------------------------------
20                  2                   0                   yes
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
---snip---
===============================================================================
```

SRv6 tunnels in EVPN-VPWS services with single-active multihoming

Single-active multihoming allows for per-service load-balancing. Single-active multihoming is configured on PE-4 and PE-5 with ES "ESI-45". Both PEs have an SDP to MTU-6, which is associated with the ES and to the Epipe service. The configuration of the local and remote AC names and Ethernet tags is identical on PE-4 and PE-5.

On PE-4, the service configuration is as follows:

```
[/]
A:admin@PE-4# configure {
    service {
        sdp 46 {
            delivery-type mpls
            far-end {
                ip-address 192.0.2.6
            }
            ldp true
            admin-state enable
        }
        system {
            bgp {
                evpn {
                    ethernet-segment "ESI-45" {
                        esi 01:00:00:00:00:45:00:00:00:01
                        df-election {
                            es-activation-timer 3
                        }
                        multi-homing-mode single-active
                        association {
                            sdp 46 {
                            }
                        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        admin-state enable
                    }
                }
            }
        }
        epipe "Epipe-2" {
            customer "1"
            service-id 2
            segment-routing-v6 1 {
                locator "loc_Epipe-2" {
                    function {
                        end-dx2 {
                        }
                    }
                }
            }
            bgp 1 {
            }
            bgp-evpn {
                local-attachment-circuit AC-ESI-45-MTU-6 {
                    eth-tag 456
                }
                remote-attachment-circuit AC-ESI-23-MTU-1 {
                    eth-tag 231
                }
                evi 20
                segment-routing-v6 1 {
                    srv6 {
                        instance 1
                        default-locator "loc_Epipe-2"
                    }
                    ecmp 2
                    # source-address 2001:db8::2:4    # defined for SRv6 on router level
                    admin-state enable
                }
            }
            spoke-sdp 46:201 {
            }
            admin-state enable
        }
    }
}
```

On PE-5, the configuration is similar, but with a different SDP:

```
[/]
A:admin@PE-5# configure {
    service {
        sdp 56 {
            delivery-type mpls
            far-end {
                ip-address 192.0.2.6
            }
            ldp true
            admin-state enable
        }
        system {
            bgp {
                evpn {
                    ethernet-segment "ESI-45" {
                        esi 01:00:00:00:00:45:00:00:00:01
                        df-election {
                            es-activation-timer 3
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        }
                        multi-homing-mode single-active
                        association {
                            sdp 56 {
                            }
                        }
                        admin-state enable
                    }
                }
            }
        }
        epipe "Epipe-2" {
            customer "1"
            service-id 2
            segment-routing-v6 1 {
                locator "loc_Epipe-2" {
                    function {
                        end-dx2 {
                        }
                    }
                }
            }
            bgp 1 {
            }
            bgp-evpn {
                local-attachment-circuit AC-ESI-45-MTU-6 {
                    eth-tag 456
                }
                remote-attachment-circuit AC-ESI-23-MTU-1 {
                    eth-tag 231
                }
                evi 20
                segment-routing-v6 1 {
                    srv6 {
                        instance 1
                        default-locator "loc_Epipe-2"
                    }
                    ecmp 2
                    # source-address 2001:db8::2:5    # defined for SRv6 on router level
                    admin-state enable
                }
            }
            spoke-sdp 56:201 {
            }
            admin-state enable
        }
    }
}
```

Three route types are exchanged between the core PEs: AD per-EVI, AD per-ES, and ES routes.

As an example, the following is the ES route with originator PE-4 sent by RR PE-2 to PE-5. It contains a target 00:00:00:00:45:00 in the extended community that is derived from the ESI:

```
# on PE-2:
45 2023/01/10 23:29:07.514 CET MINOR: DEBUG #2001 Base Peer 1: 2001:db8::2:5
"Peer 1: 2001:db8::2:5: UPDATE
Peer 1: 2001:db8::2:5 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 85
    Flag: 0x90 Type: 14 Len: 34 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.4
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        Type: EVPN-ETH-SEG Len: 23 RD: 192.0.2.4:0 ESI: 01:00:00:00:00:45:00:00:00:01, IP-Len:
 4 Orig-IP-Addr: 192.0.2.4
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0x80 Type: 9 Len: 4 Originator ID: 192.0.2.4
    Flag: 0x80 Type: 10 Len: 4 Cluster ID:
        1.1.1.1
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        df-election::DF-Type:Auto/DP:0/DF-Preference:0/AC:1
        target:00:00:00:00:45:00
"
```

The AD per-ES route has a MAX-ET tag and an ESI label in the extended community. The multihoming mode is single-active. As in the case of all-active multihoming, the ESI label is not used in Epipe services. The following BGP update with originator PE-5 is sent by RR PE-2 to its client PE-4:

```
# on PE-2:
51 2023/01/10 23:29:07.669 CET MINOR: DEBUG #2001 Base Peer 1: 2001:db8::2:4
"Peer 1: 2001:db8::2:4: UPDATE
Peer 1: 2001:db8::2:4 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 127
    Flag: 0x90 Type: 14 Len: 36 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.5
        Type: EVPN-AD Len: 25 RD: 192.0.2.5:20 ESI: 01:00:00:00:00:45:00:00:00:01, tag: MAX-ET
 Label: 0 (Raw Label: 0x0) PathId:
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0x80 Type: 9 Len: 4 Originator ID: 192.0.2.5
    Flag: 0x80 Type: 10 Len: 4 Cluster ID:
        1.1.1.1
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:20
        esi-label:3/Single-Active
    Flag: 0xc0 Type: 40 Len: 37 Prefix-SID-attr:
       SRv6 Services TLV (37 bytes):-
           Type: SRV6 L2 Service TLV (6)
                Type: 1 Len: 6
                BL:0 NL:0 FL:0 AL:0 TL:0 TO:0
"
```

The AD per-EVI route contains flags for primary and backup, which are different for routes received from PE-4 and PE-5. In this case, PE-4 is the primary in the single-active multihoming ES (P = 1):

```
# on PE-2:
53 2023/01/10 23:29:10.550 CET MINOR: DEBUG #2001 Base Peer 1: 2001:db8::2:5
"Peer 1: 2001:db8::2:5: UPDATE
Peer 1: 2001:db8::2:5 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 127
    Flag: 0x90 Type: 14 Len: 36 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.4
        Type: EVPN-AD Len: 25 RD: 192.0.2.4:20 ESI: 01:00:00:00:00:45:00:00:00:01, tag: 456
 Label: 8388416 (Raw Label: 0x7fff40) PathId:
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0x80 Type: 9 Len: 4 Originator ID: 192.0.2.4
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    Flag: 0x80 Type: 10 Len: 4 Cluster ID:
        1.1.1.1
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:20
        l2-attribute:MTU: 1514 C: 0 P: 1 B: 0
    Flag: 0xc0 Type: 40 Len: 37 Prefix-SID-attr:
       SRv6 Services TLV (37 bytes):-
           Type: SRV6 L2 Service TLV (6)
                Type: 1 Len: 6
                BL:48 NL:16 FL:20 AL:0 TL:20 TO:64
"
```

PE-5 is the backup in the single-active multihoming ES (B = 1):

```
# on PE-2:
61 2023/01/10 23:29:20.570 CET MINOR: DEBUG #2001 Base Peer 1: 2001:db8::2:5
"Peer 1: 2001:db8::2:5: UPDATE
Peer 1: 2001:db8::2:5 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 113
    Flag: 0x90 Type: 14 Len: 36 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.5
        Type: EVPN-AD Len: 25 RD: 192.0.2.5:20 ESI: 01:00:00:00:00:45:00:00:00:01, tag: 456
 Label: 8388448 (Raw Label: 0x7fff60) PathId:
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:20
        l2-attribute:MTU: 1514 C: 0 P: 0 B: 1
    Flag: 0xc0 Type: 40 Len: 37 Prefix-SID-attr:
       SRv6 Services TLV (37 bytes):-
           Type: SRV6 L2 Service TLV (6)
           Length: 34 bytes, Reserved: 0x0
         SRv6 Service Information Sub-TLV (33 bytes)
             Type: 1 Len: 30 Rsvd1: 0x0
                Type: 1 Len: 6
                BL:48 NL:16 FL:20 AL:0 TL:20 TO:64
"
```

The BGP EVPN AD routes are shown with the following command:

```
[/]
A:admin@PE-2# show router bgp routes evpn auto-disc esi 01:00:00:00:00:45:00:00:00:01
===============================================================================
 BGP Router ID:192.0.2.2        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Auto-Disc Routes
===============================================================================
Flag  Route Dist.         ESI                           NextHop
      Tag                                               Label
-------------------------------------------------------------------------------
u*>i  192.0.2.4:20        01:00:00:00:00:45:00:00:00:01 192.0.2.4
      456                                               524276

u*>i  192.0.2.4:20        01:00:00:00:00:45:00:00:00:01 192.0.2.4
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
      MAX-ET                                              0

u*>i  192.0.2.5:20        01:00:00:00:00:45:00:00:00:01 192.0.2.5
      456                                                 524278

u*>i  192.0.2.5:20        01:00:00:00:00:45:00:00:00:01 192.0.2.5
      MAX-ET                                              0

-------------------------------------------------------------------------------
Routes : 4
===============================================================================
```

For each PE in the single-active ES, there are two AD routes: the routes with MAX-ET are AD per-ES routes and the routes with a configured Ethernet tag are AD per-EVI routes.

The EVPN VPWS destination for Epipe 2 on PE-2 is ESI-45, as shown in the following output:

```
[/]
A:admin@PE-2# show service id 2 segment-routing-v6 instance 1 destinations

===============================================================================
TEP, SID
===============================================================================
Instance  TEP Address                       Segment Id
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================

===============================================================================
Segment Routing v6 Ethernet Segment Dest
===============================================================================
Instance  Eth SegId                       Num. Macs     Last Change
-------------------------------------------------------------------------------
1         01:00:00:00:00:45:00:00:00:01   0             01/10/2023 23:29:11
-------------------------------------------------------------------------------
Number of entries: 1
-------------------------------------------------------------------------------
===============================================================================
```

The ESI is resolved to the TEP address of the primary (DF) PE-4, as follows:

```
[/]
A:admin@PE-2# show service id 2 segment-routing-v6 esi 01:00:00:00:00:45:00:00:00:01

===============================================================================
Segment Routing v6 Ethernet Segment Dest
===============================================================================
Instance  Eth SegId                       Num. Macs     Last Change
-------------------------------------------------------------------------------
1         01:00:00:00:00:45:00:00:00:01   0             01/10/2023 23:29:11
-------------------------------------------------------------------------------
Number of entries: 1
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
Segment Routing v6 Dest TEP Info
===============================================================================
Instance  TEP Address                     Segment Id          Last Change
-------------------------------------------------------------------------------
1         192.0.2.4                       2001:db8:aaaa:204:* 01/10/2023 23:29:11
-------------------------------------------------------------------------------
Number of entries : 1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

The DF election is key for the forwarding and backup functions in single-active multihoming ESs. The PE elected as DF is the primary for the ES in the Epipe and unblocks its SAP and spoke SDP for upstream and downstream traffic. The rest of the PEs in the ES bring their ES SAPs or spoke SDPs operationally down.

PE-5 is a non-DF, as follows:

```
[/]
A:admin@PE-5# show service system bgp-evpn ethernet-segment name "ESI-45" evi evi-1 20

===============================================================================
EVI DF and Candidate List
===============================================================================
EVI           SvcId         Actv Timer Rem      DF  DF Last Change
-------------------------------------------------------------------------------
20            2             0                   no  01/10/2023 23:28:52
===============================================================================

===============================================================================
DF Candidates                           Time Added           Oper Pref  Do Not
                                                               Value    Preempt
-------------------------------------------------------------------------------
192.0.2.4                               01/10/2023 23:29:08  0          Disabl*
192.0.2.5                               01/10/2023 23:29:18  0          Disabl*
-------------------------------------------------------------------------------
Number of entries: 2
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

In single-active multihoming, the service SAP or spoke SDP is brought operationally down on the non-DF, as shown in the following output:

```
[/]
A:admin@PE-5# show service id 2 sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
56:201           Spok     192.0.2.6       Up      Down      524275    524275
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
```

The spoke sdp 56:201 is operationally down with a StandbyForMHProtocol flag:

```
[/]
A:admin@PE-5# show service id 2 sdp 56:201 detail | match Flag
Flags              : StandbyForMHProtocol
```

Two consecutive DF elections take place: the first DF election includes all PEs in the ES for that Epipe and determines which PE is the primary PE (flags P = 1, B = 0). The second DF election excludes this DF and determines which PE is the backup (P = 0, B = 1). All other PEs signal flags P = 0 and B = 0.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

When the primary PE fails, AD per-ES and AD per-EVI withdrawal messages are sent to the remote PE, which updates its next hop to the backup. The backup PE takes over immediately without waiting for the ES activation timer (configured with the **es-activation-timer** command) to bring up its SAP and spoke SDP.

ES failures

When the SDP toward the primary (DF) fails, the backup PE needs to take over. An SDP failure is emulated and log 99 on PE-4 shows that SDP 46 is operationally down and PE-4 is no longer the DF:

```
191 2023/01/10 23:38:55.867 CET MINOR: SVCMGR #2303 Base
"Status of SDP 46 changed to admin=up oper=down"

193 2023/01/10 23:38:55.868 CET MINOR: SVCMGR #2094 Base
"Ethernet Segment:ESI-45, EVI:20, Designated Forwarding state changed to:false"
```

Remote PEs receive route withdrawal updates (unreachable NLRI) from the former DF PE-4, for example on PE-2:

```
# on PE-2:
1 2023/01/10 23:38:55.869 CET MINOR: DEBUG #2001 Base Peer 1: 2001:db8::2:4
"Peer 1: 2001:db8::2:4: UPDATE
Peer 1: 2001:db8::2:4 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 86
    Flag: 0x90 Type: 15 Len: 82 Multiprotocol Unreachable NLRI:
        Address Family EVPN
        Type: EVPN-AD Len: 25 RD: 192.0.2.4:20 ESI: 01:00:00:00:00:45:00:00:00:01, tag: MAX-ET
 Label: 0 (Raw Label: 0x0) PathId:
        Type: EVPN-AD Len: 25 RD: 192.0.2.4:20 ESI: 01:00:00:00:00:45:00:00:00:01, tag: 456
 Label: 0 (Raw Label: 0x0) PathId:
        Type: EVPN-ETH-SEG Len: 23 RD: 192.0.2.4:0 ESI: 01:00:00:00:00:45:00:00:00:01, IP-Len:
 4 Orig-IP-Addr: 192.0.2.4
"
```

The backup PE-5 is promoted to primary (P = 1, B = 0) and sends BGP updates accordingly. The following AD per-EVI is received on PE-2:

```
# on PE-2:
4 2023/01/10 23:38:55.873 CET MINOR: DEBUG #2001 Base Peer 1: 2001:db8::2:5
"Peer 1: 2001:db8::2:5: UPDATE
Peer 1: 2001:db8::2:5 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 113
    Flag: 0x90 Type: 14 Len: 36 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.5
        Type: EVPN-AD Len: 25 RD: 192.0.2.5:20 ESI: 01:00:00:00:00:45:00:00:00:01, tag: 456
 Label: 8388448 (Raw Label: 0x7fff60) PathId:
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:20
        l2-attribute:MTU: 1514 C: 0 P: 1 B: 0
    Flag: 0xc0 Type: 40 Len: 37 Prefix-SID-attr:
       SRv6 Services TLV (37 bytes):-
           Type: SRV6 L2 Service TLV (6)
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
           Length: 34 bytes, Reserved: 0x0
             SRv6 SID Sub-Sub-TLV
                Type: 1 Len: 6
                BL:48 NL:16 FL:20 AL:0 TL:20 TO:64
"
```

PE-5 brings up its spoke SDP without waiting for the ES activation timer and takes over immediately. It is now the only DF candidate, and therefore the DF, as follows:

```
[/]
A:admin@PE-5# show service system bgp-evpn ethernet-segment name "ESI-45" evi evi-1 20

===============================================================================
EVI DF and Candidate List
===============================================================================
EVI           SvcId         Actv Timer Rem      DF  DF Last Change
-------------------------------------------------------------------------------
20            2             0                   yes 01/10/2023 23:28:52
===============================================================================

===============================================================================
DF Candidates                           Time Added           Oper Pref  Do Not
                                                               Value    Preempt
-------------------------------------------------------------------------------
192.0.2.5                               01/10/2023 23:29:18  0          Disabl*
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

BGP updates are exchanged and the remote PEs resolve the ESI to the TEP address 192.0.2.5. For example, on PE-2:

```
[/]
A:admin@PE-2# show service id 2 segment-routing-v6 esi 01:00:00:00:00:45:00:00:00:01

===============================================================================
Segment Routing v6 Ethernet Segment Dest
===============================================================================
Instance  Eth SegId                       Num. Macs     Last Change
-------------------------------------------------------------------------------
1         01:00:00:00:00:45:00:00:00:01   0             01/10/2023 23:38:56
-------------------------------------------------------------------------------
Number of entries: 1
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
Segment Routing v6 Dest TEP Info
===============================================================================
Instance  TEP Address                   Segment Id          Last Change
-------------------------------------------------------------------------------
1         192.0.2.5                     2001:db8:aaaa:205:* 01/10/2023 23:38:56
-------------------------------------------------------------------------------
Number of entries : 1
-------------------------------------------------------------------------------
===============================================================================
* indicates that the corresponding row element may have been truncated.
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Because of the default DF election algorithm, this process is revertive; as soon as the SDP 46 is operationally up again, a new DF election is triggered with two DF candidates and PE-4 is elected as DF. A non-revertive mode is also available if preference-based DF election is configured.

Troubleshooting and debugging

The following **show** and **debug** commands can be used in EVPN-VPWS:

- **show redundancy bgp-evpn-multi-homing**
- **show router bgp routes evpn** (and filters)
- **show service segment-routing-v6 [**<*ip-address*>**]**
- **show service id** <*service-id*> **bgp-evpn**
- **show service system bgp-evpn**
- **show service system bgp-evpn ethernet-segment** (and modifiers)
- **debug router bgp update**
- **show log log-id** <*log-id*>

Most of these commands have been shown in the preceding sections; some commands are shown in this section.

Information about the configured boot timers (before DF election) and ES activation timer (after the system has been elected DF) is shown as follows:

```
[/]
A:admin@PE-2# show redundancy bgp-evpn-multi-homing

===============================================================================
Redundancy BGP EVPN Multi-homing Information
===============================================================================
Boot-Timer              : 10 secs
Boot-Timer Remaining    : 0 secs
ES Activation Timer     : 3 secs
===============================================================================
```

See chapter EVPN for MPLS Tunnels for a description of these timers.

The following command shows that the BGP route type 4 (ES route) messages are only imported by the PEs in the same ES; for example, on PE-3:

```
[/]
A:admin@PE-3# show router bgp routes evpn eth-seg
===============================================================================
 BGP Router ID:192.0.2.3        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Eth-Seg Routes
===============================================================================
Flag  Route Dist.         ESI                           NextHop
      OrigAddr
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
u*>i  192.0.2.2:0         01:00:00:00:00:23:00:00:00:01 192.0.2.2
      192.0.2.2

-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

On PE-4:

```
[/]
A:admin@PE-4# show router bgp routes evpn eth-seg
===============================================================================
 BGP Router ID:192.0.2.4        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Eth-Seg Routes
===============================================================================
Flag  Route Dist.         ESI                           NextHop
      OrigAddr
-------------------------------------------------------------------------------
u*>i  192.0.2.5:0         01:00:00:00:00:45:00:00:00:01 192.0.2.5
      192.0.2.5

-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

The following command shows all the EVPN-SRv6 destinations toward TEP 192.0.2.4. Epipe 1 has an EVPN-SRv6 destination toward TEP 192.0.2.4 directly and Epipe 2 has an EVPN-SRv6 destination to ESI-45, which is resolved to TEP 192.0.2.4. This is shown in the following output:

```
[/]
A:admin@PE-2# show service segment-routing-v6 192.0.2.4

===============================================================================
SRV6 Tunnel Endpoint: 192.0.2.4
===============================================================================
Service Id               Segment Id     Type                     Srv6 Instance
-------------------------------------------------------------------------------
1                        2001:db8:aaaa* evpn                     1
-------------------------------------------------------------------------------
===============================================================================
* indicates that the corresponding row element may have been truncated.

===============================================================================
BGP EVPN SRV6 Ethernet Segment Dest
===============================================================================
Instance  Service Id     Eth Seg Id                    Segment Id
-------------------------------------------------------------------------------
1         2              01:00:00:00:00:45:00:00:00:01 2001:db8:aaaa:204:7fff:*
-------------------------------------------------------------------------------
===============================================================================
* indicates that the corresponding row element may have been truncated.
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The following command lists all configured ESs on the system:

```
[/]
A:admin@PE-2# show service system bgp-evpn ethernet-segment

===============================================================================
Service Ethernet Segment
===============================================================================
Name                            ESI                           Admin     Oper
-------------------------------------------------------------------------------
ESI-23                          01:00:00:00:00:23:00:00:00:01 Enabled   Up
-------------------------------------------------------------------------------
Entries found: 1
===============================================================================
```

In addition to the preceding commands, the following **tools dump** commands may be useful:

- **tools dump service evpn usage** - This command shows the number of EVPN-SRv6 (and EVPN-MPLS and EVPN-VXLAN) destinations in the system.
- **tools dump service system bgp-evpn ethernet-segment** <*name*> **evi** <*evi*> **df** - This command computes the DF election for a specific ESI and EVI. For all-active multihoming, there is no DF election and all PEs forward traffic. For single-active multihoming, one PE is active for a service while another PE is a backup. This command shows the DF (primary), even if it is not the local PE.

The usage of EVPN resources is shown as follows:

```
[/]
A:admin@PE-2# tools dump service evpn usage

vxlan-srv6-evpn-mpls usage statistics at 01/10/2023 23:35:05:

MPLS-TEP                                        :             0
VXLAN-TEP                                       :             0
SRV6-TEP                                        :             2
Total-TEP                                       :      2/ 16383

Mpls Dests (TEP, Egress Label + ES + ES-BMAC)   :             0
Mpls Etree Leaf Dests                           :             0
Vxlan Dests (TEP, Egress VNI + ES)              :             0
Srv6 Dests (TEP, SID  + ES)                     :             2
Total-Dest                                      :      2/196607

Sdp Bind +  Evpn Dests                          :      2/245759
ES L2/L3 PBR                                    :      0/ 32767
Evpn Etree Remote BUM Leaf Labels               :             0
```

On PE-2, there is one SRv6 TEP (192.0.2.4 in Epipe 1 and in Epipe 2) and there are two SRv6 destinations: 192.0.2.4 and ESI 01:00:00:00:00:45:00:00:00:01. PE-5 is not an SRv6 TEP for PE-2 because it is not a primary and, therefore, is not forwarding any traffic.

In all-active multihoming, the DF election is not applicable:

```
[/]
A:admin@PE-2# tools dump service system bgp-evpn ethernet-segment "ESI-23" evi 20 df

[01/10/2023 23:35:05] All Active VPWS or IP-ALIASING - DF N/A
```

In single-active multihoming, the following command shows which PE is the DF:

```
[/]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
A:admin@PE-5# tools dump service system bgp-evpn ethernet-segment "ESI-45" evi 20 df

[01/10/2023 23:35:10] Computed DF: 192.0.2.4 (Remote) (Boot Timer Expired: Yes)
[01/10/2023 23:35:10] Computed Backup: 192.0.2.5 (This Node)
```

The command is launched on PE-5, which is a backup. The computed DF is PE-4 and the boot timer has expired, meaning there is no DF re-election pending.

Conclusion

EVPN-VPWS is a simplified point-to-point version of RFC 7432. EVPN provides a unified control plane mechanism that simplifies the network deployment and operation. Single-active and all-active multihoming can be used in Epipes; EVPN-VPWS is a differentiator of EVPN compared to traditional TLDP or BGP Epipe redundancy mechanisms.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

EVPN-IFF BGP Attribute Propagation Between Families

This chapter provides information about EVPN-IFF BGP attribute propagation between families .

Topics in this chapter include:

Applicability

The information and MD-CLI configuration in this chapter are based on SR OS Release 22.7.R1. EVPN Interface-ful (EVPN-IFF) BGP attribute propagation between BGP families based on uniform propagation is supported in SR OS Release 21.2.R1 and later.

For more information on routed VPLS in EVPN, see chapters EVPN for VXLAN Tunnels (Layer 3) and EVPN for MPLS Tunnels in Routed VPLS .

Overview

SR OS allows multiple BGP owners in the same VPRN service to receive or advertise IP prefixes contained in the VPRN route table. A VPRN route table can simultaneously install and process IPv4 or IPv6 prefixes for the following owners:

- EVPN Interface-ful (EVPN-IFF)
- EVPN Interface-less (EVPN-IFL)
- VPN-IP (also referred to as IP-VPN routes)
- IP (also referred to as BGP PE-CE routes)

EVPN-IFF routes are EVPN IP-prefix routes, otherwise known as route type 5 (RT-5) routes, that are imported and exported based on the configuration of the R-VPLS services attached to the VPRN. To enable the EVPN-IFF model, the command **configure service vpls <..> bgp-evpn routes ip-prefix advertise true** needs to be configured. By default, BGP attributes are re-originated when a prefix is propagated to and from an EVPN-IFF route. However, BGP attributes can be used to influence routing (for example, local preference, Autonomous System (AS) path, communities, and so on), and therefore, SR OS supports EVPN-IFF BGP attribute propagation to other BGP families (uniform propagation), as described in *draft-ietf-bess-evpn-ipvpn-interworking*.

The following CLI command is used to enable EVPN-IFF BGP attribute propagation and EVPN-IFF best path selection:

```
[ex:/configure service system bgp evpn]
A:admin@PE-4# ip-prefix-routes ?

 ip-prefix-routes

 d-path-length-ignore  - Ignore D-PATH length for BGP path selection of EVPN-IFF
 iff-attribute-        - Enable uniform propagation of BGP attributes
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
  uniform-propagation
iff-bgp-path-         - Enable BGP path selection for EVPN-IFF routes
 selection
```

The **iff-bgp-path-selection** command cannot be enabled when **iff-attribute-uniform-propagation** is disabled.

When **iff-attribute-uniform-propagation** is enabled on a node:

- the following BGP path attributes are propagated:
  - AS path
  - domain path (D-PATH), supported in SR OS Release 21.10.R1 and later
  - IBGP-only attributes, when advertising to an IBGP neighbor: local preference, originator ID, cluster ID
  - Multiple Exit Discriminator (MED)
  - communities, large communities, extended communities
- the following BGP path attributes are not propagated across families:
  - any type 0x06 extended communities supported by RT-5 routes:
    - MAC mobility extended community
    - EVPN router MAC extended community
  - BGP encapsulation extended community
  - Route Target extended community
  - BGP tunnel encapsulation attribute
  - BGP prefix-SID attribute used in RT-5 routes and VPN-IP routes for Segment Routing over IPv6 dataplane (SRv6) services
- IBGP-only attributes are only propagated to IBGP neighbors; EBGP-only attributes only to EBGP neighbors
- routes received with well-known communities, such as no-advertise or no-export(-subconfed), are sent or not sent depending on the community values
- BGP path attributes are propagated even when doing route leaking between routing instances

If multiple EVPN-IFF routes for the same prefix are received for the same VPRN, they are by default ordered and selected based on the lowest R-VPLS Ifindex, Route Distinguisher (RD), and Ethernet tag.

When **iff-bgp-path-selection** is enabled, EVPN-IFF routes with the same or different RD are selected based on regular BGP path selection rules in the following order:

1. valid route wins over invalid route (invalid routes are looped routes or routes where the originator ID matches the receiving router)
2. lowest origin validation state (origin validation state: valid is preferred to origin validation state: not found; origin validation state: not found is preferred to origin validation state: invalid) – applicable to IPv4, IPv6, or BGP Labeled Unicast (BGP-LU) routes
3. lowest Routing Table Manager (RTM) preference
4. highest local preference
5. shortest D-PATH

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

6. lowest Accumulated Interior Gateway Protocol (AIGP) metric (AIGP is not supported for EVPN-IFL, EVPN-IFF, or IP-VPN routes)
7. shortest AS path
8. lowest origin (origin: IGP is preferred to origin: EGP; origin: EGP is preferred to origin: incomplete)
9. lowest MED (routes without MED are considered as zero or infinity based on the configuration of the **always-compare-med** command)
10. lowest owner type (owner type: BGP-label is preferred to owner type: BGP; owner type: BGP is preferred to owner type: BGP-VPN) with BGP-VPN referring to VPN-IP and EVPN-IFL
11. EBGP wins over IBGP
12. lowest route-table or tunnel-table cost to the next-hop

    Note: The **ignore-nh-metric** command is not supported for EVPN-IFF.

13. lowest next-hop type – a next-hop resolved to a tunnel-table entry is considered as a lower type than a next-hop resolved to a route-table entry
14. lowest router ID – applicable to IBGP peers
15. shortest cluster list length – applicable to IBGP peers
16. lowest IP address – IP address refers to the peer that advertised the route
17. EVPN-IFL wins over IPVPN
18. next-hop check (IPv4 next-hop wins over IPv6, then lowest next-hop wins) - The next-hop check is a tiebreaker if BGP receives the same prefix for VPN-IPv6 and EVPN-IFL. An IPv6 prefix received as VPN-IPv6 has an IPv6 next-hop whereas the same IPv6 prefix received as EVPN-IFL can have an IPv4 next-hop.
19. lowest RD for route-table selection
20. lowest path ID (add-path)

Configuration

Figure 138: Example topology shows the example topology with PE-3 as Data Center Gateway (DCGW) between an EVPN-VXLAN network and an EVPN-MPLS network. Routed VPLS is configured on PE-2, PE-3, and PE-6. Supplementary broadcast domain "SBD-12" is configured in the EVPN-VXLAN network between PE-2 and PE-3; "SBD-13" in the EVPN-MPLS network between PE-3 and PE-6. On PE-2, Ethernet VPN instance "EVI-11" is configured toward CE-1.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 138: Example topology

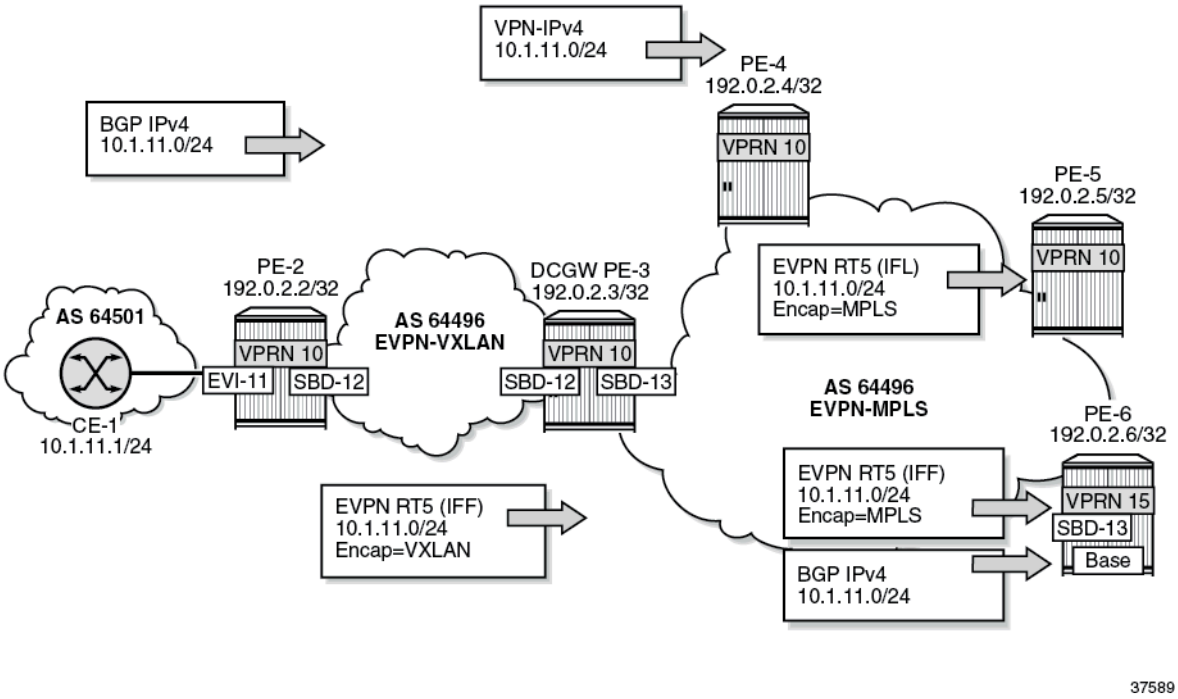

CE-1 advertises prefix 10.1.11.0/24 to BGP neighbor 10.0.0.2 in VPRN 10 on PE-2. PE-2 sends an EVPN-IFF route to DCGW PE-3. PE-3 forwards the prefix 10.1.11.0/24 as VPN-IPv4 route to PE-4, as EVPN-IFL route to PE-5, as EVPN-IFF route to PE-6, and as IPv4 route to PE-6.

The initial configuration includes the following:

- Cards, MDAs, ports
- Router interfaces on all PEs
- IS-IS on the router interfaces
- LDP on the router interfaces on PE-3, PE-4, PE-5, and PE-6

On the PEs, BGP is configured for the EVPN address family. Between PE-3 and PE-4, both the VPN-IPv4 and the EVPN address family are configured. The configuration on PE-3 is as follows:

```
# on PE-3:
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            rapid-update {
                evpn true
            }
            group "internal" {
                peer-as 64496
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            group "internal1" {
                peer-as 64496
                family {
                    evpn true
                }
            }
            neighbor "192.0.2.2" {
                group "internal1"
            }
            neighbor "192.0.2.4" {
                group "internal"
                family {
                    vpn-ipv4 true
                    evpn true
                }
            }
            neighbor "192.0.2.5" {
                group "internal"
                family {
                    evpn true
                }
            }
            neighbor "192.0.2.6" {
                group "internal"
                family {
                    evpn true
                }
            }
        }
```

On CE-1, BGP is configured in VPRN 11 for the IPv4 address family. The export policy adds communities "1:1" and "2:2" and sets the MED to a value of 81.

```
# on CE-1:
configure {
    policy-options {
        community "1:1_2:2" {
            member "1:1" { }
            member "2:2" { }
        }
        policy-statement "export-vnf-to-all" {
            entry 10 {
                from {
                    protocol {
                        name [direct direct-interface]
                    }
                }
                action {
                    action-type accept
                    bgp-med {
                        set 81
                    }
                    community {
                        add ["1:1_2:2"]
                    }
                }
            }
        }
    }
    service {
        vprn "VPRN 11" {
            admin-state enable
            service-id 11
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            customer "1"
            autonomous-system 64501
            bgp {
                split-horizon true
                export {
                    policy ["export-vnf-to-all"]
                }
                group "CE-1-PE-2" {
                    type external
                    peer-as 64496
                }
                neighbor "10.0.0.2" {
                    group "CE-1-PE-2"
                    ebgp-default-reject-policy {
                        import false
                    }
                }
            }
            interface "int-CE-1-PE-2" {
                ipv4 {
                    primary {
                        address 10.0.0.1
                        prefix-length 24
                    }
                }
                sap 1/1/2:11 {
                }
            }
            interface "test" {
                ipv4 {
                    primary {
                        address 10.1.11.1
                        prefix-length 24
                    }
                }
                sap 1/1/2:12 {
                }
            }
```

On PE-2, VPRN 10 has R-VPLS interface "int-EVI-11" toward CE-1 and R-VPLS interface "int-SBD-12" toward PE-3. BGP is configured toward neighbor 10.0.0.1 on CE-1 and the import policy sets the local preference (LP) to 200, as follows:

```
# on PE-2:
configure {
    policy-options {
        policy-statement "local-preference-200" {
            entry 10 {
                action {
                    action-type accept
                    local-preference 200
                }
            }
        }
    }
    service {
        vprn "VPRN 10" {
            admin-state enable
            service-id 10
            customer "1"
            autonomous-system 64496
            bgp {
                split-horizon true
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                local-as {
                    as-number 64496
                }
                import {
                    policy ["local-preference-200"]
                }
                group "PE-2-CE-1" {
                    type external
                    peer-as 64501
                }
                neighbor "10.0.0.1" {
                    group "PE-2-CE-1"
                    ebgp-default-reject-policy {
                        export false
                    }
                }
            }
            interface "int-EVI-11" {
                ipv4 {
                    primary {
                        address 10.0.0.2
                        prefix-length 24
                    }
                    vrrp 1 {
                        backup [10.0.0.2]
                        owner true
                        passive true
                    }
                }
                vpls "EVI-11" {
                }
            }
            interface "int-SBD-12" {
                vpls "SBD-12" {
                    evpn-tunnel {
                    }
                }
            }
        }
        vpls "EVI-11" {
            admin-state enable
            service-id 11
            customer "1"
            routed-vpls {
            }
            sap 1/1/1:11 {
            }
        }
        vpls "SBD-12" {
            admin-state enable
            service-id 12
            customer "1"
            vxlan {
                instance 1 {
                    vni 12
                }
            }
            routed-vpls {
            }
            bgp-evpn {
                evi 12
                routes {
                    mac-ip {
                        advertise false
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    }
                    ip-prefix {
                        advertise true      # enable EVPN-IFF
                    }
                }
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
        }
```

On PE-3, VPRN 10 is configured with:

- three interfaces:
  - R-VPLS interface "int-SBD-12" toward PE-2
  - R-VPLS interface "int-SBD-13" toward PE-6
  - interface "int-VPRN10-PE-3-to-PE-6" to the base router of PE-6.
- BGP-IPVPN for the exchange of VPN-IPv4 routes with PE-4
- BGP-EVPN to propagate EVPN-IFL routes to PE-5 and EVPN-IFF routes to PE-6
- BGP to propagate BGP IPv4 routes to the base router on PE-6. The export policy is only required in the BGP configuration.

```
# on PE-3:
configure {
    policy-options {
        prefix-list "10.1.0.0" {
            prefix 10.1.0.0/16 type longer {
            }
        }
        policy-statement "export-bgp" {
            entry 10 {
                from {
                    prefix-list ["10.1.0.0"]
                }
                action {
                    action-type accept
                }
            }
        }
    }
    service {
        vpls "SBD-12" {
            admin-state enable
            description "EVPN-VXLAN VPLS for EVPN tunnel to PE-2"
            service-id 12
            customer "1"
            vxlan {
                instance 1 {
                    vni 12
                }
            }
            routed-vpls {
            }
            bgp-evpn {
                evi 12
                routes {
                    mac-ip {
                        advertise false
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
                ip-prefix {
                    advertise true      # enable EVPN-IFF
                }
            }
            vxlan 1 {
                admin-state enable
                vxlan-instance 1
            }
        }
    }
    vpls "SBD-13" {
        admin-state enable
        description "EVPN-MPLS VPLS for EVPN tunnel to PE-6"
        service-id 13
        customer "1"
        routed-vpls {
        }
        bgp 1 {
        }
        bgp-evpn {
            evi 13
            routes {
                mac-ip {
                    advertise false
                }
                ip-prefix {
                    advertise true      # enable EVPN-IFF
                }
            }
            mpls 1 {
                admin-state enable
                auto-bind-tunnel {
                    resolution any
                }
            }
        }
    }
    vprn "VPRN 10" {
        admin-state enable
        service-id 10
        customer "1"
        autonomous-system 64496
        bgp-evpn {
            mpls 1 {
                admin-state enable
                route-distinguisher "192.0.2.3:10"
                vrf-target {
                    community "target:64496:10"
                }
                auto-bind-tunnel {
                    resolution any
                }
            }
        }
        bgp-ipvpn {
            mpls {
                admin-state enable
                route-distinguisher "192.0.2.3:10"
                vrf-target {
                    community "target:64496:10"
                }
                auto-bind-tunnel {
                    resolution any
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    }
                }
            }
            bgp {
                rapid-withdrawal true
                export {
                    policy ["export-bgp"]
                }
                group "base router - PE-6" {
                    family {
                        ipv4 true
                    }
                }
                neighbor "10.15.16.6" {
                    group "base router - PE-6"
                    type internal
                    peer-as 64496
                }
            }
            interface "int-SBD-12" {
                vpls "SBD-12" {
                    evpn-tunnel {
                    }
                }
            }
            interface "int-SBD-13" {
                vpls "SBD-13" {
                    evpn-tunnel {
                    }
                }
            }
            interface "int-VPRN10-PE-3-to-PE-6" {
                ipv4 {
                    primary {
                        address 10.15.16.3
                        prefix-length 24
                    }
                }
                sap 1/1/3:13 {
                }
            }
        }
```

On PE-4, VPRN 10 is configured with BGP-IPVPN, as follows. BGP between PE-3 and PE-4 is configured for the VPN-IPv4 address family.

```
# on PE-4:
configure {
    service {
        vprn "VPRN 10" {
            admin-state enable
            service-id 10
            customer "1"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "192.0.2.4:10"
                    vrf-target {
                        community "target:64496:10"
                    }
                    auto-bind-tunnel {
                        resolution any
                    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
            }
```

On PE-5, VPRN 10 is configured with BGP-EVPN, as follows:

```
# on PE-5:
configure {
    service {
        vprn "VPRN 10" {
            admin-state enable
            service-id 10
            customer "1"
            bgp-evpn {
                mpls 1 {
                    admin-state enable
                    route-distinguisher "192.0.2.5:10"
                    vrf-target {
                        community "target:64496:10"
                    }
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            bgp {
            }
```

In the base router of PE-6, BGP is configured to neighbor 10.15.16.3 on PE-3. VPRN 15 is configured with R-VPLS interface "int-SBD-13" toward PE-3. The configuration is as follows:

```
# on PE-6:
configure {
    router "Base" {
        interface "int-PE-6-to-VPRN10-PE-3" {
            port 1/1/1:13
            ipv4 {
                primary {
                    address 10.15.16.6
                    prefix-length 24
                }
            }
        }
        bgp {
            group "PE-6-CE" {
                family {
                    ipv4 true
                }
            }
            neighbor "10.15.16.3" {
                group "PE-6-CE"
                type internal
                peer-as 64496
                local-as {
                    as-number 64496
                }
            }
        }
    }
    service {
        vpls "SBD-13" {
            admin-state enable
            service-id 13
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            customer "1"
            routed-vpls {
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 13
                routes {
                    mac-ip {
                        advertise false
                    }
                    ip-prefix {
                        advertise true
                    }
                }
                mpls 1 {
                    admin-state enable
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
        }
        vprn "VPRN 15" {
            admin-state enable
            service-id 15
            customer "1"
            autonomous-system 64502
            interface "int-SBD-13" {
                vpls "SBD-13" {
                    evpn-tunnel {
                    }
                }
            }
        }
```

Default behavior

By default, BGP path attributes are re-originated when a prefix is propagated to and from an EVPN-IFF route. Figure 139: EVPN-IFF BGP path attributes are re-originated by PE-2 and PE-3 shows that PE-2 receives an IPv4 route for prefix 10.1.11.0/24 with non-default BGP path attributes, whereas PE-2 propagates the prefix as an EVPN-IFF route with default path attributes.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 139: EVPN-IFF BGP path attributes are re-originated by PE-2 and PE-3

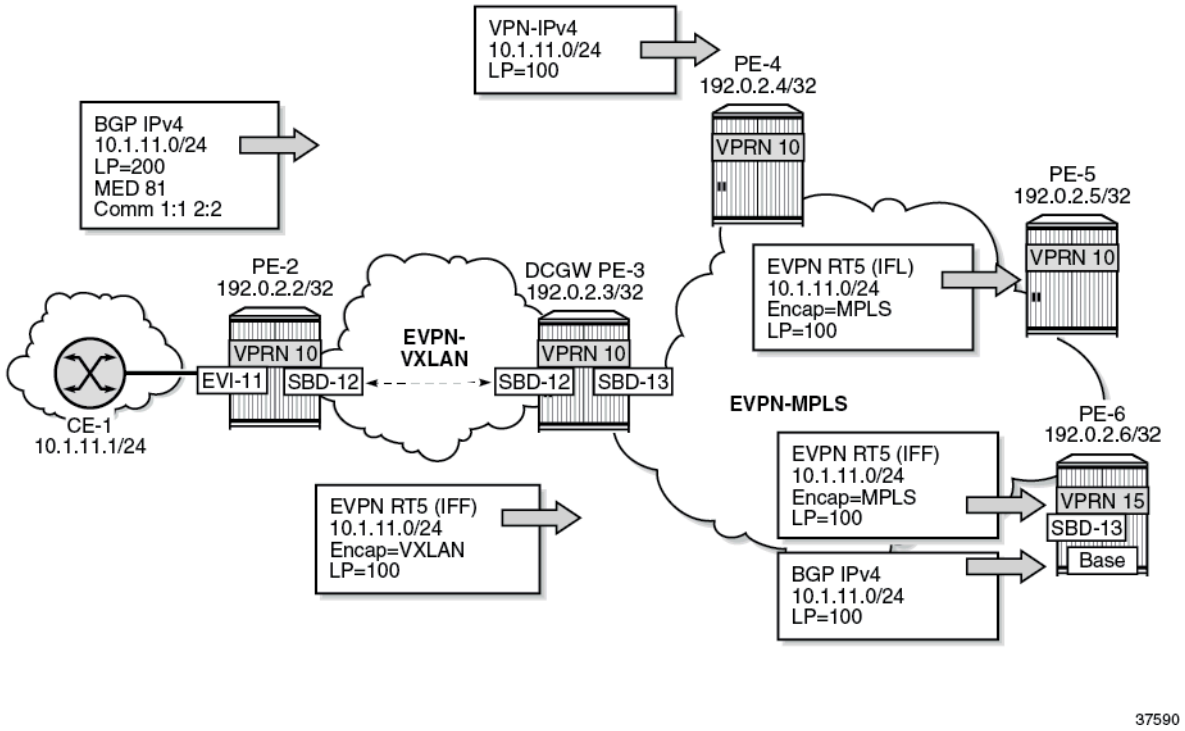

VPRN 10 on PE-2 received a BGP IPv4 route for prefix 10.1.11.0/24 with LP 200, MED 81, and communities "1:1" and "2:2":

```
[/]
A:admin@PE-2# show router 10 bgp routes 10.1.11.0/24 hunt
===============================================================================
 BGP Router ID:192.0.2.2        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP IPv4 Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : 10.1.11.0/24
Nexthop        : 10.0.0.1
Path Id        : None
From           : 10.0.0.1
Res. Protocol  : LOCAL                 Res. Metric    : 0
Res. Nexthop   : 10.0.0.1
Local Pref.    : 200                   Interface Name : int-EVI-11
Aggregator AS  : None                  Aggregator     : None
Atomic Aggr.   : Not Atomic            MED            : 81
AIGP Metric    : None                  IGP Cost       : 0
Connector      : None
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Community      : 1:1 2:2
Cluster        : No Cluster Members
Originator Id  : None                  Peer Router Id : 255.0.0.0
Fwd Class      : None                  Priority       : None
Flags          : Used Valid Best IGP In-RTM
Route Source   : External
AS-Path        : 64501
Route Tag      : 0
Neighbor-AS    : 64501
Orig Validation: NotFound
Source Class   : 0                     Dest Class     : 0
Add Paths Send : Default
RIB Priority   : Normal
Last Modified  : 00h08m14s

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

PE-2 propagates prefix 10.1.11.0/24 as an EVPN-IFF route to PE-3 with default BGP attributes: LP 100, no MED, and without the communities "1:1" and "2:2":

```
[/]
A:admin@PE-2# show router bgp routes evpn ip-prefix prefix 10.1.11.0/24 hunt
===============================================================================
 BGP Router ID:192.0.2.2        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN IP-Prefix Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
Network        : n/a
Nexthop        : 192.0.2.2
Path Id        : None
To             : 192.0.2.3
Res. Nexthop   : n/a
Local Pref.    : 100                   Interface Name : NotAvailable
Aggregator AS  : None                  Aggregator     : None
Atomic Aggr.   : Not Atomic            MED            : None
AIGP Metric    : None                  IGP Cost       : n/a
Connector      : None
Community      : target:64496:12 mac-nh:02:13:ff:ff:ff:49
                 bgp-tunnel-encap:VXLAN
Cluster        : No Cluster Members
Originator Id  : None                  Peer Router Id : 192.0.2.3
Origin         : IGP
AS-Path        : No As-Path
EVPN type      : IP-PREFIX
ESI            : ESI-0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Tag            : 0
Gateway Address: 02:13:ff:ff:ff:49
Prefix         : 10.1.11.0/24
Route Dist.    : 192.0.2.2:12
MPLS Label     : VNI 12
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                     Dest Class     : 0

-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

Uniform propagation for EVPN-IFF BGP path attributes to different BGP families

Enable **iff-attribute-uniform-propagation** and **iff-best-path-selection** on PE-2 as follows:

```
# on PE-2:
configure {
    service {
        system {
            bgp {
                evpn {
                    ip-prefix-routes
                        iff-attribute-uniform-propagation
                        iff-bgp-path-selection
                    }
                }
            }
```

In a similar configuration, **iff-attribute-uniform-propagation** and **iff-bgp-path-selection** are enabled on the other PEs.

The following command shows that uniform propagation for EVPN-IFF BGP path attributes and BGP path selection are enabled:

```
[/]
A:admin@PE-2# show service system bgp-evpn

===============================================================================
System BGP EVPN Information
===============================================================================
Eth Seg Route Dist.               : <none>
Eth Seg Oper Route Dist.          : <none>
Eth Seg Oper Route Dist Type      : none
Ad Per ES Route Target            : evi-rt
Etree
  Leaf                            : Disabled
Mcast Leave Sync Prop             : 5
Attribute Uniform Prop            : Enabled
BGP Path Selection                : Enabled
D-Path Length Ignore              : Disabled
===============================================================================
```

Figure 140: Uniform propagation for EVPN-IFF BGP path attributes between families shows the uniform propagation for EVPN-IFF BGP path attributes between families in the same Virtual Routing and Forwarding (VRF).

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 140: Uniform propagation for EVPN-IFF BGP path attributes between families

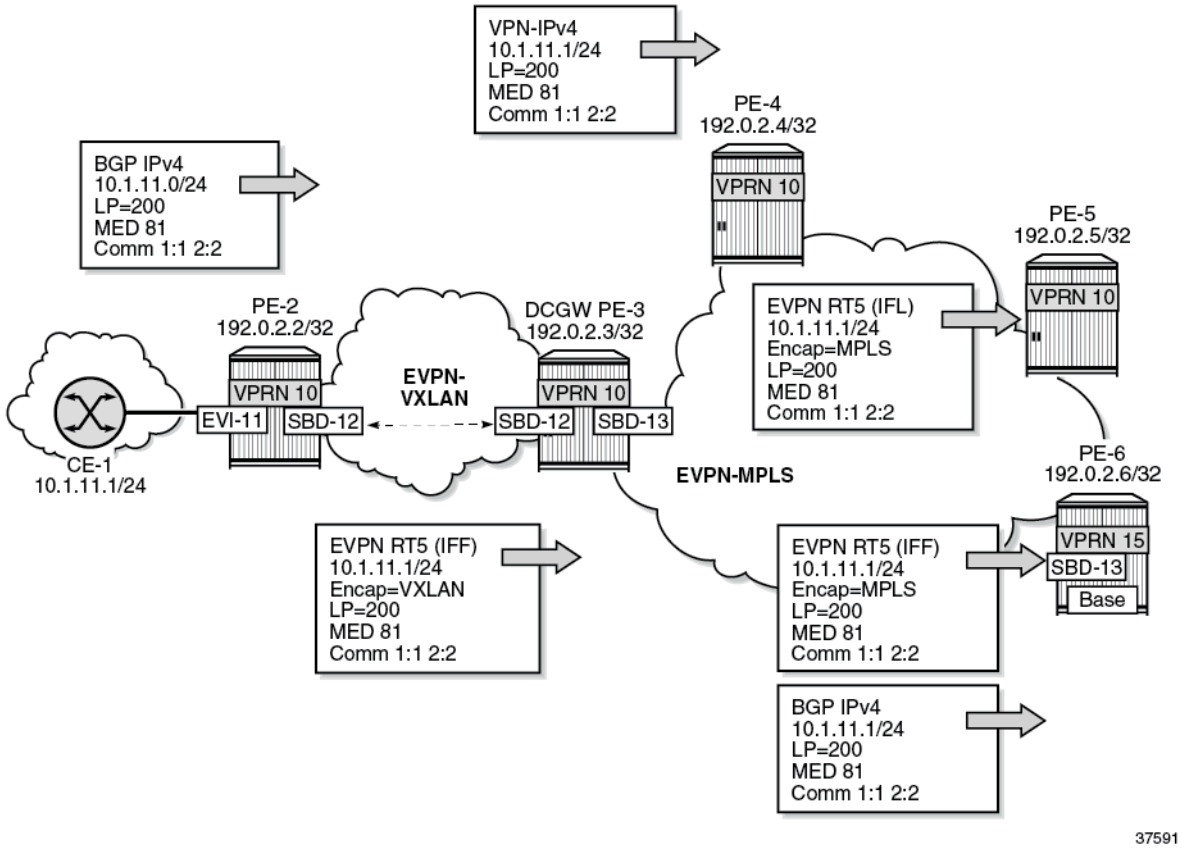

With the uniform propagation for EVPN-IFF BGP path attributes enabled, PE-2 propagates EVPN-IFF route 10.1.11.0/24 to PE-3 with LP 200, MED 81, and communities "1:1" and "2:2". The following EVPN-IFF route is received at PE-3:

```
[/]
A:admin@PE-3# show router bgp routes evpn ip-prefix prefix 10.1.11.0/24 hunt
===============================================================================
 BGP Router ID:192.0.2.3        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN IP-Prefix Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : n/a
Nexthop        : 192.0.2.2
Path Id        : None
From           : 192.0.2.2
Res. Nexthop   : 192.168.23.1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Local Pref.     : 200                    Interface Name : int-PE-3-PE-2
Aggregator AS   : None                   Aggregator     : None
Atomic Aggr.    : Not Atomic             MED            : 81
AIGP Metric     : None                   IGP Cost       : 10
Connector       : None
Community       : 1:1 2:2 target:64496:12 mac-nh:02:13:ff:ff:ff:49
                  bgp-tunnel-encap:VXLAN
Cluster         : No Cluster Members
Originator Id   : None                   Peer Router Id : 192.0.2.2
Flags           : Used Valid Best IGP
Route Source    : Internal
AS-Path         : 64501
EVPN type       : IP-PREFIX
ESI             : ESI-0
Tag             : 0
Gateway Address: 02:13:ff:ff:ff:49
Prefix          : 10.1.11.0/24
Route Dist.     : 192.0.2.2:12
MPLS Label      : VNI 12
Route Tag       : 0
Neighbor-AS     : 64501
Orig Validation: N/A
Source Class    : 0                      Dest Class     : 0
Add Paths Send : Default
Last Modified   : 00h05m09s

-------------------------------------------------------------------------------
---snip---
```

With the uniform propagation for EVPN-IFF BGP path attributes enabled, PE-3 propagates VPN-IPv4 route 10.1.11.0/24 to PE-4 with LP 200, MED 81, and communities "1:1" and "2:2". The following VPN-IPv4 route is received at PE-4:

```
[/]
A:admin@PE-4# show router bgp routes 10.1.11.0/24 vpn-ipv4 hunt
===============================================================================
 BGP Router ID:192.0.2.4        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv4 Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network         : 10.1.11.0/24
Nexthop         : 192.0.2.3
Route Dist.     : 192.0.2.3:10           VPN Label      : 524281
Path Id         : None
From            : 192.0.2.3
Res. Nexthop    : n/a
Local Pref.     : 200                    Interface Name : int-PE-4-PE-3
Aggregator AS   : None                   Aggregator     : None
Atomic Aggr.    : Not Atomic             MED            : 81
AIGP Metric     : None                   IGP Cost       : 10
Connector       : None
Community       : 1:1 2:2 target:64496:10
Cluster         : No Cluster Members
Originator Id   : None                   Peer Router Id : 192.0.2.3
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Fwd Class      : None                    Priority       : None
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : 64501
Route Tag      : 0
Neighbor-AS    : 64501
Orig Validation: N/A
Source Class   : 0                       Dest Class     : 0
Add Paths Send : Default
Last Modified  : 00h06m17s
VPRN Imported  :  10

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

PE-3 propagates EVPN-IFL route 10.1.11.0/24 to PE-5 with LP 200, MED 81, and communities "1:1" and "2:2". The following EVPN-IFL route is received at PE-5:

```
[/]
A:admin@PE-5# show router bgp routes evpn ip-prefix prefix 10.1.11.0/24 hunt
===============================================================================
 BGP Router ID:192.0.2.5        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN IP-Prefix Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : n/a
Nexthop        : 192.0.2.3
Path Id        : None
From           : 192.0.2.3
Res. Nexthop   : 192.168.35.1
Local Pref.    : 200                     Interface Name : int-PE-5-PE-3
Aggregator AS  : None                    Aggregator     : None
Atomic Aggr.   : Not Atomic              MED            : 81
AIGP Metric    : None                    IGP Cost       : 10
Connector      : None
Community      : 1:1 2:2 target:64496:10 bgp-tunnel-encap:MPLS
Cluster        : No Cluster Members
Originator Id  : None                    Peer Router Id : 192.0.2.3
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : 64501
EVPN type      : IP-PREFIX
ESI            : ESI-0
Tag            : 0
Gateway Address: 00:00:00:00:00:00
Prefix         : 10.1.11.0/24
Route Dist.    : 192.0.2.3:10
MPLS Label     : LABEL 524280
Route Tag      : 0
Neighbor-AS    : 64501
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Orig Validation: N/A
Source Class   : 0                     Dest Class     : 0
Add Paths Send : Default
Last Modified  : 00h06m52s

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

PE-3 propagates EVPN-IFF route 10.1.11.0/24 to PE-6 with LP 200, MED 81, and communities "1:1" and "2:2". The following EVPN-IFF route is received at PE-6:

```
[/]
A:admin@PE-6# show router bgp routes evpn ip-prefix prefix 10.1.11.0/24 hunt
===============================================================================
 BGP Router ID:192.0.2.6        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN IP-Prefix Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : n/a
Nexthop        : 192.0.2.3
Path Id        : None
From           : 192.0.2.3
Res. Nexthop   : 192.168.36.1
Local Pref.    : 200                   Interface Name : int-PE-6-PE-3
Aggregator AS  : None                  Aggregator     : None
Atomic Aggr.   : Not Atomic            MED            : 81
AIGP Metric    : None                  IGP Cost       : 10
Connector      : None
Community      : 1:1 2:2 target:64496:13 mac-nh:02:17:ff:ff:ff:4a
                 bgp-tunnel-encap:MPLS
Cluster        : No Cluster Members
Originator Id  : None                  Peer Router Id : 192.0.2.3
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : 64501
EVPN type      : IP-PREFIX
ESI            : ESI-0
Tag            : 0
Gateway Address: 02:17:ff:ff:ff:4a
Prefix         : 10.1.11.0/24
Route Dist.    : 192.0.2.3:13
MPLS Label     : LABEL 524283
Route Tag      : 0
Neighbor-AS    : 64501
Orig Validation: N/A
Source Class   : 0                     Dest Class     : 0
Add Paths Send : Default
Last Modified  : 00h07m20s

-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
RIB Out Entries
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

PE-3 propagates BGP IPv4 route 10.1.11.0/24 to PE-6 with LP 200, MED 81, and communities "1:1" and "2:2". The following IPv4 route is received at PE-6:

```
[/]
A:admin@PE-6# show router bgp routes 10.1.11.0/24 hunt
===============================================================================
 BGP Router ID:192.0.2.6        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP IPv4 Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : 10.1.11.0/24
Nexthop        : 10.15.16.3
Path Id        : None
From           : 10.15.16.3
Res. Protocol  : LOCAL                 Res. Metric    : 0
Res. Nexthop   : 10.15.16.3
Local Pref.    : 200                   Interface Name : int-PE-6-to-VPRN10-PE*
Aggregator AS  : None                  Aggregator     : None
Atomic Aggr.   : Not Atomic            MED            : 81
AIGP Metric    : None                  IGP Cost       : 0
Connector      : None
Community      : 1:1 2:2
Cluster        : No Cluster Members
Originator Id  : None                  Peer Router Id : 192.0.2.3
Fwd Class      : None                  Priority       : None
Flags          : Used Valid Best IGP In-RTM
Route Source   : Internal
AS-Path        : 64501
Route Tag      : 0
Neighbor-AS    : 64501
Orig Validation: NotFound
Source Class   : 0                     Dest Class     : 0
Add Paths Send : Default
RIB Priority   : Normal
Last Modified  : 00h07m37s

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
* indicates that the corresponding row element may have been truncated.
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

EVPN-IFF BGP path attributes exported to leaked EVPN routes

Figure 141: Example topology shows the example topology with two VPRNs on DCGW PE-3 where routes are leaked.

Figure 141: Example topology

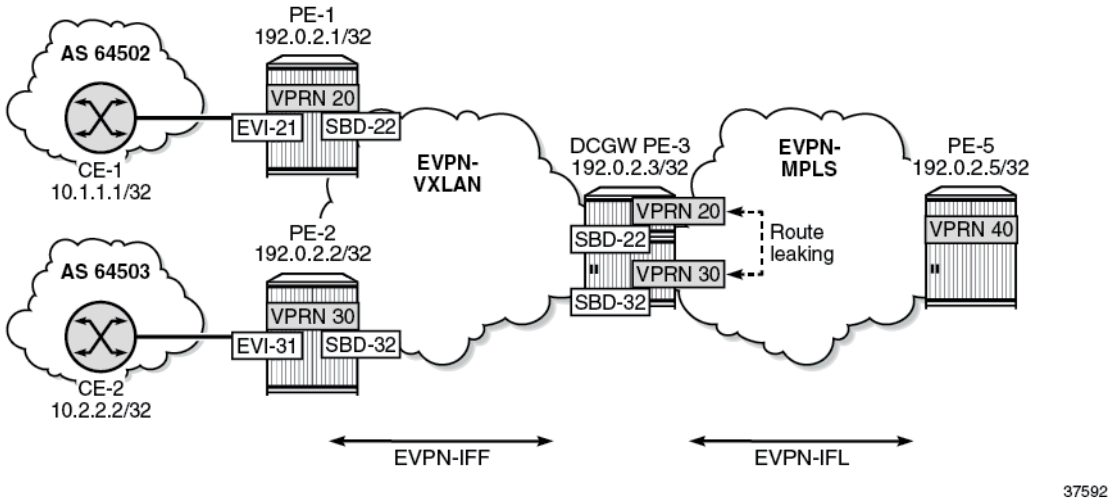

The uniform propagation for EVPN-IFF BGP path attributes is enabled on all PEs.

Figure 142: BGP path attributes are propagated in leaked EVPN routes shows that CE-1 exports an IPv4 route for prefix 10.1.1.1/32 to PE-1. This route has non-default BGP attributes; for example, MED 88, AS path 64502 64505 64504, and community "11:11" "color:00:55". PE-1 exports this route as an EVPN-IFF route to PE-3. PE-3 forwards this route as EVPN-IFL route to PE-5. On PE-3, the route is leaked from VPRN 20 to VPRN 30. The BGP path attributes are propagated to the leaked EVPN routes, except those attributes that are not expected to be propagated, such as the router's MAC extended community. PE-3 advertises an EVPN-IFF route for prefix 10.1.1.1/32 to PE-2.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 142: BGP path attributes are propagated in leaked EVPN routes

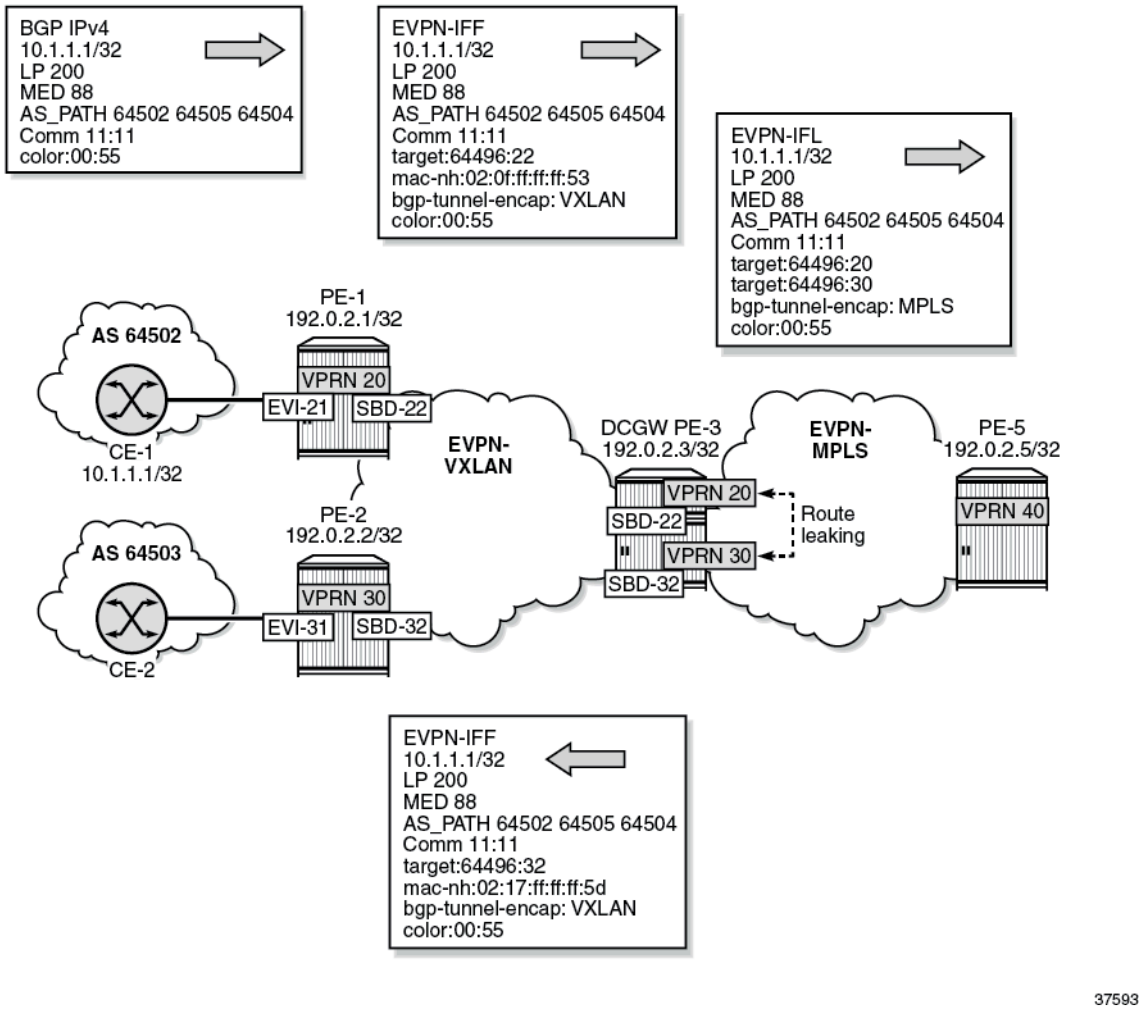

In a similar way, CE-2 exports IPv4 prefix 10.2.2.2/32 to PE-2 with non-default BGP path attributes. PE-2 advertises this prefix as an EVPN-IFF route with the same BGP path attributes. PE-3 leaks the route from VPRN 30 to VPRN 20 while preserving the BGP path attributes. PE-3 advertises an EVPN-IFF route for prefix 10.2.2.2/32 to PE-1 with the same BGP path attributes. PE-3 also advertises the prefix as EVPN-IFL route to PE-5 with the same BGP path attributes. For brevity, the routes for prefix 10.2.2.2/32 are not shown here.

In this example, VPRN "CE-1" is configured as follows. The export policy sets the MED, prepends some AS numbers to the AS path, and adds the communities "11:11" and "color:00:55".

```
# on CE-1:
configure {
    policy-options {
        community "11:11" {
            member "11:11" { }
        }
        community "color:00:55" {
            member "color:00:55" { }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
        policy-statement "export-vnf-to-all-2" {
            entry 10 {
                from {
                    protocol {
                        name [direct direct-interface]
                    }
                }
                action {
                    action-type next-entry
                    as-path-prepend {
                        as-path 64504
                    }
                    bgp-med {
                        set 88
                    }
                    community {
                        add ["11:11" "color:00:55"]
                    }
                }
            }
            entry 20 {
                from {
                    protocol {
                        name [direct direct-interface]
                    }
                }
                action {
                    action-type accept
                    as-path-prepend {
                        as-path 64505
                    }
                }
            }
        }
    }
    service {
        vprn "CE-1" {
            admin-state enable
            service-id 23
            customer "1"
            autonomous-system 64502
            bgp {
                local-as {
                    as-number 64502
                }
                export {
                    policy ["export-vnf-to-all-2"]
                }
                group "PE-1-CE-1" {
                }
                neighbor "10.2.0.254" {
                    group "PE-1-CE-1"
                    type external
                    peer-as 64496
                    ebgp-default-reject-policy {
                        import false
                    }
                }
            }
            interface "int-CE-1-PE-1" {
                ipv4 {
                    primary {
                        address 10.2.0.1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        prefix-length 24
                    }
                }
                sap 1/2/2:21 {
                }
            }
            interface "loopback" {
                loopback true
                ipv4 {
                    primary {
                        address 10.1.1.1
                        prefix-length 32
                    }
                }
            }
        }
```

On PE-1, an import policy sets the LP to a value of 200. VPRN 20 has R-VPLS interface "int-EVI-21" toward CE-1 and R-VPLS interface "int-SBD-22" toward PE-2.

```
# on PE-1:
configure {
    policy-options {
        policy-statement "local-preference-200" {
            entry 10 {
                action {
                    action-type accept
                    local-preference 200
                }
            }
        }
    }
    service {
        vpls "EVI-21" {
            admin-state enable
            service-id 21
            customer "1"
            routed-vpls {
            }
            sap 1/2/1:21 {
            }
        }
        vpls "SBD-22" {
            admin-state enable
            service-id 22
            customer "1"
            vxlan {
                instance 1 {
                    vni 22
                }
            }
            routed-vpls {
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 22
                routes {
                    mac-ip {
                        advertise false
                    }
                    ip-prefix {
                        advertise true
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
            }
            vxlan 1 {
                admin-state enable
                vxlan-instance 1
            }
        }
    }
    vprn "VPRN 20" {
        admin-state enable
        service-id 20
        customer "1"
        autonomous-system 64496
        bgp {
            local-as {
                as-number 64496
            }
            import {
                policy ["local-preference-200"]
            }
            group "PE-1-CE" {
                type external
                peer-as 64502
            }
            neighbor "10.2.0.1" {
                group "PE-1-CE"
                ebgp-default-reject-policy {
                    export false
                }
            }
        }
        interface "int-EVI-21" {
            ipv4 {
                primary {
                    address 10.2.0.254
                    prefix-length 24
                }
                vrrp 1 {
                    backup [10.2.0.254]
                    owner true
                    passive true
                }
            }
            vpls "EVI-21" {
            }
        }
        interface "int-SBD-22" {
            vpls "SBD-22" {
                evpn-tunnel {
                }
            }
        }
    }
```

The configuration on PE-2 is similar with VPRN 30, R-VPLS "EVI-31", and R-VPLS "SBD-32".

PE-3 has two VPRNs: "VPRN 20" and "VPRN 30". Export policy "leak-color-55-into-30" is used to leak routes with color community "color:00:55" from VPRN 20 to VPRN 30. The configuration is as follows:

```
# on PE-3:
configure {
    policy-options {
        community "RT64496:20" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            member "target:64496:20" { }
        }
        community "RT64496:30" {
            member "target:64496:30" { }
        }
        community "color:00:55" {
            member "color:00:55" { }
        }
        policy-statement "leak-color-55-into-20" {
            entry 10 {
                from {
                    community {
                        name "color:00:55"
                    }
                }
                action {
                    action-type accept
                    community {
                        add ["RT64496:20" "RT64496:30"]
                    }
                }
            }
        }
        policy-statement "leak-color-55-into-30" {
            entry 10 {
                from {
                    community {
                        name "color:00:55"
                    }
                }
                action {
                    action-type accept
                    community {
                        add ["RT64496:20" "RT64496:30"]
                    }
                }
            }
        }
    }
    service {
        vpls "SBD-22" {
            admin-state enable
            service-id 22
            customer "1"
            vxlan {
                instance 1 {
                    vni 22
                }
            }
            routed-vpls {
            }
            bgp-evpn {
                evi 22
                routes {
                    mac-ip {
                        advertise false
                    }
                    ip-prefix {
                        advertise true
                    }
                }
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
            }
        }
        vprn "VPRN 20" {
            admin-state enable
            service-id 20
            customer "1"
            autonomous-system 64496
            bgp-evpn {
                mpls 1 {
                    admin-state enable
                    route-distinguisher "192.0.2.3:20"
                    vrf-export {
                        policy ["leak-color-55-into-30"]
                    }
                    vrf-target {
                        import-community "target:64496:20"
                    }
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            interface "int-SBD-22" {
                vpls "SBD-22" {
                    evpn-tunnel {
                    }
                }
            }
        }
        vpls "SBD-32" {
            admin-state enable
            service-id 32
            customer "1"
            vxlan {
                instance 1 {
                    vni 32
                }
            }
            routed-vpls {
            }
            bgp-evpn {
                evi 32
                routes {
                    mac-ip {
                        advertise false
                    }
                    ip-prefix {
                        advertise true
                    }
                }
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
        }
        vprn "VPRN 30" {
            admin-state enable
            service-id 30
            customer "1"
            autonomous-system 64496
            bgp-evpn {
                mpls 1 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    admin-state enable
                    route-distinguisher "192.0.2.3:30"
                    vrf-export {
                        policy ["leak-color-55-into-20"]
                    }
                    vrf-target {
                        import-community "target:64496:30"
                    }
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            interface "int-SBD-32" {
                vpls "SBD-32" {
                    evpn-tunnel {
                    }
                }
            }
        }
```

PE-3 exports the prefix route as EVPN-IFL to PE-5. On PE-5, VPRN 40 is configured as follows:

```
# on PE-5:
configure {
    policy-options {
        community "RT64496:20" {
            member "target:64496:20" { }
        }
        community "RT64496:30" {
            member "target:64496:30" { }
        }
        policy-statement "vrf-40-export" {
            entry 10 {
                from {
                    protocol {
                        name [direct direct-interface]
                    }
                }
                action {
                    action-type accept
                    community {
                        add ["RT64496:20" "RT64496:30"]
                    }
                }
            }
        }
        policy-statement "vrf-40-import" {
            entry 10 {
                from {
                    community {
                        name "RT64496:20"
                    }
                }
                action {
                    action-type accept
                }
            }
            entry 20 {
                from {
                    community {
                        name "RT64496:30"
                    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
                action {
                    action-type accept
                }
            }
        }
    }
    service {
        vprn "VPRN 40" {
            admin-state enable
            service-id 40
            customer "1"
            autonomous-system 64496
            bgp-evpn {
                mpls 1 {
                    admin-state enable
                    route-distinguisher "192.0.2.5:40"
                    vrf-export {
                        policy ["vrf-40-export"]
                    }
                    vrf-import {
                        policy ["vrf-40-import"]
                    }
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            interface "loopback" {
                loopback true
                ipv4 {
                    primary {
                        address 10.5.5.5
                        prefix-length 32
                    }
                }
            }
```

CE-1 exports an IPv4 route for prefix 10.1.1.1/32 to PE-1 with community "color:00:55" and other non-default BGP path attributes. The route table for VPRN 20 on PE-1 includes an BGP IPv4 route for prefix 10.1.1.1/32:

```
[/]
A:admin@PE-1# show router 20 route-table 10.1.1.1/32

===============================================================================
Route Table (Service: 20)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
10.1.1.1/32                                   Remote  BGP       00h03m25s  170
       10.2.0.1                                                      0
-------------------------------------------------------------------------------
No. of Routes: 1
```

PE-1 propagates prefix 10.1.1.1/32 in an EVPN-IFF route. On PE-3, the route table includes an EVPN-IFF route for prefix 10.1.1.1/32:

```
[/]
A:admin@PE-3# show router 20 route-table 10.1.1.1/32
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Route Table (Service: 20)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
10.1.1.1/32                                     Remote  EVPN-IFF  00h03m28s  169
       int-SBD-22 (ET-02:0f:ff:ff:ff:53)                            0
-------------------------------------------------------------------------------
No. of Routes: 1
```

PE-3 forwards prefix 10.1.1.1/32 as an EVPN-IFL to PE-5. On PE-5, the route table includes an EVPN-IFL route for prefix 10.1.1.1/32:

```
[/]
A:admin@PE-5# show router 40 route-table

===============================================================================
Route Table (Service: 40)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
10.1.1.1/32                                     Remote  EVPN-IFL  00h03m59s  170
       192.0.2.3 (tunneled)                                         10
10.2.2.2/32                                     Remote  EVPN-IFL  00h03m56s  170
       192.0.2.3 (tunneled)                                         10
10.5.5.5/32                                     Local   Local     00h04m03s  0
       loopback                                                     0
-------------------------------------------------------------------------------
No. of Routes: 3
```

In a similar way, PE-5 received an EVPN-IFL route for prefix 10.2.2.2/32. Prefix 10.5.5.5/32 is local to VPRN 40 on PE-5 and is advertised to PE-3 as EVPN-IFL route.

On PE-3, routes with community "color:00:55" are leaked between VPRN 20 and VPRN 30. PE-1 and PE-3 have forwarded the route with the original BGP path attributes, so this community is preserved and the route for prefix 10.1.1.1/32 is leaked to VPRN 30, as shown in the following route table. The next hop is R-VPLS "SBD-22" in local VPRN 20.

```
[/]
A:admin@PE-3# show router 30 route-table

===============================================================================
Route Table (Service: 30)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
10.1.1.1/32                                     Remote  EVPN-IFL  00h04m23s  169
       Local VRF [20:int-SBD-22]                                    0
10.2.2.2/32                                     Remote  EVPN-IFF  00h04m15s  169
       int-SBD-32 (ET-02:13:ff:ff:ff:5d)                            0
10.3.0.0/24                                     Remote  EVPN-IFF  00h04m34s  169
       int-SBD-32 (ET-02:13:ff:ff:ff:5d)                            0
10.5.5.5/32                                     Remote  EVPN-IFL  00h04m22s  170
       192.0.2.5 (tunneled)                                         10
-------------------------------------------------------------------------------
No. of Routes: 4
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

PE-3 propagates prefix 10.1.1.1/32 as an EVPN-IFF route to PE-2, so the route table for VPRN 30 on PE-2 includes an entry for prefix 10.1.1.1/32 with next hop "SBD-32" toward VPRN 30 on PE-3:

```
[/]
A:admin@PE-2# show router 30 route-table

===============================================================================
Route Table (Service: 30)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
10.1.1.1/32                                   Remote  EVPN-IFF  00h06m05s  169
       int-SBD-32 (ET-02:17:ff:ff:ff:5d)                          0
10.2.2.2/32                                   Remote  BGP       00h05m57s  170
       10.3.0.1                                                   0
10.3.0.0/24                                   Local   Local     00h06m57s  0
       int-EVI-31                                                 0
10.5.5.5/32                                   Remote  EVPN-IFF  00h06m04s  169
       int-SBD-32 (ET-02:17:ff:ff:ff:5d)                          0
-------------------------------------------------------------------------------
No. of Routes: 4
```

The following show commands illustrate that the BGP path attributes are propagated. VPRN 20 on PE-1 receives an IPv4 route for prefix 10.1.1.1/32 from CE-1 with LP 200, MED 88, AS path 64502 64505 64504, and communities "1:1" "color:00:55", as follows:

```
[/]
A:admin@PE-1# show router 20 bgp routes 10.1.1.1/32 hunt
===============================================================================
 BGP Router ID:192.0.2.1        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP IPv4 Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : 10.1.1.1/32
Nexthop        : 10.2.0.1
Path Id        : None
From           : 10.2.0.1
Res. Protocol  : LOCAL                 Res. Metric    : 0
Res. Nexthop   : 10.2.0.1
Local Pref.    : 200                   Interface Name : int-EVI-21
Aggregator AS  : None                  Aggregator     : None
Atomic Aggr.   : Not Atomic            MED            : 88
AIGP Metric    : None                  IGP Cost       : 0
Connector      : None
Community      : 11:11 color:00:55
Cluster        : No Cluster Members
Originator Id  : None                  Peer Router Id : 192.0.2.1
Fwd Class      : None                  Priority       : None
Flags          : Used Valid Best IGP In-RTM
Route Source   : External
AS-Path        : 64502 64505 64504
Route Tag      : 0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Neighbor-AS    : 64502
Orig Validation: NotFound
Source Class   : 0                     Dest Class     : 0
Add Paths Send : Default
RIB Priority   : Normal
Last Modified  : 00h06m22s

-------------------------------------------------------------------------------
---snip---
```

PE-1 forwards an EVPN-IFF route to PE-3 for prefix 10.1.1.1/32 with the original BGP path attributes, as follows:

```
[/]
A:admin@PE-1# show router bgp routes 10.1.1.1/32 evpn ip-prefix hunt
===============================================================================
 BGP Router ID:192.0.2.1        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN IP-Prefix Routes
===============================================================================
---snip---
-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
---snip---

Network        : n/a
Nexthop        : 192.0.2.1
Path Id        : None
To             : 192.0.2.3
Res. Nexthop   : n/a
Local Pref.    : 200                   Interface Name : NotAvailable
Aggregator AS  : None                  Aggregator     : None
Atomic Aggr.   : Not Atomic            MED            : 88
AIGP Metric    : None                  IGP Cost       : n/a
Connector      : None
Community      : 11:11 target:64496:22 mac-nh:02:0f:ff:ff:ff:53
                 bgp-tunnel-encap:VXLAN color:00:55
Cluster        : No Cluster Members
Originator Id  : None                  Peer Router Id : 192.0.2.3
Origin         : IGP
AS-Path        : 64502 64505 64504
EVPN type      : IP-PREFIX
ESI            : ESI-0
Tag            : 0
Gateway Address: 02:0f:ff:ff:ff:53
Prefix         : 10.1.1.1/32
Route Dist.    : 192.0.2.1:22
MPLS Label     : VNI 22
Route Tag      : 0
Neighbor-AS    : 64502
Orig Validation: N/A
Source Class   : 0                     Dest Class     : 0
---snip---
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

PE-3 forwards an EVPN-IFL route for prefix 10.1.1.1/32 to PE-5, so PE-5 receives the following route with the original BGP path attributes:

```
[/]
A:admin@PE-5# show router bgp routes evpn ip-prefix prefix 10.1.1.1/32 hunt
===============================================================================
 BGP Router ID:192.0.2.5        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN IP-Prefix Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : n/a
Nexthop        : 192.0.2.3
Path Id        : None
From           : 192.0.2.3
Res. Nexthop   : 192.168.35.1
Local Pref.    : 200                    Interface Name : int-PE-5-PE-3
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : 88
AIGP Metric    : None                   IGP Cost       : 10
Connector      : None
Community      : 11:11 target:64496:20 target:64496:30
                 bgp-tunnel-encap:MPLS color:00:55
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.3
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : 64502 64505 64504
EVPN type      : IP-PREFIX
ESI            : ESI-0
Tag            : 0
Gateway Address: 00:00:00:00:00:00
Prefix         : 10.1.1.1/32
Route Dist.    : 192.0.2.3:20
MPLS Label     : LABEL 524282
Route Tag      : 0
Neighbor-AS    : 64502
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
Add Paths Send : Default
Last Modified  : 00h10m47s

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

On PE-3, the route for prefix 10.1.1.1/32 is leaked from VPRN 20 to VPRN 30. Prefix 10.1.1.1/32 is then advertised to PE-2 in the new context but preserves the BGP path attributes, so PE-2 receives the following route:

```
[/]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
A:admin@PE-2# show router bgp routes evpn ip-prefix prefix 10.1.1.1/32 hunt
===============================================================================
 BGP Router ID:192.0.2.2        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN IP-Prefix Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
---snip---
Network        : n/a
Nexthop        : 192.0.2.3
Path Id        : None
From           : 192.0.2.3
Res. Nexthop   : 192.168.23.2
Local Pref.    : 200                    Interface Name : int-PE-2-PE-3
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : 88
AIGP Metric    : None                   IGP Cost       : 10
Connector      : None
Community      : 11:11 target:64496:32 mac-nh:02:17:ff:ff:ff:5d
                 bgp-tunnel-encap:VXLAN color:00:55
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.3
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : 64502 64505 64504
EVPN type      : IP-PREFIX
ESI            : ESI-0
Tag            : 0
Gateway Address: 02:17:ff:ff:ff:5d
Prefix         : 10.1.1.1/32
Route Dist.    : 192.0.2.3:32
MPLS Label     : VNI 32
Route Tag      : 0
Neighbor-AS    : 64502
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
Add Paths Send : Default
Last Modified  : 00h08m17s
---snip---
```

Conclusion

SR OS nodes can be configured to propagate EVPN-IFF BGP path attributes between families to influence the path selection, as per *draft-ietf-bess-evpn-ipvpn-interworking*.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

EVPN-MPLS E-Tree

This chapter provides information about EVPN-MPLS E-Tree.

Topics in this chapter include:

Applicability

This chapter was initially written for SR OS Release 15.0.R6, but the CLI in the current edition is based on SR OS Release 23.7.R1. VPLS E-Tree without EVPN is supported in SR OS Release 12.0.R4, and later. EVPN-MPLS E-Tree is supported in SR OS Release 15.0.R1, and later.

Overview

Ethernet Tree (E-Tree) is a rooted multipoint Ethernet service defined by the Metro Ethernet Forum (MEF). E-Tree can be implemented based on the following:

- RFC 7796, *Ethernet-Tree Support in Virtual Private LAN Services* (VPLS E-Tree without EVPN)
- RFC 8317, *E-Tree Support in EVPN and PBB-EVPN* (EVPN-MPLS E-Tree)

VPLS E-Tree without EVPN

The E-Tree implementation is based on RFC 7796 and is supported for unicast and broadcast, unknown unicast, and multicast (BUM) traffic. Interfaces can be defined as root attachment circuit (AC) or leaf AC, or both, as described in Table 8: Interfaces in E-Tree. A VPLS E-Tree can have multiple root ACs. Access and network interfaces are both supported on SAPs and SDP bindings.

Table 8: Interfaces in E-Tree

Interface	Tag
Access interface (user-to-network interface - UNI)	Root tag
	Leaf tag
Network interface (network-to-network interface - NNI)	Root-leaf tag

On the ingress access interfaces, all frames are tagged and forwarded. On the network interfaces, no traffic is dropped based on the root or leaf tag. On the egress access interfaces, all traffic toward a root AC is forwarded, whereas traffic toward a leaf AC is only forwarded when it originates from a root AC, as summarized in Table 9: E-Tree Forwarding on Access Interfaces. Traffic from leaf AC to leaf AC is blocked.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Table 9: E-Tree Forwarding on Access Interfaces

	To root AC	To leaf AC
From root AC	Allowed	Allowed
From leaf AC	Allowed	Not allowed

Within an E-Tree, the split horizon group capability is inherent for leaf SAPs and leaf SDP bindings and extends to all the remote nodes that are part of the same VPLS E-Tree service.

Ingress Tagging and Egress Filtering

Figure 143: Frame Forwarding in a VPLS E-Tree without EVPN shows how frames are forwarded in an E-Tree. The ingress node PE-2 knows whether the frame comes from a leaf AC or a root AC and adds a tag indicating "from root" or "from leaf". Specific VLAN IDs are used to indicate "from root" or "from leaf". The egress node PE-4 forwards the frame based on the "from root" or "from leaf" tag, as follows:

- A frame with the "from root" tag can be forwarded to any AC, leaf or root.
- A frame with the "from leaf" tag can only be forwarded to a root AC, not to a leaf AC.

Figure 143: Frame Forwarding in a VPLS E-Tree without EVPN

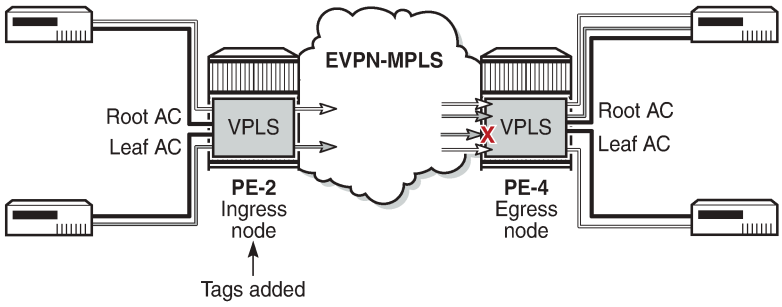

SAPs and SDP bindings are considered as root AC automatically (in the following example, SAP 1/2/c1/1:4 is a root AC); leaf ACs get the keyword **leaf-ac**, and NNI SAPs and SDP bindings get the keyword **root-leaf-tag**. The root tag equals the service delimiting VLAN ID (VID) in the SAP and the leaf tag can only be configured with a different value.

```
On PE-2:
configure {
    service {
        vpls "VPLS 4" {
            admin-state enable
            service-id 4
            customer "1"
            etree true
            sap 1/2/c1/1:4 { }
            sap 1/2/c3/1:4 {
                etree-leaf true
            }
            sap 1/2/c5/1:4 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                etree-root-leaf-tag {
                    leaf 44
                }
            }
            spoke-sdp 24:4 {
                etree-root-leaf-tag true
                vc-type vlan
            }
            spoke-sdp 210:4 {
                etree-leaf true
            }
        }
    }
}
```

VLAN ranges are not allowed in a VPLS E-Tree, as shown for the following connection profile VLAN, which is configured on PE-2:

```
On PE-2:
configure {
    connection-profile {
        vlan 10 {
            qtag-range 10 {
                end 19
            }
            qtag-range 110 {
                end 110
            }
        }
    }
}
```

The following error is raised when attempting to configure a SAP with VLAN range cp-10:

```
configure {
    service {
        vpls "VPLS 4" {
            sap 1/2/c3/1:cp-10 { }
MINOR: MGMT_CORE #4001: configure service vpls "VPLS 4" sap 1/2/c3/1:cp-10
 - vlan-range not allowed with etree vpls

configure {
    service {
        vpls "VPLS 4" {
            sap 1/2/c3/1:cp-10.* {
                etree-leaf true
            }
MINOR: MGMT_CORE #4001: configure service vpls "VPLS 4" sap 1/2/c3/1:cp-10.*
 - SAP and port encapsulation values are incompatible
 - configure port 1/2/c3/1 ethernet encap-type
MINOR: MGMT_CORE #4001: configure service vpls "VPLS 4" sap 1/2/c3/1:cp-10.*
 - vlan-range not allowed with etree vpls
 - configure service vpls "VPLS 4" etree
```

All incoming frames on a SAP or SDP binding in a VPLS have their dot1q/qinq encapsulation removed by the local PE. In a VPLS E-Tree, the local PE then adds a VLAN tag with a dedicated VID indicating whether the frame originates from a root AC or a leaf AC.

- For dot1q/qinq-based L2 services, a VLAN tag with VID x is added for root and VID y for leaf. Frames with VID x are forwarded to any type of AC, while frames with VID y are only forwarded to root ACs at

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

the remote node, as shown in Figure 144: VLAN Tags Added by Ingress Node and Filtered by Egress Node in VPLS E-Tree.

- For pseudowire-based L2 services, a VLAN tag with VID 1 is hard-coded for frames received on a root AC and a VLAN tag with VID 2 for frames received on a leaf AC.

Figure 144: VLAN Tags Added by Ingress Node and Filtered by Egress Node in VPLS E-Tree

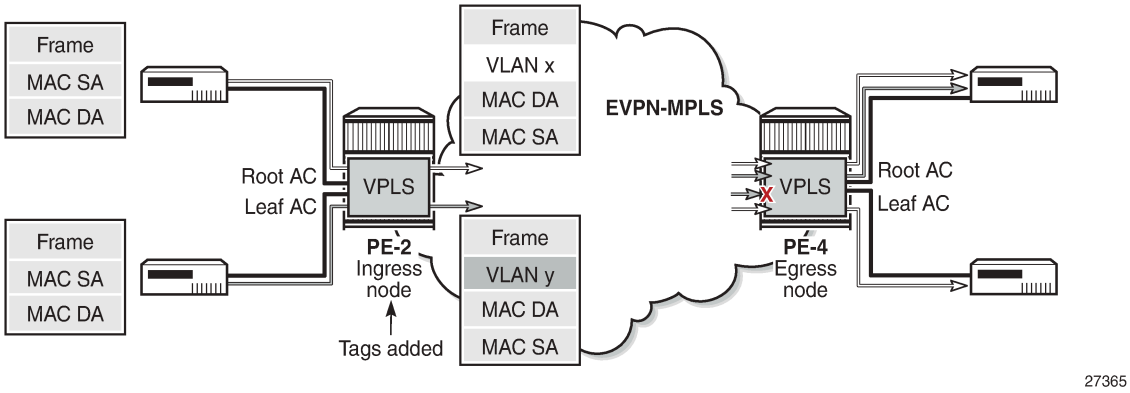

EVPN-MPLS E-Tree

Operators migrate their regular VPLS services to EVPN services because of the advantages offered by EVPN, such as all-active multi-homing, scalability, and easy provisioning. EVPN-MPLS E-Trees block leaf-to-leaf traffic, while allowing all traffic from and to root ACs. The following is a configuration example of an EVPN-MPLS E-Tree. The **evpn-etree-leaf-label** command is only relevant for EVPN E-Tree services and allocates an E-Tree leaf label on the system, which is used for egress filtering of BUM traffic.

```
configure {
    service {
        system {
            bgp {
                evpn {
                    etree-leaf-label true
                }
            }
        }
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            etree true
            bgp 1 { }
            bgp-evpn {
                evi 1
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            sap 1/2/c1/1:1 { }
            sap 1/2/c3/1:1 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                etree-leaf true
            }
            spoke-sdp 210:1 {
                etree-leaf true
            }
        }
    }
}
```

SAPs or SDP bindings are by default root AC objects. MAC addresses learned on root AC objects are advertised as usual, while MAC addresses learned on a SAP or SDP binding configured as leaf AC are advertised with an BGP EVPN E-Tree extended community with leaf indication bit L=1.

BGP EVPN VXLAN is not supported in E-Tree services; only EVPN-MPLS E-Tree is supported. The following error is raised when attempting to configure VXLAN in an E-Tree enabled service:

```
configure {
    service {
        vpls "VPLS 3" {
            admin-state enable
            service-id 3
            customer "1"
            etree true
            vxlan {
                instance 1
MINOR: MGMT_CORE #2203: configure service vpls "VPLS 3" vxlan instance 1
 - Invalid element - etree or m-vpls must be unset
```

In an EVPN-MPLS E-Tree, it is not required and not even possible to configure the **etree-root-leaf-tag** option on interfaces. The following error is raised when attempting to configure a spoke SDP or SAP with **etree-root-leaf-tag** option:

```
configure {
    service {
        vpls "VPLS 1" {
            spoke-sdp 24:1 {
                etree-root-leaf-tag true
                vc-type vlan
            }
MINOR: MGMT_CORE #4001: configure service vpls "VPLS 1" spoke-sdp 24:1 etree-root-leaf-tag
 - Not supported with bgp-evpn
 - configure service vpls "VPLS 1" bgp-evpn


configure {
    service {
        vpls "VPLS 1" {
            sap 1/2/c3/1:200 {
                etree-root-leaf-tag {
                    leaf 22
                }
MINOR: SVCMGR #12: configure service vpls "VPLS 1" sap 1/2/c3/1:200 etree-root-leaf-tag
 - Inconsistent Value error
 - not supported with bgp-evpn
 - configure service vpls "VPLS 1" bgp-evpn
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

BGP EVPN Control Plane for EVPN E-Tree

No leaf tag needs to be added to frames forwarded to EVPN destinations. Instead, the BGP EVPN control plane for EVPN E-Tree advertises a leaf indication bit and a leaf label in the E-tree extended community, as shown in Figure 145: BGP EVPN Control Plane for EVPN E-Tree.

Figure 145: BGP EVPN Control Plane for EVPN E-Tree

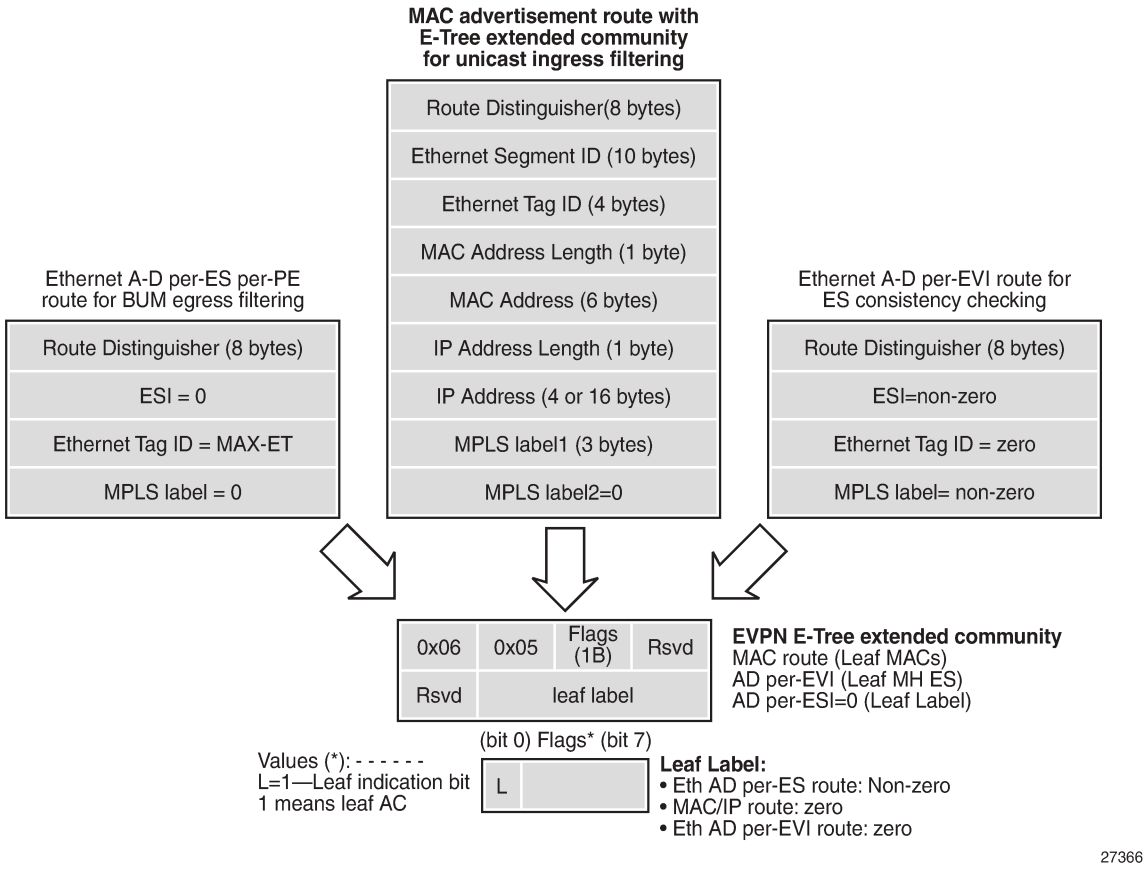

The BGP EVPN control plane is extended with the EVPN E-Tree extended community, as per RFC 8317. The low-order bit of the flags field contains the L-bit (L=1 indicates a leaf AC). The leaf label contains a 20-bit MPLS label that is non-zero for Ethernet Auto Discovery (AD) per Ethernet Segment (per-ES) routes (tag MAX-ET), but it equals zero for MAC/IP routes and Ethernet AD per EVPN Instance (per-EVI) routes (tag 0). The following BGP EVPN AD per-ES route contains an EVPN E-Tree extended community with L=0 and leaf label 524282, and is used for egress BUM filtering. RFC 8317 states that the leaf indication bit L must be ignored on reception and should be zero on transmission.

```
On PE-2:
1 2023/08/01 14:28:10.587 CEST MINOR: DEBUG #2001 Base Peer 1: 192.0.2.4
"Peer 1: 192.0.2.4: UPDATE
Peer 1: 192.0.2.4 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 81
    Flag: 0x90 Type: 14 Len: 36 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        Type: EVPN-AD Len: 25 RD: 192.0.2.2:1 ESI: ESI-0, tag: MAX-ET Label: 0 (Raw Label: 0x0)
 PathId:
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64496:1
        etree::L:0/Leaf-Label:524282
        bgp-tunnel-encap:MPLS
"
```

The following BGP EVPN MAC route contains an EVPN E-Tree extended community with L=1 and leaf label 0, and is used for known unicast ingress filtering:

```
On PE-2:
4 2023/08/01 14:28:14.229 CEST MINOR: DEBUG #2001 Base Peer 1: 192.0.2.4
"Peer 1: 192.0.2.4: UPDATE
Peer 1: 192.0.2.4 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 89
    Flag: 0x90 Type: 14 Len: 44 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
        Type: EVPN-MAC Len: 33 RD: 192.0.2.2:1 ESI: ESI-0, tag: 0, mac len: 48 mac:
 ca:fe:09:29:29:29, IP len: 0, IP: NULL, label1: 8388496 (Raw Label: 0x7fff90)
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64496:1
        etree::L:1/Leaf-Label:0
        bgp-tunnel-encap:MPLS
"
```

The following BGP EVPN AD per-EVI route contains an EVPN E-Tree extended community with L=1 and leaf label 0, and is used for ES consistency checking:

```
On PE-4:
80 2023/08/01 15:12:36.304 CEST MINOR: DEBUG #2001 Base Peer 1: 192.0.2.5
"Peer 1: 192.0.2.5: UPDATE
Peer 1: 192.0.2.5 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 81
    Flag: 0x90 Type: 14 Len: 36 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.5
        Type: EVPN-AD Len: 25 RD: 192.0.2.5:2 ESI: 01:00:00:00:00:45:01:00:00:01, tag: 0 Label:
 8388480 (Raw Label: 0x7fff80) PathId:
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64496:2
        etree::L:1/Leaf-Label:0
        bgp-tunnel-encap:MPLS
"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

When PE-2 receives a BGP EVPN MAC route with an E-Tree extended community with leaf indication bit L=1, the PE imports the route and installs the MAC address in the forwarding database (FDB) with an EVPN leaf (Lf) flag, as follows:

```
[/]
A:admin@PE-2# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
1          ca:fe:01:01:01:01 sdp:210:1               L/0      08/01/23 14:33:40
1          ca:fe:06:46:46:46 mpls-1:                 Evpn     08/01/23 14:31:37
                             192.0.2.4:524281
           ldp:65538
1          ca:fe:07:47:47:47 mpls-1:                 Evpn,Lf  08/01/23 14:31:37
                             192.0.2.4:524281
           ldp:65538
1          ca:fe:08:28:28:28 sap:1/2/c1/1:1          LT/0     08/01/23 14:28:14
1          ca:fe:09:29:29:29 sap:1/2/c3/1:1          LT/0     08/01/23 14:28:14
-------------------------------------------------------------------------------
No. of MAC Entries: 5
-------------------------------------------------------------------------------
Legend:L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf T=Trusted
===============================================================================
```

If receiving the same MAC route as root from PE-1 and as leaf from PE-2, the MAC route from PE-1 is selected: root MAC routes have higher priority than leaf MAC routes. Root static MAC routes take precedence over leaf static MAC routes.

EVPN MAC routes with a higher sequence number have a higher priority than root or leaf MAC routes. MAC mobility procedures take precedence to first identify the location of the MAC before associating that MAC with a root or a leaf site. The EVPN MAC route selection criteria in tie-break order are as follows:

1. Conditional static MACs (local protected MACs)
2. Auto-learned protected MACs (locally learned MACs on SAPs or mesh/spoke SDPs because of the configuration of auto-learn-mac-protect)
3. EVPN ES PBR MACs
4. EVPN static MACs (remote protected MACs)
5. Data plane learned MACs (regular MAC learning on SAPs/SDP-bindings)
6. EVPN MACs with a higher sequence number
7. EVPN E-Tree root MACs
8. Lowest IP (next-hop IP of the EVPN NLRI)
9. Lowest Ethernet tag (Ethernet tag is zero for MPLS and non-zero for VXLAN)
10. Lowest RD

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Ingress Leaf Filtering for Unicast Traffic

EVPN-MPLS E-Tree is the only E-Tree technology able to do unicast ingress filtering, as opposed to the usual unicast egress filtering that, for example, VPLS does. Remote MAC addresses are learned in the control plane, so EVPN can optimize the forwarding by filtering known unicast traffic at the ingress:

- Unicast frames entering a root AC at the ingress PE are not filtered. The MAC destination address (DA) is looked up in the FDB and the frames are forwarded. The MAC source address (SA) is learned and advertised in BGP EVPN without the E-Tree extended community.
- Unicast frames entering a leaf AC at the ingress PE are filtered. The MAC DA is looked up in the FDB. When the MAC DA is learned from an EVPN leaf (or a leaf AC), the frame is dropped. When the MAC DA is learned from an EVPN root (or root AC), the frame is forwarded. The MAC SA is learned and advertised in BGP EVPN with leaf indication bit L=1.

Figure 146: Ingress Leaf Filtering for Known Unicast Traffic shows that PE-4 advertises MAC2 with leaf indication bit L=1. When a frame is sent with MAC SA MAC1 on a leaf AC of PE-2, PE-2 does a MAC lookup in the FDB to find out that the DA MAC2 is learned from an EVPN leaf. Therefore, PE-2 does not forward the frame to PE-4, but drops it at the ingress.

Figure 146: Ingress Leaf Filtering for Known Unicast Traffic

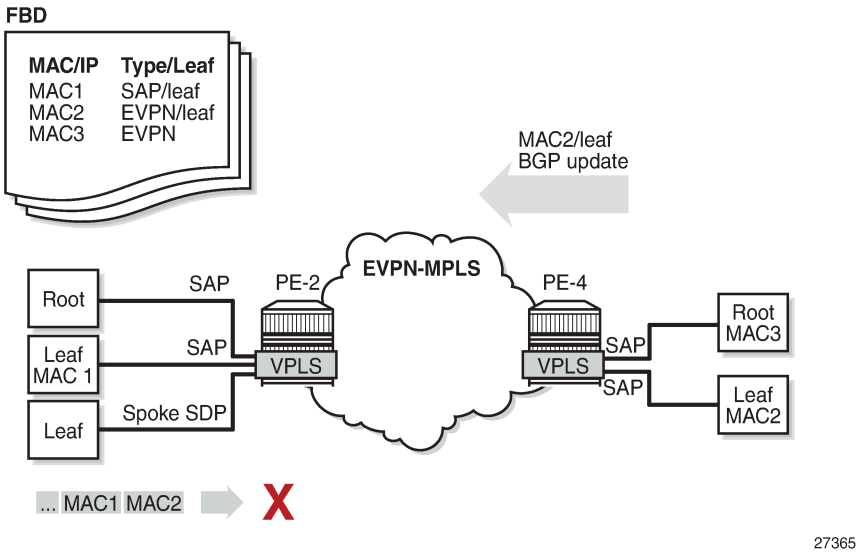

The ingress filtering blocks E-Tree leaf-to-leaf traffic and requires the implementation of an extra leaf EVPN-MPLS destination per remote PE containing leaf ACs per E-Tree service. Therefore, a dedicated EVPN-MPLS binding is created per leaf unicast traffic in the service. This additional internal EVPN-MPLS destination is created per remote PE that contains a leaf and that advertises at least one leaf MAC. The MPLS E-Tree leaf destination is created when a MAC route with L=1 is received. Any EVPN E-Tree service could potentially use one additional EVPN-MPLS destination for leaf unicast traffic per remote PE. This additional EVPN-MPLS leaf destination in the E-Tree is only unicast and not part of the flooding list. The EVPN-MPLS leaf destination consumes EVPN resources, as can be verified as follows:

```
[/]
A:admin@PE-2# tools dump service evpn usage | match "Mpls Etree"
Mpls Etree Leaf Dests                          :             1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

All MAC addresses received with L=1 point to this EVPN-MPLS E-Tree leaf destination, whereas root MAC addresses point to the root destination.

Egress Leaf Filtering for BUM Traffic

Figure 147: Egress Leaf Filtering for BUM Traffic shows that leaf-to-leaf BUM traffic is filtered at the egress, based on the EVPN leaf label advertised in the E-Tree extended community of the zero ESI AD per-ES route (tag=MAX-ET).

- BUM frames that enter a root AC at the ingress PE are not filtered; the BUM frames follow regular EVPN data plane procedures.
- BUM frames that enter a leaf AC at the ingress PE are marked as leaf and forwarded or replicated to the egress IOM. At the egress IOM, the frame is flooded in the default multicast list, subject to the following:
  - Leaf entries are skipped when BUM traffic is forwarded, so no BUM traffic is forwarded to local leaf ACs.
  - BUM traffic to remote BGP EVPN PEs is encapsulated with the EVPN label stack.
    - If the remote PE has advertised an AD per-ES route with E-Tree leaf label L1, this leaf label L1 is added at the bottom of the stack. At the egress PE, when the leaf label L1 matches the leaf label of the PE, the BUM traffic is only forwarded to the root ACs, not to the leaf ACs.
    - If the egress PE does not have any E-Tree enabled service, it has not advertised any AD per-ES route with E-Tree leaf label. The local PE forwards the BUM traffic with BGP EVPN encapsulation, but without an additional label. Even when the egress PE does not have E-Tree enabled, it can still work with the VPLS E-Tree service available in the ingress PE. No traffic is dropped at the egress PE where no E-Tree is configured.

Figure 147: Egress Leaf Filtering for BUM Traffic

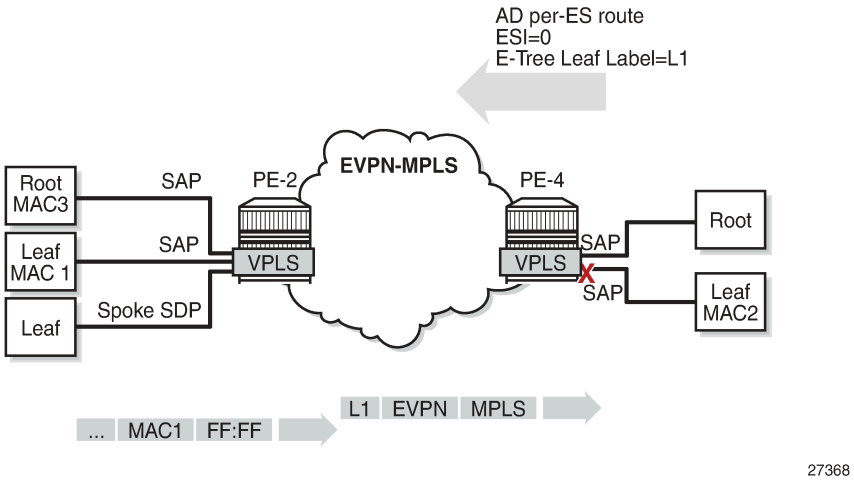

The following command is used to monitor the ESI label entries consumed by the EVPN E-Tree application:

```
[/]
A:admin@PE-2# tools dump service evpn usage | match "BUM"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Evpn Etree Remote BUM Leaf Labels               :            1
```

Configuration

The initial configuration on the nodes includes the following:

- Cards, MDAs, ports
- Router interfaces
- IS-IS (alternatively, OSPF can be used)
- LDP between the PEs
- BGP for the EVPN address family (between the PEs)

In this section, the following cases are described:

- EVPN-MPLS E-Tree without multi-homing
- EVPN-MPLS E-Tree with all-active and single-active multi-homing

EVPN-MPLS E-Tree without Multi-homing

Figure 148: Example Topology for EVPN-MPLS E-Tree without Multi-homing shows an example topology with two PEs in an EVPN-MPLS network with VPLS 1 configured as E-Tree. CE-6 and CE-8 have root ACs and are able to send and receive traffic to and from all other CEs, whereas CE-7, CE-9, and CE-1 are only able to exchange traffic with CE-6 and CE-8, but not with each other. The video source can be connected to CE-8 (root AC) or CE-9 (leaf AC).

Figure 148: Example Topology for EVPN-MPLS E-Tree without Multi-homing

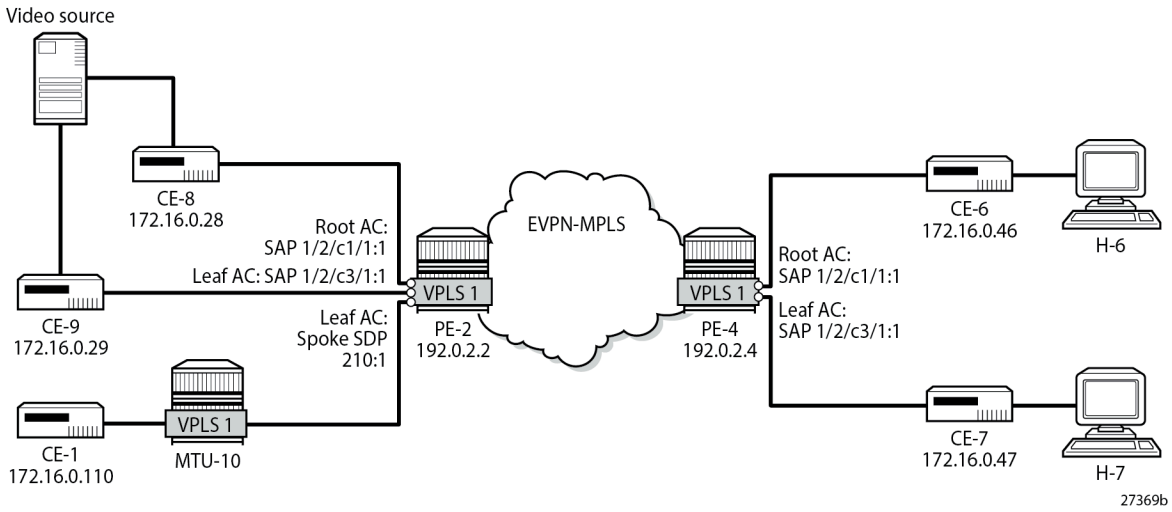

The service configuration on PE-2 is as follows:

```
On PE-2:
configure {
    service {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        sdp 210 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.10
            }
        }
        system {
            bgp {
                evpn {
                    etree-leaf-label true
                }
            }
        }
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            etree true
            bgp 1 { }
            bgp-evpn {
                evi 1
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            sap 1/2/c1/1:1 { }
            sap 1/2/c3/1:1 {
                etree-leaf true
            }
            spoke-sdp 210:1 {
                etree-leaf true
            }
        }
    }
}
```

The service configuration on PE-4 is similar, with SAP 1/2/c1/1:1 as root AC and SAP 1/2/c3/1:1 as leaf AC.

The following command on PE-2 shows that SAP 1/2/c1/1:1 is a root AC (default), SAP 1/2/c3/1:1 is a leaf AC (indicated by "L"), and spoke SDP 210:1 is also a leaf AC.

```
[/]
A:admin@PE-2# show service id 1 etree

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 1                   Vpn Id            : 0
Service Type      : VPLS
---snip---
Etree Mode        : Enabled
Admin State       : Up                  Oper State        : Up
---snip---
-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Identifier                                Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/2/c1/1:1                            q-tag        8936    8936    Up   Up
sap:1/2/c3/1:1 (L)                        q-tag        8936    8936    Up   Up
sdp:210:1 (L) S(192.0.2.10)               Spok         0       8910    Up   Up
-------------------------------------------------------------------------------
Legend: (L): Leaf-Ac, (RL): Root-Leaf-Tag
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

The following command on PE-2 shows that SAP 1/2/c1/1:1 is not configured as a leaf AC (Leaf-Ac Disabled), while SAP 1/2/c3/1:1 is configured as a leaf AC. Root-leaf tag cannot be configured on objects in an EVPN-MPLS E-Tree, so this is always disabled and no leaf tag is defined.

```
[/]
A:admin@PE-2# show service sap-using etree

===============================================================================
Etree SAP Information
===============================================================================
Svc Id      SAP                               Leaf-Tag    Root-     Leaf-Ac
                                                          leaf-tag
-------------------------------------------------------------------------------
1           1/2/c1/1:1                        0           Disabled  Disabled
1           1/2/c3/1:1                        0           Disabled  Enabled
-------------------------------------------------------------------------------
Number of etree saps: 2
===============================================================================
```

Likewise, the following command shows that spoke SDP 210:1 is configured as a leaf AC. Again, root-leaf tag cannot be configured on an object in an EVPN-MPLS E-Tree.

```
[/]
A:admin@PE-2# show service sdp-using etree

===============================================================================
Etree SDP-BIND Information
===============================================================================
Svc Id      SDP-BIND                          Type      Root-     Leaf-Ac
                                                        leaf-tag
-------------------------------------------------------------------------------
1           210:1                             Spoke     Disabled  Enabled
-------------------------------------------------------------------------------
Number of etree sdp-binds: 1
===============================================================================
```

EVPN E-Tree Known Unicast Ingress Filtering

Unicast traffic can be exchanged between CE-8 (root AC) and any other CE. However, unicast traffic from CE-9 on leaf AC can only be exchanged with CE-8 and CE-6 on root ACs, but not with CE-7 (via leaf AC SAP 1/2/c3/1:1) or CE-1 (via leaf AC spoke SDP 210:1), as follows:

```
[/]
A:admin@CE-9# ping 172.16.0.28 interval 0.1 count 20 output-format summary     # succeeds -
 leaf AC can send to root AC
PING 172.16.0.28 56 data bytes
!!!!!!!!!!!!!!!!!!!!
---- 172.16.0.28 PING Statistics ----
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
20 packets transmitted, 20 packets received, 0.00% packet loss
round-trip min = 2.45ms, avg = 2.66ms, max = 2.95ms, stddev = 0.101ms

[/]
A:admin@CE-9# ping 172.16.0.46 interval 0.1 count 20 output-format summary     # succeeds -
 leaf AC can send to root AC
PING 172.16.0.46 56 data bytes
!!!!!!!!!!!!!!!!!!!!
---- 172.16.0.46 PING Statistics ----
20 packets transmitted, 20 packets received, 0.00% packet loss
round-trip min = 3.48ms, avg = 3.76ms, max = 5.80ms, stddev = 0.477ms

[/]
A:admin@CE-9# ping 172.16.0.47 interval 0.1 count 20 output-format summary     # fails - leaf
 AC cannot send to leaf AC!
PING 172.16.0.47 56 data bytes
...   ...   ...   ...   . ....................
---- 172.16.0.47 PING Statistics ----
20 packets transmitted, 0 packets received, 100% packet loss

[/]
A:admin@CE-9# ping 172.16.0.110 interval 0.1 count 20 output-format summary    # fails - leaf
 AC cannot send to leaf AC!
PING 172.16.0.110 56 data bytes
...   ...   ...   ...   . ....................
---- 172.16.0.110 PING Statistics ----
20 packets transmitted, 0 packets received, 100% packet loss
```

The following FDB for VPLS 1 on PE-2 shows that MAC address ca:fe:07:47:47:47 of CE-7 is learned as EVPN leaf, whereas MAC address ca:fe:01:01:01:01 of CE-1 is learned on the local root spoke SDP.

```
[/]
A:admin@PE-2# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
1          ca:fe:01:01:01:01 sdp:210:1               L/0      08/01/23 14:33:40
1          ca:fe:06:46:46:46 mpls-1:                 Evpn     08/01/23 14:31:37
                             192.0.2.4:524281
           ldp:65538
1          ca:fe:07:47:47:47 mpls-1:                 Evpn,Lf  08/01/23 14:31:37
                             192.0.2.4:524281
           ldp:65538
1          ca:fe:08:28:28:28 sap:1/2/c1/1:1          LT/0     08/01/23 14:28:14
1          ca:fe:09:29:29:29 sap:1/2/c3/1:1          LT/0     08/01/23 14:28:14
-------------------------------------------------------------------------------
No. of MAC Entries: 5
-------------------------------------------------------------------------------
Legend:L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf T=Trusted
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

EVPN E-Tree BUM Egress Filtering

When multicast traffic is sent from a video source via CE-8 (root AC), both CE-6 and CE-7 receive this traffic; for multicast traffic sent via CE-9 (leaf AC), only CE-6 (root AC) receives this traffic. PE-2 received leaf label 524282 in an AD per-ES route from PE-4, as follows:

```
[/]
A:admin@PE-2# show router bgp routes evpn auto-disc rd 192.0.2.4:1 detail
===============================================================================
 BGP Router ID:192.0.2.2        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Auto-Disc Routes
===============================================================================
Original Attributes

Network        : n/a
Nexthop        : 192.0.2.4
Path Id        : None
From           : 192.0.2.4
Res. Nexthop   : 192.168.24.2
Local Pref.    : 100                   Interface Name : int-PE-2-PE-4
---snip---
Community      : target:64496:1 etree::L:0/Leaf-Label:524282
                 bgp-tunnel-encap:MPLS
---snip---
EVPN type      : AUTO-DISC
ESI            : ESI-0
Tag            : MAX-ET
Route Dist.    : 192.0.2.4:1
MPLS Label     : LABEL 0
---snip---
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

Multicast traffic is sent with three labels: MPLS (LDP), EVPN, and leaf label. The EVPN label is 524280 for multicast, as follows:

```
[/]
A:admin@PE-2# show service id 1 evpn-mpls

===============================================================================
BGP EVPN-MPLS Dest (Instance 1)
===============================================================================
TEP Address                     Transport:Tnl     Egr Label  Oper  Mcast  Num
                                                             State        MACs
-------------------------------------------------------------------------------
192.0.2.4                       ldp:65538         524280     Up    bum    0
192.0.2.4                       ldp:65538         524281     Up    none   2
-------------------------------------------------------------------------------
Number of entries: 2
-------------------------------------------------------------------------------
---snip---
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The MPLS transport label is 524287, as follows:

```
[/]
A:admin@PE-2# show router ldp bindings active prefixes prefix 192.0.2.4/32

===============================================================================
---snip---
===============================================================================
LDP IPv4 Prefix Bindings (Active)
===============================================================================
Prefix                                  Op
IngLbl                                  EgrLbl
EgrNextHop                              EgrIf/LspId
-------------------------------------------------------------------------------
192.0.2.4/32                            Push
  --                                    524287
192.168.24.2                            1/1/c1/1

192.0.2.4/32                            Swap
524285                                  524287
192.168.24.2                            1/1/c1/1

-------------------------------------------------------------------------------
No. of IPv4 Prefix Active Bindings: 2
===============================================================================
```

The video source sends the following multicast stream via CE-9 (leaf AC):

```
[/]
A:admin@CE-9# show router pim group detail

===============================================================================
PIM Source Group ipv4
===============================================================================
Group Address      : 232.1.1.1
Source Address     : 192.168.55.2
---snip---
Rpf Neighbor       : 192.168.19.1
Incoming Intf      : int-CE-9-CE-1
Outgoing Intf List : int-CE-9-PE-2

Curr Fwding Rate   : 9751.560 kbps
Forwarded Packets  : 29664              Discarded Packets  : 0
Forwarded Octets   : 43962048           RPF Mismatches     : 0
Spt threshold      : 0 kbps             ECMP opt threshold : 7
Admin bandwidth    : 1 kbps
-------------------------------------------------------------------------------
Groups : 1
===============================================================================
```

Receiver H-6 has joined the multicast stream and CE-6 (root AC) receives the following multicast group:

```
[/]
A:admin@CE-6# show router pim group detail

===============================================================================
PIM Source Group ipv4
===============================================================================
Group Address      : 232.1.1.1
Source Address     : 192.168.55.2
---snip---
Rpf Neighbor       : 172.16.0.29
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Incoming Intf      : int-CE-6-PE-4
Outgoing Intf List : int-CE-6-H-6

Curr Fwding Rate   : 9751.560 kbps
Forwarded Packets  : 168421            Discarded Packets  : 0
Forwarded Octets   : 249599922         RPF Mismatches     : 0
Spt threshold      : 0 kbps            ECMP opt threshold : 7
Admin bandwidth    : 1 kbps
-------------------------------------------------------------------------------
Groups : 1
===============================================================================
```

Receiver H-7 has also joined the multicast stream, but CE-7 (leaf AC) cannot receive BUM traffic from a leaf AC, so the forwarding rate is 0 kbps, as follows:

```
[/]
A:admin@CE-7# show router pim group detail

===============================================================================
PIM Source Group ipv4
===============================================================================
Group Address      : 232.1.1.1
Source Address     : 192.168.55.2
---snip---
Rpf Neighbor       :
Incoming Intf      :
Outgoing Intf List : int-CE-7-H-7

Curr Fwding Rate   : 0.000 kbps
Forwarded Packets  : 0                 Discarded Packets  : 0
Forwarded Octets   : 0                 RPF Mismatches     : 0
Spt threshold      : 0 kbps            ECMP opt threshold : 7
Admin bandwidth    : 1 kbps
-------------------------------------------------------------------------------
Groups : 1
===============================================================================
```

EVPN E-Tree Egress Filtering Based on MAC SA

Egress filtering on MAC SA is required to cover cases when the ingress PE sends traffic received on a leaf AC, but without leaf indication. Figure 149: EVPN E-Tree Egress Filtering Based on MAC SA shows that CE-9 sends traffic with MAC SA ca:fe:09:29:29:29 on a leaf AC.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 149: EVPN E-Tree Egress Filtering Based on MAC SA

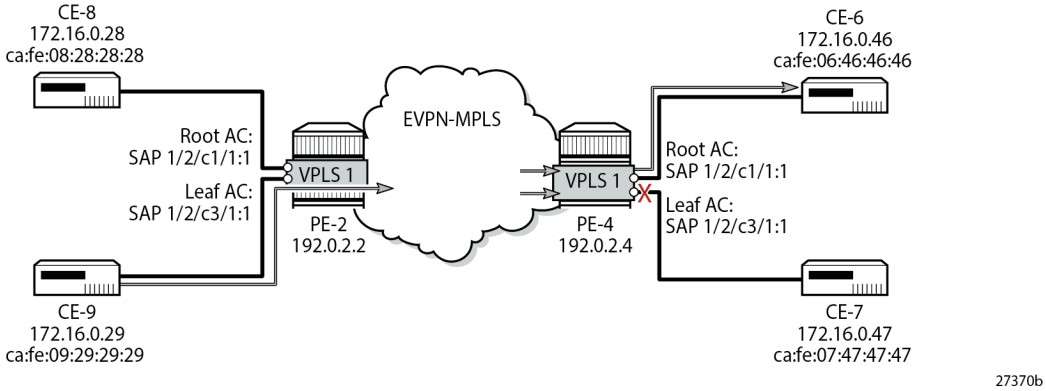

When CE-9 sends unicast traffic to CE-6 with root MAC DA ca:fe:06:46:46:46, the ingress PE-2 forwards the frames to this root MAC DA to egress PE-4. However, if PE-4 does not have the MAC DA in its FDB (because of aging or MAC flush and the MAC route has not made it yet to PE-2), it may flood the frame to all the root and leaf ACs, even if the frame originated from a leaf AC. EVPN E-Tree egress filtering based on MAC SA prevents this from happening, so the traffic is only forwarded to the root AC.

The data path does the egress filtering based on MAC SA as follows:

- First, frames are identified as leaf frames in one of the following cases:
  - Frames arriving on a leaf SAP
  - EVPN traffic arriving with a leaf label
  - Frames arriving with a MAC SA that is flagged as being a leaf SA
- At the egress PE, frames identified as leaf are filtered depending on the type of traffic:
  - For known unicast traffic, the FDB is consulted. If the MAC DA FDB entry is marked as being from a leaf, the frame is dropped to prevent leaf-to-leaf forwarding.
  - For BUM traffic, the leaf frames are filtered at the egress IOM to suppress leaf-to-leaf forwarding.

EVPN-MPLS E-Tree with Multi-homing

Figure 150: Example Topology with All-active ESs and Single-active ES shows the example topology with two all-active multi-homing vESs on PE-2 and PE-3 and one single-active multi-homing ES on PE-4 and PE-5.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 150: Example Topology with All-active ESs and Single-active ES

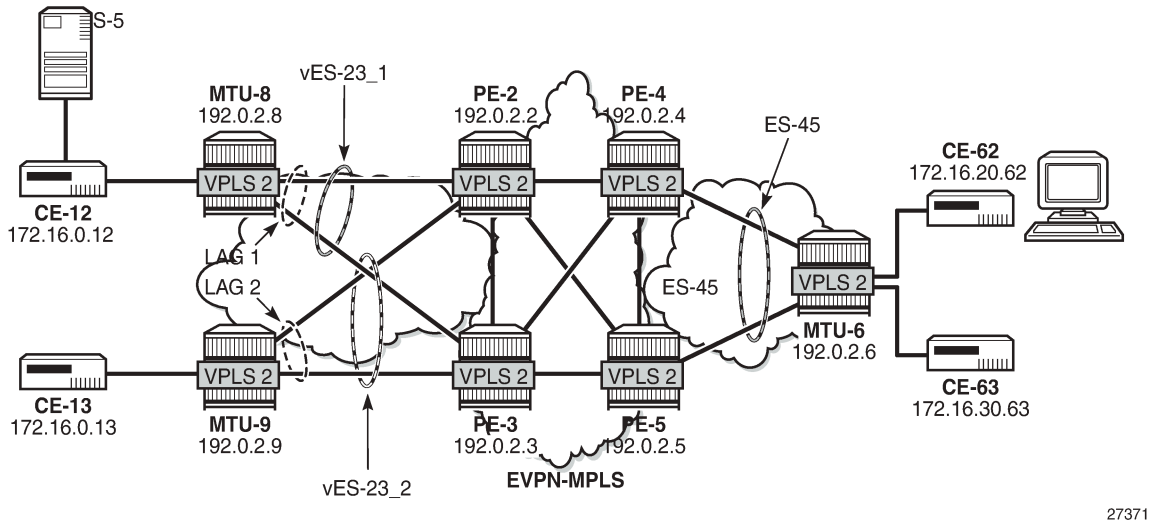

On PE-2, two all-active multi-homing vESs are configured. VPLS 2 is configured as EVPN-MPLS E-Tree with LAG 1 as root AC and LAG 2 as leaf AC. RD 2.2.2.2 is configured and used in the non-zero AD per-ES routes, while the zero ESI routes (AD per-ES) use the IP address 192.0.2.2. The service configuration on PE-2 is as follows:

```
On PE-2:
configure {
    service {
        system {
            bgp {
                evpn {
                    ad-per-es-route {
                        route-target-type evi-route-target-set
                        route-distinguisher-ip-address 2.2.2.2
                    }
                    etree-leaf-label true
                    ethernet-segment "vESI-23_1" {
                        admin-state enable
                        type virtual
                        esi 0x01000000002301000001
                        multi-homing-mode all-active
                        df-election {
                            es-activation-timer 3
                        }
                        association {
                            lag "lag-1" {
                                virtual-ranges {
                                    dot1q {
                                        q-tag 2 {
                                            end 2
                                        }
                                    }
                                }
                            }
                        }
                    }
                    ethernet-segment "vESI-23_2" {
                        admin-state enable
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        type virtual
                        esi 0x01000000002302000001
                        multi-homing-mode all-active
                        df-election {
                            es-activation-timer 3
                        }
                        association {
                            lag "lag-2" {
                                virtual-ranges {
                                    dot1q {
                                        q-tag 2 {
                                            end 2
                                        }
                                    }
                                }
                            }
                        }
                    }
                }
            }
        }
        vpls "VPLS 2" {
            admin-state enable
            service-id 2
            customer "1"
            etree true
            bgp 1 { }
            bgp-evpn {
                evi 2
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            sap lag-1:2 { }
            sap lag-2:2 {
                etree-leaf true
            }
        }
    }
}
```

The service configuration on PE-3 is identical, but with **evi-rt-set route-distinguisher** 3.3.3.3 instead.

Note:
The command **config service system bgp-evpn ad-per-es-route-target evi-rt-set** is not supported for EVPN E-Tree services. When the command is configured on a router, the AD per-ES routes (with ESI=0) used for EVPN E-Tree services are always advertised with the service route target and route distinguisher, regardless of the **ad-per-es-route-target** configuration. AD per-ES routes for non-zero ESIs (used for regular multi-homing) is usually sent using either **evi-rt-set** or **evi-rt**, based on the router configuration.

It is important that all the ACs in each EVI for an ES must either be root ACs or leaf ACs in both PEs where the ES is defined, not a mix. In this example, SAP lag-1:2 is assigned to vES-23_1 and defined as root AC in both PE-2 and PE-3. Likewise, SAP lag-2:2 is assigned to vES-23_2 and configured as leaf AC in PE-2 and PE-3. However, if the configuration were a mix of root and leaf ACs in different PEs of the same ES, a

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

remote PE (PE-4 or PE-5) would receive the AD per-EVI routes with inconsistent leaf indication and would treat the AC as root AC.

PE-2 sends the following BGP EVPN AD routes: an AD per-ES route with zero ESI and RD 192.0.2.2:2 (for egress filtering of BUM traffic) and an EVPN AD per-EVI route with non-zero ESI and RD 2.2.2.2:1 (to verify the ES consistency).

```
On PE-2:
20 2023/08/01 15:11:48.381 CEST MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 81
    Flag: 0x90 Type: 14 Len: 36 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
        Type: EVPN-AD Len: 25 RD: 192.0.2.2:2 ESI: ESI-0, tag: MAX-ET Label: 0 (Raw Label: 0x0)
 PathId:
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64496:2
        etree::L:0/Leaf-Label:524282
        bgp-tunnel-encap:MPLS
"

28 2023/08/01 15:11:48.384 CEST MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 73
    Flag: 0x90 Type: 14 Len: 36 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
        Type: EVPN-AD Len: 25 RD: 2.2.2.2:1 ESI: 01:00:00:00:00:23:01:00:00:01, tag: MAX-ET
 Label: 0 (Raw Label: 0x0) PathId:
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64496:2
        esi-label:524275/All-Active
"
```

The following command shows the EVI RT set RD ranging from 2.2.2.2:1 to 2.2.2.2:512. In VPLS VPLS 2, the configured EVI is 2 and needs to be divided by 128, the number of EVI RT sets that are advertised. This value is rounded up to 1; therefore, the RD in the preceding AD per-EVI equals 2.2.2.2:1. The minimum EVI RT set RD equals 2.2.2.2:1 and the maximum is 2.2.2.2:512, because the EVI ranges from 1 to 65535 and 65536/128=512.

```
[/]
A:admin@PE-2# show service system bgp-evpn

===============================================================================
System BGP EVPN Information
===============================================================================
Eth Seg Route Dist.               : <none>
Eth Seg Oper Route Dist.          : 192.0.2.2:0
Eth Seg Oper Route Dist Type      : default
Ad Per ES Route Target            : evi-rt-set
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
EVI RT set Route Dist.            : 2.2.2.2:1 - 2.2.2.2:512
Extended Evi Range                : Disabled
Etree
  Leaf                            : Enabled
  Leaf Label                      : 524282 (dynamic)
---snip---
===============================================================================
```

Remote PE-4 received the following EVPN AD per-ES routes from PE-2: two non-zero ESI routes (for vES-23_1 and vES-23_2) and a zero ESI route.

```
[/]
A:admin@PE-4# show router bgp routes evpn auto-disc tag MAX-ET next-hop 192.0.2.2
===============================================================================
 BGP Router ID:192.0.2.4        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Auto-Disc Routes
===============================================================================
Flag  Route Dist.         ESI                           NextHop
      Tag                                               Label
-------------------------------------------------------------------------------
u*>i  2.2.2.2:1           01:00:00:00:00:23:01:00:00:01 192.0.2.2
      MAX-ET                                            LABEL 0

u*>i  2.2.2.2:1           01:00:00:00:00:23:02:00:00:01 192.0.2.2
      MAX-ET                                            LABEL 0

u*>i  192.0.2.2:2         ESI-0                         192.0.2.2
      MAX-ET                                            LABEL 0

-------------------------------------------------------------------------------
Routes : 3
===============================================================================
```

On PE-4 and PE-5, ES-45 is configured in single-active mode. The service configuration on PE-4 is as follows:

```
On PE-4:
configure {
    service {
        system {
            bgp {
                evpn {
                    ad-per-es-route {
                        route-target-type evi-route-target-set
                        route-distinguisher-ip-address 4.4.4.4
                    }
                    etree-leaf-label true
                    ethernet-segment "ES-45" {
                        admin-state enable
                        esi 0x01000000004501000001
                        multi-homing-mode single-active
                        df-election {
                            es-activation-timer 3
                            service-carving-mode manual
                            manual {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                                preference {
                                    value 10000
                                }
                            }
                        }
                        association {
                            sdp 46 { }
                        }
                    }
                }
            }
        }
        vpls "VPLS 2" {
            admin-state enable
            service-id 2
            customer "1"
            etree true
            bgp 1 { }
            bgp-evpn {
                evi 2
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            spoke-sdp 46:2 {
                etree-leaf true
            }
        }
    }
}
```

The service configuration is similar on PE-5, but with a lower preference for the ES, so PE-4 is the DF, as follows.

```
[/]
A:admin@PE-4# show service id 2 ethernet-segment
No sap entries

===============================================================================
SDP Ethernet-Segment Information
===============================================================================
SDP                  Eth-Seg                         Status
-------------------------------------------------------------------------------
46:2                 ES-45                           DF
===============================================================================
No vxlan instance entries
```

For the all-active multi-homing vESs, PE-2 is the DF, as follows:

```
[/]
A:admin@PE-2# show service id 2 ethernet-segment

===============================================================================
SAP Ethernet-Segment Information
===============================================================================
SAP                  Eth-Seg                         Status
-------------------------------------------------------------------------------
lag-1:2              vESI-23_1                       DF
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
lag-2:2              vESI-23_2                       DF
===============================================================================
No sdp entries
No vxlan instance entries
```

Ingress Filtering for Unicast Traffic

Traffic can be sent between CE-12 (root AC lag-1:2) and CE-62 (leaf AC spoke SDP 46:2), but traffic between CE-13 (leaf AC lag-2:2) and CE-63 (leaf AC spoke SDP 46:2) is filtered. The following FDB for VPLS 2 on PE-2 shows two EVPN leaf MAC addresses: ca:fe:06:00:20:62 for CE-62 and ca:fe:06:00:30:63 for CE-63.

```
[/]
A:admin@PE-2# show service id 2 fdb detail

===============================================================================
Forwarding Database, Service 2
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
2          ca:fe:01:00:20:12 sap:lag-1:2             L/0      08/01/23 15:14:06
2          ca:fe:01:00:30:13 sap:lag-2:2             Evpn     08/01/23 15:14:09
2          ca:fe:06:00:20:62 eES:                    Evpn,Lf  08/01/23 15:12:37
                             01:00:00:00:00:45:01:00:00:01
2          ca:fe:06:00:30:63 eES:                    Evpn,Lf  08/01/23 15:14:31
                             01:00:00:00:00:45:01:00:00:01
-------------------------------------------------------------------------------
No. of MAC Entries: 4
-------------------------------------------------------------------------------
Legend:L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf T=Trusted
===============================================================================
```

The FDB for VPLS 2 on PE-3 shows the same EVPN leaf MAC addresses. For all PEs in the all-active MH ESs, the MAC addresses ca:fe:06:00:20:12 and ca:fe:06:00:30:13 from the locally attached ACs can be learned on the SAPs or via EVPN from the ES peer where they are learned on the SAPs. In this case, they are learned on the SAPs on PE-2 and PE-3.

```
[/]
A:admin@PE-3# show service id 2 fdb detail

===============================================================================
Forwarding Database, Service 2
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
2          ca:fe:01:00:20:12 sap:lag-1:2             L/0      08/01/23 15:12:17
2          ca:fe:01:00:30:13 sap:lag-2:2             L/0      08/01/23 15:14:09
2          ca:fe:06:00:20:62 eES:                    Evpn,Lf  08/01/23 15:12:36
                             01:00:00:00:00:45:01:00:00:01
2          ca:fe:06:00:30:63 eES:                    Evpn,Lf  08/01/23 15:14:31
                             01:00:00:00:00:45:01:00:00:01
-------------------------------------------------------------------------------
No. of MAC Entries: 4
-------------------------------------------------------------------------------
Legend:L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf T=Trusted
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The following FDB for VPLS 2 on DF PE-4 shows one EVPN leaf MAC address: ca:fe:01:00:30:13 for CE-13 on a remote ES.

```
[/]
A:admin@PE-4# show service id 2 fdb detail

===============================================================================
Forwarding Database, Service 2
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
2          ca:fe:01:00:20:12 eES:                    Evpn     08/01/23 15:12:25
                             01:00:00:00:00:23:01:00:00:01
2          ca:fe:01:00:30:13 eES:                    Evpn,Lf  08/01/23 15:14:09
                             01:00:00:00:00:23:02:00:00:01
2          ca:fe:06:00:20:62 sdp:46:2                L/0      08/01/23 15:12:36
2          ca:fe:06:00:30:63 sdp:46:2                L/0      08/01/23 15:14:31
-------------------------------------------------------------------------------
No. of MAC Entries: 4
-------------------------------------------------------------------------------
Legend:L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf T=Trusted
===============================================================================
```

PE-5 is NDF, and the following FDB shows three MAC routes of type EVPN leaf, for CE-13, CE-62, and CE-63.

```
[/]
A:admin@PE-5# show service id 2 fdb detail

===============================================================================
Forwarding Database, Service 2
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
2          ca:fe:01:00:20:12 eES:                    Evpn     08/01/23 15:12:38
                             01:00:00:00:00:23:01:00:00:01
2          ca:fe:01:00:30:13 eES:                    Evpn,Lf  08/01/23 15:14:09
                             01:00:00:00:00:23:02:00:00:01
2          ca:fe:06:00:20:62 eES:                    Evpn,Lf  08/01/23 15:12:38
                             01:00:00:00:00:45:01:00:00:01
2          ca:fe:06:00:30:63 eES:                    Evpn,Lf  08/01/23 15:14:31
                             01:00:00:00:00:45:01:00:00:01
-------------------------------------------------------------------------------
No. of MAC Entries: 4
-------------------------------------------------------------------------------
Legend:L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf T=Trusted
===============================================================================
```

Egress Filtering for BUM Traffic

Each PE advertises zero ESI AD per-ES routes (with tag MAX-ET) that are needed for egress BUM filtering.

BUM frames received on an ES root AC are flooded to the EVPN, based on regular EVPN procedures. The regular ESI label is sent for split horizon when frames are sent to the DF or NDF PEs in the same ES.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

BUM frames received on an ES leaf AC are flooded in the default multicast list. The egress PE does not forward BUM traffic to any leaf ACs, including the ES leaf ACs. However, in the unlikely event that some ACs in a specific ES for an EVI have an inconsistent E-Tree configuration, these ACs are treated as root ACs, and the traffic is forwarded.

The remote PE-4 receives the following EVPN AD routes from DF PE-2: a zero ESI AD per-ES (tag MAX-ET), two AD per-EVI (tag 0) routes with a non-zero label, and two AD per-ES routes (tag MAX-ET).

```
[/]
A:admin@PE-4# show router bgp routes evpn auto-disc next-hop 192.0.2.2
===============================================================================
 BGP Router ID:192.0.2.4        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Auto-Disc Routes
===============================================================================
Flag  Route Dist.         ESI                           NextHop
      Tag                                               Label
-------------------------------------------------------------------------------
u*>i  2.2.2.2:1           01:00:00:00:00:23:01:00:00:01 192.0.2.2
      MAX-ET                                            LABEL 0

u*>i  2.2.2.2:1           01:00:00:00:00:23:02:00:00:01 192.0.2.2
      MAX-ET                                            LABEL 0

u*>i  192.0.2.2:2         ESI-0                         192.0.2.2
      MAX-ET                                            LABEL 0

u*>i  192.0.2.2:2         01:00:00:00:00:23:01:00:00:01 192.0.2.2
      0                                                 LABEL 524277

u*>i  192.0.2.2:2         01:00:00:00:00:23:02:00:00:01 192.0.2.2
      0                                                 LABEL 524277

-------------------------------------------------------------------------------
Routes : 5
===============================================================================
```

The same remote PE-4 receives similar EVPN AD routes from NDF PE-3: a zero ESI AD per-ES (tag MAX-ET), two AD per-EVI (tag 0) routes with a non-zero label, and two AD per-ES routes (tag MAX-ET).

```
[/]
A:admin@PE-4# show router bgp routes evpn auto-disc next-hop 192.0.2.3
===============================================================================
 BGP Router ID:192.0.2.4        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Auto-Disc Routes
===============================================================================
Flag  Route Dist.         ESI                           NextHop
      Tag                                               Label
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
u*>i  3.3.3.3:1           01:00:00:00:00:23:01:00:00:01 192.0.2.3
      MAX-ET                                          LABEL 0

u*>i  3.3.3.3:1           01:00:00:00:00:23:02:00:00:01 192.0.2.3
      MAX-ET                                          LABEL 0

u*>i  192.0.2.3:2         ESI-0                         192.0.2.3
      MAX-ET                                          LABEL 0

u*>i  192.0.2.3:2         01:00:00:00:00:23:01:00:00:01 192.0.2.3
      0                                               LABEL 524280

u*>i  192.0.2.3:2         01:00:00:00:00:23:02:00:00:01 192.0.2.3
      0                                               LABEL 524280

-------------------------------------------------------------------------------
Routes : 5
===============================================================================
```

The following detailed information about the AD per-ES route (tag MAX-ET) for mass withdraw on PE-4 shows that no E-Tree extended community is sent by PE-2; only the ESI-label extended community is sent.

```
[/]
A:admin@PE-4# show router bgp routes evpn auto-disc rd 2.2.2.2:1 tag MAX-ET esi
 01:00:00:00:00:23:01:00:00:01 detail
===============================================================================
 BGP Router ID:192.0.2.4        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Auto-Disc Routes
===============================================================================
Original Attributes

Network        : n/a
Nexthop        : 192.0.2.2
Path Id        : None
From           : 192.0.2.2
Res. Nexthop   : 192.168.24.1
---snip---
Community      : target:64496:2 esi-label:524275/All-Active
---snip---
EVPN type      : AUTO-DISC
ESI            : 01:00:00:00:00:23:01:00:00:01
Tag            : MAX-ET
Route Dist.    : 2.2.2.2:1
MPLS Label     : LABEL 0
---snip---
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

A similar result is seen for the other vES:

```
[/]
A:admin@PE-4# show router bgp routes evpn auto-disc rd 2.2.2.2:1 tag MAX-ET esi
 01:00:00:00:00:23:02:00:00:01 detail
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
 BGP Router ID:192.0.2.4        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Auto-Disc Routes
===============================================================================
Original Attributes

Network        : n/a
Nexthop        : 192.0.2.2
Path Id        : None
From           : 192.0.2.2
Res. Nexthop   : 192.168.24.1
---snip---
Community      : target:64496:2 esi-label:524274/All-Active
---snip---
EVPN type      : AUTO-DISC
ESI            : 01:00:00:00:00:23:02:00:00:01
Tag            : MAX-ET
Route Dist.    : 2.2.2.2:1
MPLS Label     : LABEL 0
---snip---
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

The following detailed information about the AD per-EVI (tag 0) on PE-4 shows that if the ES is root (as for vES-23_1), the regular extended community is sent, not the E-Tree extended community.

```
[/]
A:admin@PE-4# show router bgp routes evpn auto-disc rd 192.0.2.2:2 tag 0 esi
 01:00:00:00:00:23:01:00:00:01 detail
===============================================================================
 BGP Router ID:192.0.2.4        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Auto-Disc Routes
===============================================================================
Original Attributes

Network        : n/a
Nexthop        : 192.0.2.2
Path Id        : None
From           : 192.0.2.2
Res. Nexthop   : 192.168.24.1
---snip---
Community      : target:64496:2 bgp-tunnel-encap:MPLS
---snip---
EVPN type      : AUTO-DISC
ESI            : 01:00:00:00:00:23:01:00:00:01
Tag            : 0
Route Dist.    : 192.0.2.2:2
MPLS Label     : LABEL 524277
---snip---
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

The following detailed information about the AD per-EVI (tag 0) on PE-4 shows that if the ES is leaf (as for vES-23_2), the E-Tree extended community is sent, along with the regular extended community.

```
[/]
A:admin@PE-4# show router bgp routes evpn auto-disc rd 192.0.2.2:2 tag 0 esi
 01:00:00:00:00:23:02:00:00:01 detail
===============================================================================
 BGP Router ID:192.0.2.4        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Auto-Disc Routes
===============================================================================
Original Attributes

Network        : n/a
Nexthop        : 192.0.2.2
Path Id        : None
From           : 192.0.2.2
Res. Nexthop   : 192.168.24.1
---snip---
Community      : target:64496:2 etree::L:1/Leaf-Label:0
                 bgp-tunnel-encap:MPLS
---snip---
EVPN type      : AUTO-DISC
ESI            : 01:00:00:00:00:23:02:00:00:01
Tag            : 0
Route Dist.    : 192.0.2.2:2
MPLS Label     : LABEL 524277
---snip---
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

The **tools dump service evpn usage** command shows that there are three EVPN E-Tree remote BUM leaf labels:

```
[/]
A:admin@PE-2# tools dump service evpn usage | match "BUM"
Evpn Etree Remote BUM Leaf Labels               :            3
```

This corresponds to the following three ESI-0 AD per-ES routes (tag MAX-ET) on PE-2:

```
[/]
A:admin@PE-2# show router bgp routes evpn auto-disc esi ESI-0
===============================================================================
 BGP Router ID:192.0.2.2        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
BGP EVPN Auto-Disc Routes
===============================================================================
Flag  Route Dist.         ESI                         NextHop
      Tag                                             Label
-------------------------------------------------------------------------------
u*>i  192.0.2.3:2         ESI-0                       192.0.2.3
      MAX-ET                                          LABEL 0

u*>i  192.0.2.4:2         ESI-0                       192.0.2.4
      MAX-ET                                          LABEL 0

u*>i  192.0.2.5:2         ESI-0                       192.0.2.5
      MAX-ET                                          LABEL 0

-------------------------------------------------------------------------------
Routes : 3
===============================================================================
```

Conclusion

E-Trees can be used for enterprise business services, for the distribution of IPTV multicast content, for centralized backup BNGs, and so on. In a VPLS E-Tree, leaf SAPs or leaf SDP bindings cannot exchange traffic with each other, similar to split horizon group behavior. The E-Tree restrictions apply to all remote PEs that are part of the same service. E-Trees can be applied in an EVPN-MPLS VPLS as well as in a regular VPLS.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

EVPN-MPLS Interconnect for EVPN-VXLAN VPLS Services

This chapter provides information about EVPN-MPLS Interconnect for EVPN-VXLAN VPLS Services.

Topics in this chapter include:

Applicability

This chapter was initially written for SR OS Release 14.0.R5, but the MD-CLI in the current edition is based on SR OS Release 21.2.R1.

Chapters EVPN for MPLS Tunnels and EVPN for VXLAN Tunnels (Layer 2) are prerequisite reading.

Overview

When EVPN-MPLS is deployed in the WAN, many service providers are looking for a way to integrate existing Layer 2 EVPN-VXLAN based data center services into the WAN, while keeping the end-to-end advantages of EVPN. The IETF *draft-ietf-bess-dci-evpn-overlay* describes how to provide Layer 2 connectivity for EVPN-overlay data centers in different ways. This chapter follows section 4.4 of that document, in which EVPN-MPLS is used in the same VPLS service that terminates overlay (VXLAN) tunnels.

To provide EVPN-MPLS connectivity to VPLS services terminating EVPN-VXLAN, SR OS supports the configuration of BGP-EVPN MPLS and BGP-EVPN VXLAN at the same time by adding two BGP instances to the service. Two BGP instances are supported in the same VPLS at most. BGP-EVPN MPLS and BGP-EVPN VXLAN can both use BGP instance 1 or 2, but they must use different instances.

In a service with EVPN-VXLAN and EVPN-MPLS, the **config>service>vpls>bgp-evpn>mpls 2** command allows the user to associate BGP-EVPN MPLS to BGP instance 2, while BGP-EVPN VXLAN is associated to BGP instance 1, and therefore, have both encapsulations simultaneously enabled in the same service. Either BGP instance 1 or 2 can be associated to BGP-EVPN VXLAN or MPLS, but they must be different. When the two BGP instances are successfully added to the same VPLS service, the service behaves as follows:

- MAC/IP routes received on one instance will be "consumed" (accepted, imported, and installed in FDB) and re-advertised in the other instance, as long as the route is the best route for a specific MAC or MAC/IP.
- Inclusive multicast routes are independently generated for each BGP instance.
- From a data plane perspective, EVPN-MPLS and EVPN-VXLAN destinations are instantiated in different implicit Split-Horizon Groups (SHGs) so that traffic can be forwarded between the two SHGs, but not between destinations of the same kind. For example, traffic coming from EVPN-MPLS cannot be forwarded to other destinations in the EVPN-MPLS SHG.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The following example shows a VPLS service configured on PE-2 with two BGP instances and both encapsulations, VXLAN and MPLS, configured at the same time:

```
configure {
    service {
        vpls "VPLS 1" {
            description "evpn-mpls and evpn-vxlan in the same service"
            admin-state enable
            service-id 1
            customer "1"
            vxlan {
                instance 1 {
                    vni 1
                }
            }
            bgp 1 {
                route-distinguisher "10:1"
                route-target {
                    export "target:64500:1"
                    import "target:64500:1"
                }
            }
            bgp 2 {
                route-distinguisher "10:2"
                route-target {
                    export "target:64500:1"
                    import "target:64500:1"
                }
            }
            bgp-evpn {
                evi 1
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
                mpls 2 {
                    admin-state enable
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
```

In the preceding example

- **bgp 1** is the default BGP instance.
- **bgp 2** is the additional instance that is required when both BGP-EVPN VXLAN and BGP-EVPN MPLS are enabled in the service.
- The same commands supported under BGP instance 1 exist for this second BGP instance, with the following considerations:
  - **pw-template-binding** – the pseudowire (PW) template binding can only exist in BGP instance 1; it is not supported in BGP instance 2. Because no SDP-bindings can exist in a VPLS service with two BGP instances, the **pw-template-binding** command is ineffective in this configuration.
  - **route-distinguisher** – the route distinguisher in both BGP instances must be different.
  - **route-target** – the route target in both instances can be the same or different.
  - **vsi-import** and **vsi-export** – import and export policies can also be defined for either BGP instance.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- The **mpls 2** command will assign BGP instance 2 to MPLS. The VPLS configuration can only be committed if the BGP instance associated with MPLS has a different route distinguisher than the BGP instance associated with VXLAN.
- The **evi** can still be used for auto-derivation of RD/RT on BGP instance 1 and auto-derivation of RT (not RD) on BGP instance 2. Auto-RD or an explicitly configured RD is needed in BGP instance 2.

Configuration

Figure 151: EVPN-MPLS interconnect for EVPN-VXLAN - example topology shows the example topology that will be used throughout this chapter, as well as the BGP peering topology. PE-1, PE-2, and PE-3 simulate a data center, shown as Overlay-Network-1, where PE-2 and PE-3 are DC GWs. In the same way, PE-4, PE-5, and PE-6 simulate a remote data center, Overlay-Network-2. Inside each DC, EVPN-VXLAN is used.

The two DC GW pairs are connected by EVPN-MPLS; therefore, CE-1 and CE-6 are end-to-end connected by EVPN without any VLAN or PW hand-off, maintaining all the EVPN advantages across the DC Interconnect (DCI) network.

Figure 151: EVPN-MPLS interconnect for EVPN-VXLAN - example topology

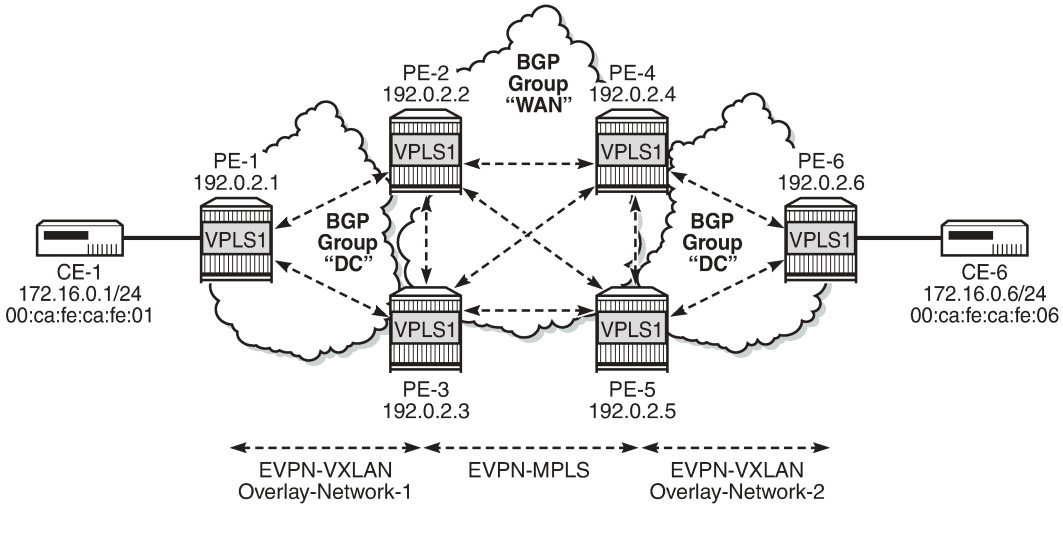

The example topology consists of six 7750 SR routers with the following initial configuration:

- Hybrid ports (they could have been network type too) are interconnecting the six PEs with configured router interfaces.
- The six PEs are running IS-IS and creating point-to-point adjacencies.
- Link LDP is configured in the core, among PE-2, PE-3, PE-4, and PE-5, while PE-1 and PE-6 are only running VXLAN.
- EVPN uses MP-BGP for exchanging reachability at service level. Therefore, BGP peering sessions must be established among the PEs for the EVPN family. Figure 151: EVPN-MPLS interconnect for EVPN-VXLAN - example topology shows the peering sessions established among the six PEs. Although usually a Route-Reflector (RR) is used in each DC and another RR in the WAN, in this example, there are direct peering sessions in each DC and in the WAN.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The following output shows the BGP configuration of PE-2. The BGP configuration on the rest of the DC GWs (PE-3, PE-4, and PE-5) is similar:

```
# on PE-2:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            family {
                ipv4 false
                evpn true
            }
            rapid-update {
                evpn true
            }
            group "DC" {
                type internal
                import {
                    policy ["drop SOO-DCGW-23"]
                }
                export {
                    policy ["allow only vxlan and add SOO"]
                }
            }
            group "WAN" {
                type internal
                import {
                    policy ["drop SOO-DCGW-23"]
                }
                export {
                    policy ["allow only mpls and add SOO"]
                }
            }
            neighbor "192.0.2.1" {
                group "DC"
            }
            neighbor "192.0.2.3" {
                group "DC"
            }
            neighbor "192.0.2.4" {
                group "WAN"
            }
            neighbor "192.0.2.5" {
                group "WAN"
            }
        }
```

Two different BGP groups are configured: DC and WAN. The DC group contains the DC neighbors (including the peer DC GW) and the WAN group contains the WAN neighbors. This grouping makes the use of policies easier. These policies will be explained in the section The mandatory use of BGP policies in the multi-homed anycast solution.

The following output shows the BGP configuration of PE-1. PE-6 has a similar BGP configuration.

```
# on PE-1:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            rapid-withdrawal true
            family {
                ipv4 false
                evpn true
            }
            rapid-update {
                evpn true
            }
            group "DC" {
                type internal
            }
            neighbor "192.0.2.2" {
                group "DC"
            }
            neighbor "192.0.2.3" {
                group "DC"
            }
        }
```

VPLS service configuration

After the base infrastructure (interfaces, IGP, LDP in the core, and BGP) is configured, the services can be added. The configuration example in this section will use VPLS 1 as the service to be interconnected across the two DCs.

PE-1 and PE-6 have a regular EVPN-VXLAN configuration; DCI connectivity provided by EVPN-MPLS is completely transparent to them. The configuration of VPLS 1 in PE-1 is as follows:

```
# on PE-1:
configure {
    service {
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            vxlan {
                instance 1 {
                    vni 1
                }
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 1
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
            sap 1/2/1:1 {
            }
        }
```

See the EVPN for VXLAN Tunnels (Layer 2) chapter for a complete description of the EVPN-VXLAN commands.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The configuration on PE-2, PE-3, PE-4, and PE-5 (see Figure 151: EVPN-MPLS interconnect for EVPN-VXLAN - example topology) enables EVPN-VXLAN and EVPN-MPLS in the same VPLS service. As an example, the VPLS 1 configuration on PE-2 is as follows:

```
# on PE-2:
configure {
    service {
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            vxlan {
                instance 1 {
                    vni 1
                }
            }
            bgp 1 {
                route-distinguisher "64500:1"
            }
            bgp 2 {
                route-distinguisher "64500:2"
            }
            bgp-evpn {
                evi 1
                incl-mcast-orig-ip 23.23.23.23
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
                mpls 2 {
                    admin-state enable
                    ingress-replication-bum-label true
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
        }
```

As described in the Overview section, the preceding configuration enables the router to create EVPN-VXLAN and EVPN-MPLS destinations in the same VPLS service, but in different SHGs. In addition to the **bgp 2** commands already described in the Overview section, the **incl-mcast-orig-ip** command is added in the configuration. If configured, this command will change the originating IP address in the inclusive multicast routes (from the default system IP) for both BGP instances. The section Multi-homed anycast configuration for dual BGP-instance VPLS services describes why this command is added.

The following section provides a detailed description of the expected behavior for EVPN routes that are imported and exported on dual BGP instance VPLS services.

EVPN route handling in dual BGP-instance VPLS services

This section describes how the BGP-EVPN routes are processed in dual BGP instance services.

Usually, the router validates the received tunnel encapsulation (from the RFC 5512 Extended Community) with the configured encapsulation of the service/BGP-instance. Therefore, an EVPN-VXLAN route will not get imported into the BGP-EVPN MPLS instance and vice-versa. This is also how the different EVPN route types are handled in dual BGP instance services:

- **Route type 1 - auto-discovery routes**

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

AD per-EVI routes are never generated by services with two BGP instances (because no Ethernet Segment (ES) can be associated with the dual BGP instance service). However, AD per-EVI routes can still be received from the EVPN-MPLS peers and are processed as usual. Therefore, a VPLS service with two BGP instances will still support aliasing/backup and AD per-ES checking procedures for a remote multi-homed ES, as described in the EVPN for MPLS Tunnels chapter. However, in the example in Figure 151: EVPN-MPLS interconnect for EVPN-VXLAN - example topology, PE-6 does not have any local multi-homed ES configured; therefore, no AD per-EVI routes are present in this example.

- **Route type 2 - MAC/IP routes**

  MAC/IP routes received on one of the two BGP instances will be imported and the MAC addresses added to the FDB according to the existing selection rules. If the MAC address is active (therefore installed in the FDB), it will be re-advertised in the other BGP instance with the BGP attributes of the other BGP instance (new route target if different, new route distinguisher, and so on). The **bgp-evpn>routes>mac-ip>advertise** command will govern the advertisement of MAC addresses in either BGP instance.

  The MAC/IP route redistribution across BGP instances is performed according to the following rules:

  – A MAC route is redistributed only if it is the best route according to the EVPN selection rules in the EVPN for MPLS Tunnels chapter.

  – Assuming a specific MAC route is the best one and has to be redistributed, the MAC/IP information along with the sticky bit is propagated in the redistribution.

  – A change in the MAC/IP route sequence number or sticky bit in one instance is updated in the other instance, as long as that route is the best MAC route for the route key.

  – When a MAC address moves within the EVPN-VXLAN (or the EVPN-MPLS) network, the MAC route is received on the same BGP instance where it was previously received, but now with a higher sequence number. In this case, the MAC route will be redistributed with the new sequence number. However, a router with two BGP instances in the same service will not detect any duplicate MAC on the EVPN-VXLAN and EVPN-MPLS networks.

  As an example, the following output shows the debug of a MAC/IP route received on PE-2, on the BGP instance for EVPN-VXLAN on VPLS 1, and how the route is re-advertised to the BGP instance used for MPLS (with a different next-hop, route distinguisher, label, and BGP tunnel encapsulation):

```
# on PE-2:
14 2021/03/16 17:31:10.452 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.1
"Peer 1: 192.0.2.1: UPDATE
Peer 1: 192.0.2.1 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 81
    Flag: 0x90 Type: 14 Len: 44 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.1
        Type: EVPN-MAC Len: 33 RD: 192.0.2.1:1 ESI: ESI-0, tag: 0, mac len: 48
                       mac: 00:ca:fe:ca:fe:01, IP len: 0, IP: NULL, label1: 1
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:1
        bgp-tunnel-encap:VXLAN
"
```

```
15 2021/03/16 17:31:10.453 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.4
"Peer 1: 192.0.2.4: UPDATE
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Peer 1: 192.0.2.4 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 89
    Flag: 0x90 Type: 14 Len: 44 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
        Type: EVPN-MAC Len: 33 RD: 64500:2 ESI: ESI-0, tag: 0, mac len: 48
                       mac: 00:ca:fe:ca:fe:01, IP len: 0, IP: NULL, label1: 8388528
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        origin:64500:23
        target:64500:1
        bgp-tunnel-encap:MPLS
"
```

```
16 2021/03/16 17:31:10.453 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.5
"Peer 1: 192.0.2.5: UPDATE
Peer 1: 192.0.2.5 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 89
    Flag: 0x90 Type: 14 Len: 44 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
        Type: EVPN-MAC Len: 33 RD: 64500:2 ESI: ESI-0, tag: 0, mac len: 48
                       mac: 00:ca:fe:ca:fe:01, IP len: 0, IP: NULL, label1: 8388528
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        origin:64500:23
        target:64500:1
        bgp-tunnel-encap:MPLS
"
```

- **Route type 3 - inclusive multicast routes**

  EVPN Inclusive Multicast Ethernet Tag (IMET) routes are generated independently for each BGP instance with the correct BGP tunnel encapsulation extended community and the tunnel type associated to the BGP instance; for example, Ingress Replication (IR), P2MP mLDP, or Assisted Replication (AR):

  – On the EVPN-VXLAN BGP instance, IR or AR IMET routes are supported.

    - When **assisted-replication replicator** is enabled and the received VXLAN broadcast and multicast packets contain an IP DA = AR-IP, the DC GW will send the packets back to VXLAN (but not to the VXLAN termination end-point (VTEP) from where the packet is received) in addition to the EVPN-MPLS destinations.
    - If **assisted-replication replicator** is used on the DC GWs, the AR-IP (**configure>service>system>vxlan>assisted-replication>ip-address**) must be a loopback different from the router's system IP and the configured **bgp-evpn>incl-mcast-orig-ip**. The two AR-IP addresses in the DC GW pair do not need to be the same IP address.

  – On the EVPN-MPLS BGP instance, IR, P2MP mLDP, or composite IMET routes are supported.

  – Following is the behavior when the **incl-mcast-orig-ip** command is used:

    - The configured IP in the **incl-mcast-orig-ip** command is encoded in the originating IP field of the IMET routes for IR, P2MP, and composite routes for both BGP instances.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- The originating IP field of the IMET AR routes is still derived from the configured **service>system>vxlan>assisted-replication>ip-address** value.

– The received IMET routes will be processed in the following way depending on their type:
  - IMET-IR routes: the EVPN destination (MPLS or VXLAN) is set up based on the NLRI next-hop.
  - IMET-P2MP routes: the Provider Multicast Service Interface (PMSI) Tunnel Attribute (PTA) tunnel ID will be used to join the mLDP tree (as mLDP FEC in the LDP mapping messages).
  - IMET-P2MP-IR (composite) routes: the PTA tunnel ID is used to join the mLDP tree. The NLRI next-hop is used to build the EVPN destination.
  - IMET-AR routes: the NLRI next-hop is used to build the EVPN-VXLAN destination.

– Upon reception of two IMET routes with similar information, the router behaves as follows:
  - If the router receives two IMET routes with the same originating IP, different RDs, and different NLRI next-hops, it will set up two EVPN destinations, one to each next-hop.
  - If the router gets two IMET routes with the same originating IP, different RDs, but the same next-hop, it will set up only one EVPN destination.
  - The router will not set up an EVPN destination to its DC GW peer if the received originating IP matches its own originating IP, regardless of whether the local RD and the remote RD are the same or different. This enables the use of the redundant anycast solution that is described in the following section: Multi-homed anycast configuration for dual BGP-instance VPLS services.

- **Route type 4 - ES routes**

  ESs are supported in routers where dual BGP-instance services exist. However, because dual BGP-instance VPLS services do not support SDP-bindings, ESs and ES routes are not relevant to these types of services.

- **Route type 5 - IP-prefix routes**

  R-VPLS services are not supported along with dual BGP instances; therefore, IP-prefix routes are neither generated nor processed by the service.

Multi-homed anycast configuration for dual BGP-instance VPLS services

Services with EVPN-MPLS and EVPN-VXLAN SHGs are specified in *draft-ietf-bess-dci-evpn-overlay* and the associated multi-homing solution is also described in the same draft. That multi-homing solution is based on an interconnect ES that allows all-active and single-active multi-homed EVPN networks as well as local attachment circuits in the DC GWs (SAP/SDP-bindings).

This chapter was initially written for SR OS Release 14.0.R5 and interconnect ESs were not supported in that release. Therefore, an anycast solution is used to provide redundancy. This anycast solution is based on the two PE DC GWs in the redundant pair being configured to advertised MAC/IP and IMET routes with the same route key, so that the remote PEs will only pick up one of the two anycast DC GWs when sending unicast or BUM traffic, and no loop or packet duplication is created.

Figure 152: EVPN destinations created on multi-homed anycast DC GWs is an example of how multi-homing can be achieved for dual BGP-instance VPLS services. The figure also shows the EVPN destinations created and their direction (see the arrows). For instance, only one EVPN multicast destination is created for PE-1, PE-2, or PE-4. Therefore, BUM traffic sent by CE-1 will be sent via PE-2, PE-4, and PE-6 only, and no duplication or loops occur.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 152: EVPN destinations created on multi-homed anycast DC GWs

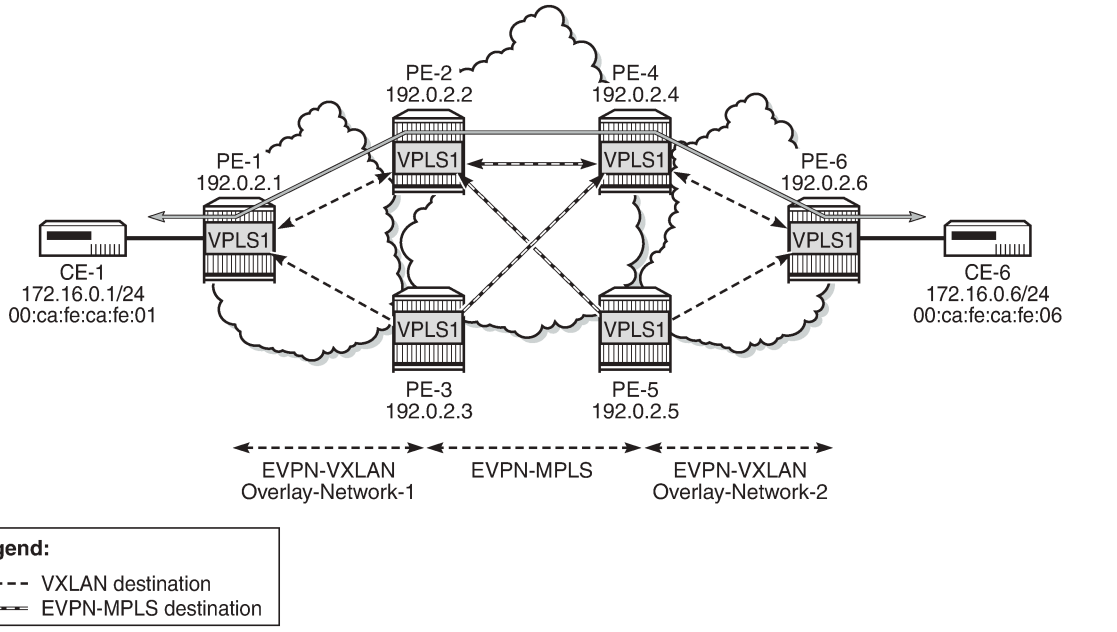

The following output shows the VPLS 1 configuration on PE-2 and PE-3 so that this anycast redundancy can be realized. The route distinguishers as well as the **incl-mcast-orig-ip** addresses must match between the two PEs in the redundant pair. VPLS 1 is configured on PE-2 as follows:

```
# on PE-2:
configure {
    service {
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            vxlan {
                instance 1 {
                    vni 1
                }
            }
            bgp 1 {
                route-distinguisher "64500:1"
            }
            bgp 2 {
                route-distinguisher "64500:2"
            }
            bgp-evpn {
                evi 1
                incl-mcast-orig-ip 23.23.23.23
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
                mpls 2 {
                    admin-state enable
                    ingress-replication-bum-label true
                    auto-bind-tunnel {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        resolution any
                    }
                }
            }
        }
```

The VPLS 1 configuration on PE-3 is as follows:

```
# on PE-3:
configure {
    service {
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            vxlan {
                instance 1 {
                    vni 1
                }
            }
            bgp 1 {
                route-distinguisher "64500:1"
            }
            bgp 2 {
                route-distinguisher "64500:2"
            }
            bgp-evpn {
                evi 1
                incl-mcast-orig-ip 23.23.23.23
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
                mpls 2 {
                    admin-state enable
                    ingress-replication-bum-label true
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
        }
```

The VPLS 1 configuration on PE-4 is as follows:

```
# on PE-4:
configure {
    service {
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            vxlan {
                instance 1 {
                    vni 1
                }
            }
            bgp 1 {
                route-distinguisher "64501:1"
            }
            bgp 2 {
                route-distinguisher "64501:2"
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            bgp-evpn {
                evi 1
                incl-mcast-orig-ip 45.45.45.45
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
                mpls 2 {
                    admin-state enable
                    ingress-replication-bum-label true
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
        }
```

The VPLS 1 configuration on PE-5 is as follows:

```
# on PE-5:
configure {
    service {
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            vxlan {
                instance 1 {
                    vni 1
                }
            }
            bgp 1 {
                route-distinguisher "64501:1"
            }
            bgp 2 {
                route-distinguisher "64501:2"
            }
            bgp-evpn {
                evi 1
                incl-mcast-orig-ip 45.45.45.45
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
                mpls 2 {
                    admin-state enable
                    ingress-replication-bum-label true
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
        }
```

Based on the preceding configuration example, the DC GWs behavior in this scenario is as follows:

- PE-2 and PE-3 both send IMET IR routes to the other PEs with the same route key but a different next-hop. The route key in IMET routes comprises [RD, Ethernet tag, originator-IP/length], which in this case will be [64500:1, 0, 23.23.23.23/32] for the EVPN-VXLAN IMET routes and [64500:2, 0, 23.23.23.23/32] for the EVPN-MPLS IMET routes.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- In the same way, PE-2 and PE-3 both send MAC/IP routes to the other PEs with the same route key but a different next-hop. The route key comprises [RD, Ethernet tag, MAC/MAC-length, IP/IP-length].

The configuration of the same **incl-mcast-orig-ip** address and RDs in both DC GWs enables the anycast solution due to the following:

- The configured originating IP (for example, 23.23.23.23 in PE-2 and PE-3) is not required to be a reachable IP address, which forces the remote PEs (or RRs if they exist) to select only one of the two DC GWs for BUM traffic (based on regular BGP selection). In this example, the remote PEs will select the PE-2 IMET route and create only one destination. The following output shows the IMET routes received by PE-1 (only the PE-2 route is used) and the created EVPN-VXLAN destination to PE-2. The same behavior could have been shown in the rest of the PEs.

```
[/]
A:admin@PE-1# show router bgp routes evpn incl-mcast
===============================================================================
 BGP Router ID:192.0.2.1        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Inclusive-Mcast Routes
===============================================================================
Flag  Route Dist.          OrigAddr
      Tag                  NextHop
-------------------------------------------------------------------------------
u*>i  64500:1              23.23.23.23
      0                    192.0.2.2

*>i   64500:1              23.23.23.23
      0                    192.0.2.3

-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

```
[/]A:admin@PE-1# show service id 1 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
Instance    VTEP Address                          Egress VNI  EvpnStatic Num
 Mcast       Oper State                            L2 PBR     SupBcasDom MACs
-------------------------------------------------------------------------------
1           192.0.2.2                             1           evpn       1
 BUM         Up                                    No          No
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 1
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN-VXLAN Ethernet Segment Dest
===============================================================================
Instance  Eth SegId                       Num. Macs     Last Change
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- Due to the same RD and originating IP configured on PE-2 and PE3 (similarly in PE-4 and PE-5), the DC GW redundant PEs will never establish an EVPN destination between each other. PE-2 only sets up EVPN multicast destinations to PE-1 and PE-4, as follows:

```
[/]
A:admin@PE-2# show service id 1 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
Instance    VTEP Address                          Egress VNI  EvpnStatic Num
 Mcast       Oper State                            L2 PBR     SupBcasDom MACs
-------------------------------------------------------------------------------
1           192.0.2.1                             1           evpn       1
 BUM         Up                                    No          No
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 1
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN-VXLAN Ethernet Segment Dest
===============================================================================
Instance  Eth SegId                       Num. Macs     Last Change
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

```
[/]
A:admin@PE-2# show service id 1 evpn-mpls

===============================================================================
BGP EVPN-MPLS Dest
===============================================================================
TEP Address     Egr Label     Num. MACs   Mcast           Last Change
                Transport:Tnl                             Sup BCast Domain
-------------------------------------------------------------------------------
192.0.2.4       524282        0           bum             03/16/2021 17:29:38
                ldp:65538                                 No
192.0.2.4       524283        1           none            03/16/2021 17:31:34
                ldp:65538                                 No
-------------------------------------------------------------------------------
Number of entries : 2
-------------------------------------------------------------------------------
===============================================================================
---snip---
```

- Likewise, when the two redundant PEs receive the same MAC/IP route, they will both re-advertise it with the same route key, forcing the remote PEs to pick up only one of the two (based on regular BGP selection) and create only one EVPN destination (if different from the multicast destination). In the following example, PE-6 advertised the CE-6 MAC address, that is, re-advertised by PE-4/PE-5 and then by PE-2/PE-3, but only one of the routes is selected at each hop. The following output shows that PE-1 selects the PE-2 MAC/IP route (see the "used" flag) and uses the existing EVPN destination to PE-2:

```
[/]
A:admin@PE-1# show router bgp routes evpn mac
===============================================================================
 BGP Router ID:192.0.2.1        AS:64500       Local AS:64500
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN MAC Routes
===============================================================================
Flag  Route Dist.         MacAddr           ESI
      Tag                 Mac Mobility      Label1
                          Ip Address
                          NextHop
-------------------------------------------------------------------------------
u*>i  64500:1             00:ca:fe:ca:fe:06 ESI-0
      0                   Seq:0             VNI 1
                          n/a
                          192.0.2.2

*>i   64500:1             00:ca:fe:ca:fe:06 ESI-0
      0                   Seq:0             VNI 1
                          n/a
                          192.0.2.3

-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

```
[/]
A:admin@PE-1# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
1          00:ca:fe:ca:fe:01 sap:1/2/1:1             L/30     03/16/21 17:31:10
1          00:ca:fe:ca:fe:06 vxlan-1:                Evpn     03/16/21 17:31:34
                             192.0.2.2:1
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

```
[/]
A:admin@PE-1# show service id 1 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
Instance    VTEP Address                          Egress VNI  EvpnStatic Num
 Mcast       Oper State                            L2 PBR     SupBcasDom MACs
-------------------------------------------------------------------------------
1           192.0.2.2                             1           evpn       1
 BUM         Up                                    No          No
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 1
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
BGP EVPN-VXLAN Ethernet Segment Dest
===============================================================================
Instance  Eth SegId                      Num. Macs     Last Change
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

- As shown in the preceding outputs, the EVPN destinations are always created to the IMET or MAC/IP route's BGP next-hops, which are still the system IP address of the routers (they could have also been a loopback address). The BGP next-hops need to be reachable in their respective network: DC or WAN.

The mandatory use of BGP policies in the multi-homed anycast solution

BGP policies must be configured in a multi-homed anycast solution, such as the one described in the previous section. Without policies, the following undesired behavior would happen:

- IMET routes with VXLAN encapsulation would be sent to the BGP peers in the MPLS network and IMET routes with MPLS encapsulation sent to BGP peers in the DC. The configured BGP policies will avoid that and make sure that the VXLAN routes are only sent to the DC and MPLS routes only to the WAN.
- MAC/IP routes received in the VXLAN BGP instance of a DC GW would be re-advertised to the redundant DC GW in the MPLS BGP instance and the redundant DC GW would re-advertise the same MAC again into the VXLAN instance, creating a control plane loop. The same thing would happen for MAC/IP routes received in an MPLS BGP instance. The configured BGP policies will prevent a DC GW from re-advertising MAC/IP routes received from the redundant DC GW.

While service-level BGP policies (**config>service>vpls>bgp>vsi-import/export**) may have been configured to prevent these loops and misbehavior, the use of BGP peer-level policies (**config>router>bgp>group>import/export**) is recommended due to the following reasons:

- Simplicity - BGP peer-level policies do not require any extra configuration at the service level, only at the BGP level.
- Scalability - BGP peer-level policies scale better than VSI-level policies, because the number of services where the VSI policies should be configured may be significant.

The following policies are configured in the example used in this chapter. No policies are needed in PE-1 and PE-6; only the DC GWs must be configured.

Following are the policies and how they are applied in PE-2 and PE-3:

```
# on PE-2, PE-3:
configure {
    policy-options {
        community "SOO-DCGW-23" {
            member "origin:64500:23" { }
        }
        community "mpls" {
            member "bgp-tunnel-encap:MPLS" { }
        }
        community "vxlan" {
            member "bgp-tunnel-encap:VXLAN" { }
        }

/* "drop SOO-DCGW-23" will drop any EVPN route that is received from PE-3,
the other DC GW in the pair. */

        policy-statement "drop SOO-DCGW-23" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            entry 10 {
                from {
                    family [evpn]
                    community {
                        name "SOO-DCGW-23"
                    }
                }
                action {
                    action-type reject
                }
            }
        }

/* "allow only mpls and add SOO" has a twofold objective: avoids sending EVPN-VXLAN
routes to the MPLS network and marks the advertised EVPN routes with a Site-Of-Origin extended
 community that identifies the DC GW pair. */

        policy-statement "allow only mpls and add SOO" {
            entry 10 {
                from {
                    family [evpn]
                    community {
                        name "vxlan"
                    }
                }
                action {
                    action-type reject
                }
            }
            entry 20 {
                from {
                    family [evpn]
                }
                action {
                    action-type accept
                    community {
                        add ["SOO-DCGW-23"]
                    }
                }
            }
        }

/* In the same way, "allow only vxlan and add SOO" avoids sending EVPN-MPLS routes
to the VXLAN network and marks the EVPN routes with a Site-Of-Origin extended
community that identifies the DC GW pair. */

        policy-statement "allow only vxlan and add SOO" {
            entry 10 {
                from {
                    family [evpn]
                    community {
                        name "mpls"
                    }
                }
                action {
                    action-type reject
                }
            }
            entry 20 {
                from {
                    family [evpn]
                }
                action {
                    action-type accept
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    community {
                        add ["SOO-DCGW-23"]
                    }
                }
            }
        }

/* The policies are properly applied at group level, as follows: */

# on PE-2, PE-3:
configure {
    router "Base" {
        bgp {
            ---snip---
            group "DC" {
                type internal
                import {
                    policy ["drop SOO-DCGW-23"]
                }
                export {
                    policy ["allow only vxlan and add SOO"]
                }
            }
            group "WAN" {
                type internal
                import {
                    policy ["drop SOO-DCGW-23"]
                }
                export {
                    policy ["allow only mpls and add SOO"]
                }
            }
            ---snip---
```

PE-4 and PE-5 use the same BGP peer policies, but using a Site Of Origin extended community identifying the PE-4/PE-5 pair instead of the PE-2/PE-3 pair:

```
# on PE-4, PE-5:
configure {
    policy-options {
        community "SOO-DCGW-45" {
            member "origin:64500:45" { }
        }
        community "mpls" {
            member "bgp-tunnel-encap:MPLS" { }
        }
        community "vxlan" {
            member "bgp-tunnel-encap:VXLAN" { }
        }
---snip---
```

Dual BGP instance VPLS service caveats

When two BGP instances are enabled on the same VPLS service, the following considerations apply:

- SDP-bindings are not supported (therefore, no pw-template-binding is needed in the service). Any attempt to add an SDP-binding to a service with two BGP instances will be blocked by the CLI, as follows:

```
*[ex:/configure service vpls "VPLS 1" spoke-sdp 21:1]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
A:admin@PE-2# commit
MINOR: MGMT_CORE #4001: configure service vpls "VPLS 1" - multiple bgp-evpn instances not
 supported with local mesh or spoke sdp
```

- Services that are not supported: R-VPLS, M-VPLS, I-VPLS, B-VPLS, or E-Tree VPLS
  - A consequence of not supporting R-VPLS is that no routes type 5 (IP-Prefix routes) are supported on dual BGP-instance services.
- Proxy-ARP/ND is not supported.
- BGP multi-homing is not supported.
- Although the Assisted-Replication feature is supported on dual BGP-instance VPLS services, the Assisted-Replication configuration is only relevant to the VXLAN destinations. See section EVPN route handling in dual BGP-instance VPLS services for some considerations about how EVPN handles IMET AR routes.

In addition to the preceding restrictions, some commands have a specific behavior when two BGP instances are configured:

- **config>service>vpls>bgp-evpn>routes>mac-ip>advertise** enables/disables the re-advertisement of MAC/IP routes in a BGP instance for MAC addresses that have been learned in the other BGP instance in the service.
- **config>service>vpls>bgp-evpn>routes>mac-ip>unknown-mac <boolean>** enables/disables the advertisement of the unknown MAC route (MAC 00:...:00) on the BGP-EVPN VXLAN instance. The unknown MAC route is never sent to the BGP-EVPN MPLS instance.

The use of provider tunnels on multi-homed anycast solutions

The use of provider tunnels in dual BGP-instance VPLS services connecting multiple DCs is not recommended. Figure 153: Use of provider-tunnels between anycast DC GWs create packet duplication shows the case where the same BGP-EVPN service is configured in redundant anycast DC GWs and mLDP is used in the MPLS instance. In this case, packet duplication may occur if the configuration is not done carefully.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 153: Use of provider-tunnels between anycast DC GWs create packet duplication

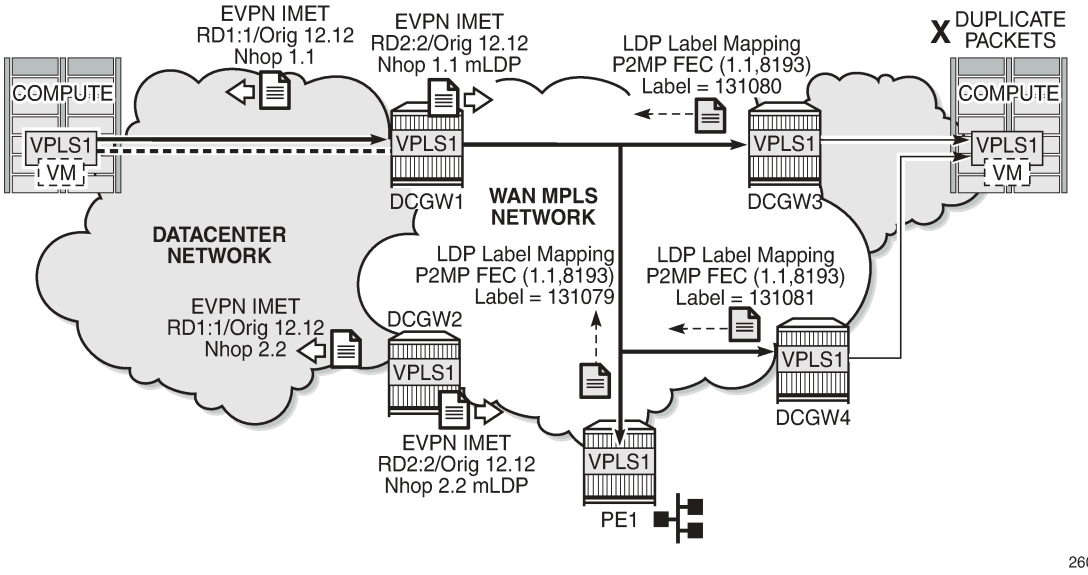

When mLDP is used along with multiple anycast multi-homing DC GWs to send BUM traffic to remote PEs, but no BUM traffic between DCs is needed, the same originating IP must be used on all the DC GWs; otherwise, packet duplication may happen. In the example in Figure 153: Use of provider-tunnels between anycast DC GWs create packet duplication, each pair of DC GWs, DCGW1/DCGW2 and DCGW3/DCGW4, is configured with a different originating IP (**config>service>vpls>bgp-evpn>incl-mcast-orig-ip**):

- DCGW3 and DCGW4 will receive the IMET route with the same route key from DCGW1 and DCGW2.
- DCGW3 and DCGW4 will select only one route, which will usually be the same; for example, the DCGW1 IMET route.
- Because of that, both DCGW3 and DCGW4 will join the mLDP tree with root in DCGW1, creating packet duplication when DCGW1 sends BUM traffic.
- Remote PE nodes with a single MPLS instance will join the mLDP tree without any issue.

To avoid the packet duplication shown by the example of Figure 153: Use of provider-tunnels between anycast DC GWs create packet duplication, the same originating IP may be configured in the four DCGWs, while the RD is still different per pair. By doing that:

- In the example of Figure 153: Use of provider-tunnels between anycast DC GWs create packet duplication, DCGW3 and DCGW4 will never join any mLDP tree sourced from DCGW1 or DCGW2. This will prevent any packet duplication because a router will ignore IMET routes received with its own originating IP, regardless of the RD.
- PE-1 (a remote EVPN-MPLS PE) will still join the mLDP trees from the two DCs.
- The preceding configuration allows the use of mLDP as long as no BUM traffic is required between the two DCs. If BUM traffic is required between DCs, IR must be used.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Troubleshooting and debugging

The following show and debug commands can be used in dual BGP-instance VPLS services:

- show router bgp routes evpn (and filters)
- show service evpn-mpls [<TEP ip-address>]
- show service vxlan [<TEP ip-address>]
- show service id bgp-evpn
- show service id evpn-mpls (and modifiers)
- show service id vxlan destinations
- debug router bgp update (in classic CLI)
- show log log-id 99

See chapter EVPN for MPLS Tunnels and EVPN for VXLAN Tunnels (Layer 2) for a detailed description of these commands.

Also, in dual BGP-instance VPLS services, the **show service id bgp <bgp-instance>** command may help see the BGP parameters of each individual BGP instance:

```
[/]
A:admin@PE-2# show service id 1 bgp ?

 bgp [<number>]

 [bgp-instance] <number>
 <number>  - <1..2>

    <number> - <1..2>
```

```
[/]
A:admin@PE-2# show service id 1 bgp 1

===============================================================================
BGP Information
===============================================================================
Vsi-Import           : None
Vsi-Export           : None
Route Dist           : 64500:1
Oper Route Dist      : 64500:1
Oper RD Type         : configured
Rte-Target Import    : None                Rte-Target Export: None
Oper RT Imp Origin   : derivedEvi          Oper RT Import   : 64500:1
Oper RT Exp Origin   : derivedEvi          Oper RT Export   : 64500:1
PW-Template Id       : None
-------------------------------------------------------------------------------
===============================================================================
```

```
[/]
A:admin@PE-2# show service id 1 bgp 2

===============================================================================
BGP Information
===============================================================================
Vsi-Import           : None
Vsi-Export           : None
Route Dist           : 64500:2
Oper Route Dist      : 64500:2
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Oper RD Type         : configured
Rte-Target Import    : None                Rte-Target Export: None
Oper RT Imp Origin   : derivedEvi          Oper RT Import   : 64500:1
Oper RT Exp Origin   : derivedEvi          Oper RT Export   : 64500:1
-------------------------------------------------------------------------------
===============================================================================
```

Conclusion

As service providers deploy EVPN-MPLS in the network for Ethernet local area network (E-LAN) and Ethernet point-to-point (E-Line) services, the use of EVPN-MPLS to interconnect data centers is becoming a popular option. Based on *draft-ietf-bess-dci-evpn-overlay*, SR OS supports the connectivity of Layer 2 EVPN-VXLAN services to an EVPN-MPLS network. To implement that EVPN-MPLS Data Center Interconnect (DCI) solution, VPLS services support dual BGP instances, where EVPN-VXLAN and EVPN-MPLS can coexist simultaneously in the same VPLS service. This chapter describes the configuration of such dual BGP-instance VPLS services and how to deploy them in a redundant anycast DC GW configuration.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

EVPN-VXLAN VPWS

This chapter provides information about EVPN-VXLAN VPWS.

Topics in this chapter include:

Applicability

This chapter was initially written for SR OS Release 16.0.R7, but the MD-CLI in the current edition is based on SR OS Release 21.5.R2.

Overview

Some service providers use VXLAN as a next-generation access technology between Multi-Service Access Node (MSAN) PE and core PE routers. VXLAN allows any IP router in the aggregation core and provides a simple alternative to MPLS. Static VXLAN bindings can be used when the MSAN PEs do not support any control plane. However, EVPN offers a control plane protocol for the VXLAN bindings for faster convergence and fault propagation. In this chapter, the focus is on EVPN-VPWS, which provides a lighter control plane compared to full-blown EVPN when point-to-point services need to be extended to the Data Center (DC).

EVPN-VXLAN VPWS is similar to EVPN-MPLS VPWS, including support of Equal Cost Multi-Path (ECMP), and EVPN All-Active (AA) and Single-Active (SA) Multi-Homing (MH). The configuration resembles the EVPN-MPLS Epipe configuration, as described in the EVPN for MPLS Tunnels in Epipe Services (EVPN-VPWS) chapter. As an example, the following configures EVPN-VXLAN Epipe 4 with SA MH.

```
# on PE-4:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "ES45" {
                        admin-state enable
                        esi 01:00:00:00:00:45:00:00:00:04
                        multi-homing-mode single-active
                        df-election {
                            es-activation-timer 3
                        }
                        association {
                            sdp 460 {
                            }
                        }
                    }
                }
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        epipe "Epipe-4" {
            admin-state enable
            service-id 4
            customer "1"
            bgp 1 {
            }
            spoke-sdp 460:4 {
            }
            vxlan {
                instance 1 {
                    vni 4
                }
            }
            bgp-evpn {
                evi 4
                local-attachment-circuit "AC-45" {
                    eth-tag 145
                }
                remote-attachment-circuit "AC-23" {
                    eth-tag 123
                }
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                    ecmp 2
                }
            }
        }
        sdp 460 {
            admin-state enable
            description "GRE SDP for SA MH"
            far-end {
                ip-address 192.0.2.6
            }
        }
```

The SDP is a GRE SDP, because no MPLS is configured in the network. The VNI is 4, and the local Attachment Circuit (AC) name is "AC-45" with Ethernet tag 145, whereas the remote AC name is "AC-23" with Ethernet tag 123. An ES can contain up to four nodes. Each of these nodes will have the same local Ethernet tag.

On Epipe services, the BGP instance is 1 and the VXLAN instance is 1. ECMP is configured with a value of 2, so the traffic flows can be sprayed over two paths with equal cost (a value greater than 2 can be configured if aliasing to more than two nodes is needed). By default, **send-tunnel-encap** is enabled, which determines whether the RFC 5512 encapsulation extended community is sent with VXLAN value (if enabled) or not sent.

EVPN-VPWS uses BGP-EVPN route type 1 (autodiscovery (AD) per-EVI routes and AD per-ES routes) and route type 4 (Ethernet Segment (ES) routes); it does not use route types 2 (MAC/IP routes), 3 (Inclusive Multicast routes), or 5 (IP Prefix routes). Figure 154: BGP-EVPN AD per-EVI route shows the fields in a BGP-EVPN AD per-EVI route.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 154: BGP-EVPN AD per-EVI route

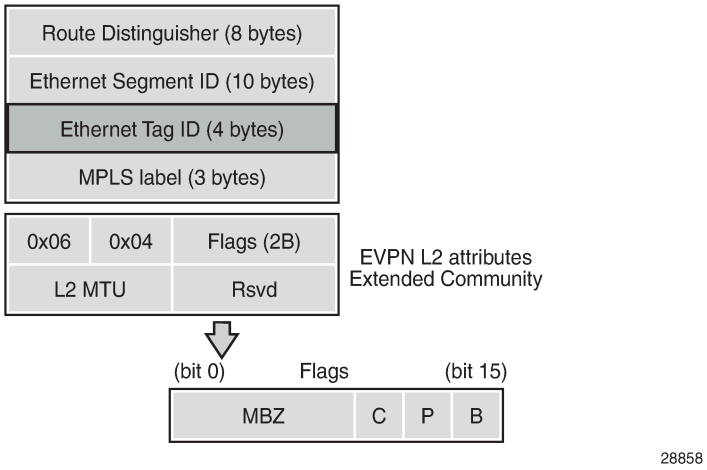

The Route Distinguisher (RD) is encoded as specified in RFC 7432; in this example, the system IP address is followed by the service ID, such as 192.0.2.2:1 for Epipe 1 on PE-2. The MPLS label field is encoded as the VXLAN Network Identifier (VNI) and the Ethernet tag field defines the local Attachment Circuit (AC) ID. The ES ID (ESI) is the 10 bytes configured ESI for MH and equals zero for single-homed services.

The EVPN L2 attributes extended community has type 0x06 (EVPN) and subtype 0x04 (EVPN L2 attributes). The flags are defined as follows:

- Flag C (control word) is set if control word is configured in the service. For EVPN-MPLS VPWS, the control word can be configured in the **bgp-evpn>mpls** context, but for EVPN-VXLAN VPWS, the control word cannot be configured in the **bgp-evpn>vxlan** context, so flag C is always zero (C=0).
- Flag P (primary) is set in MH scenarios: all nodes in an AA MH ES send P=1, but in an SA MH ES, only the Designated Forwarder (DF) sends P=1, while the NDFs send P=0. In single-homed scenarios, all nodes send P=0.
- Flag B (backup) is set in SA MH scenarios: the NDF that will take the primary role after the original primary node has failed is the backup, so it sends B=1. All other NDFs have B=0. In AA MH scenarios, all nodes send B=0. Also, in single-homed scenarios, all nodes except for the backup DF send B=0.

If the received L2 MTU does not match the configured service MTU, the EVPN binding is not set up. However, if the received L2 MTU is zero, the MTU is ignored.

AD per-EVI routes are responsible for aliasing. The following BGP update shows an AD per-EVI route received from DF 192.0.2.4 (PE-4) in an SA MH ES with ESI 01:00:00:00:00:45:00:00:00:04, Ethernet tag 145 for the local AC on PE-4, and MPLS label 4 for Epipe 4. The primary flag is set: P=1.

```
102 2021/06/30 17:30:58.043 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.4
"Peer 1: 192.0.2.4: UPDATE
Peer 1: 192.0.2.4 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 81
    Flag: 0x90 Type: 14 Len: 36 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.4
        Type: EVPN-AD Len: 25 RD: 192.0.2.4:4 ESI: 01:00:00:00:00:45:00:00:00:04,
                      tag: 145 Label: 4
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:4
        l2-attribute:MTU: 1514 C: 0 P: 1 B: 0
        bgp-tunnel-encap:VXLAN
"
```

As per RFC 8214, in an AD per-ES route, the Ethernet tag is MAX-ET (all bits are set), the MPLS label is zero, and the BGP extended community contains the single-active flag (1 for SA and 0 for AA) and ESI label. Figure 155: BGP-EVPN AD per-ES route shows the fields in a BGP-EVPN AD per-ES route.

Figure 155: BGP-EVPN AD per-ES route

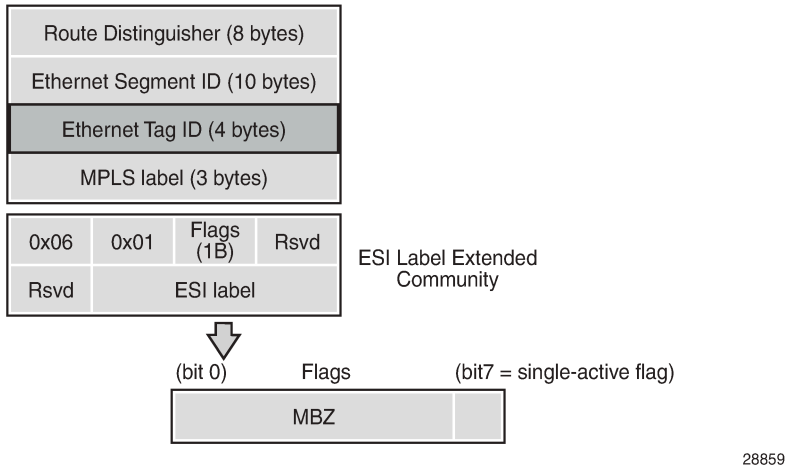

The following AD per-ES route is received by PE-2 from PE-4, which is in an SA MH ES with ESI 01:00:00:00:00:45:00:00:00:04.

```
62 2021/06/30 17:30:25.151 CEST MINOR: DEBUG #2001 Base Peer 1: 192.0.2.4
"Peer 1: 192.0.2.4: UPDATE
Peer 1: 192.0.2.4 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 73
    Flag: 0x90 Type: 14 Len: 36 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.4
        Type: EVPN-AD Len: 25 RD: 192.0.2.4:4 ESI: 01:00:00:00:00:45:00:00:00:04,
                      tag: MAX-ET Label: 0
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:4
        esi-label:524284/Single-Active
"
```

Figure 156: BGP-EVPN ES route shows a BGP-EVPN route type 4 (ES route) that is used for MH ES discovery and DF election.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 156: BGP-EVPN ES route

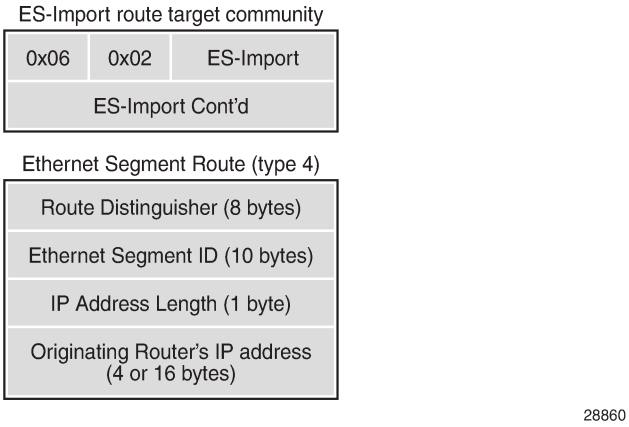

The RD is taken from the system level RD; by default, the RD is derived as system-IP:0, such as 192.0.2.4:0 for PE-4. The ESI contains the 10-byte identifier as configured in the ES. The ES import route target community has type 0x06 (EVPN) and subtype 0x02 (ES import route target), and is derived from the MAC address portion of the ESI. This extended community is treated as a route target, such as: target:00:00:00:00:45:00. Only the PEs attached to the ES will import the ES route.

The following BGP update shows a BGP-EVPN ES route sent by PE-4. The RD is defined as 192.0.2.4:0, the ESI is 01:00:00:00:00:45:00:00:00:04, and the originating IP address is 192.0.2.4 for PE-4. The ES import route target is target:00:00:00:00:45:00.

```
45 2021/06/30 17:30:55.107 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 71
    Flag: 0x90 Type: 14 Len: 34 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.4
        Type: EVPN-ETH-SEG Len: 23 RD: 192.0.2.4:0 ESI: 01:00:00:00:00:45:00:00:00:04,
                           IP-Len: 4 Orig-IP-Addr: 192.0.2.4
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        df-election::DF-Type:Auto/DP:0/DF-Preference:0/AC:1
        target:00:00:00:00:45:00
"
```

By default, the system IP addresses are used for the VXLAN tunnel termination. However, it is possible to use non-system IPv4 or IPv6 termination for EVPN-VXLAN VPWS, both for single-homed and multi-homed services. In that case, Forwarding Path Extension (FPE) needs to be defined with VXLAN termination, as described in chapter Static VXLAN Termination in Epipe Services.

The following shows the configuration of the single-homed Epipe 2 using non-system IPv4 source VXLAN Tunnel Endpoint (VTEP) 10.0.3.1 on PE-3. Likewise, it is possible to use a non-system IPv6 source VTEP, such as **vxlan>source-vtep 2001::3:1**. Unlike the source VTEP, the egress VTEP cannot be configured when BGP-EVPN is enabled. The egress VTEP is dynamically learned via BGP instead.

```
# on PE-3:
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
configure {
    service {
        epipe "Epipe-2" {
            admin-state enable
            service-id 2
            customer "1"
            bgp 1 {
            }
            sap 1/1/1:2 {
            }
            vxlan {
                source-vtep 10.0.3.1
                instance 1 {
                    vni 2
                }
            }
            bgp-evpn {
                evi 2
                local-attachment-circuit "AC-3" {
                    eth-tag 103
                }
                remote-attachment-circuit "AC-5" {
                    eth-tag 105
                }
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
        }
```

Configuration

The following use cases are included in the configuration section:

- Single-homed EVPN-VXLAN Epipe using IPv4 system addresses
- Single-homed EVPN-VXLAN Epipe using non-system IPv4 addresses
- Single-homed EVPN-VXLAN Epipe using non-system IPv6 addresses
- AA and SA multi-homed EVPN-VXLAN Epipe using IPv4 system addresses
- AA and SA multi-homed EVPN-VXLAN Epipe using non-system IPv4 addresses
- AA and SA multi-homed EVPN-VXLAN Epipe using non-system IPv6 addresses

Figure 157: Example topology shows the example topology with six PEs. EVPN-VXLAN Epipe services will be configured on the core PEs PE-2, PE-3, PE-4, and PE-5. On the access nodes PE-1 and PE-6, ordinary Epipe services will be configured, without EVPN-VXLAN. The CEs are emulated by VPRN services configured on PE-1 or PE-6.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 157: Example topology

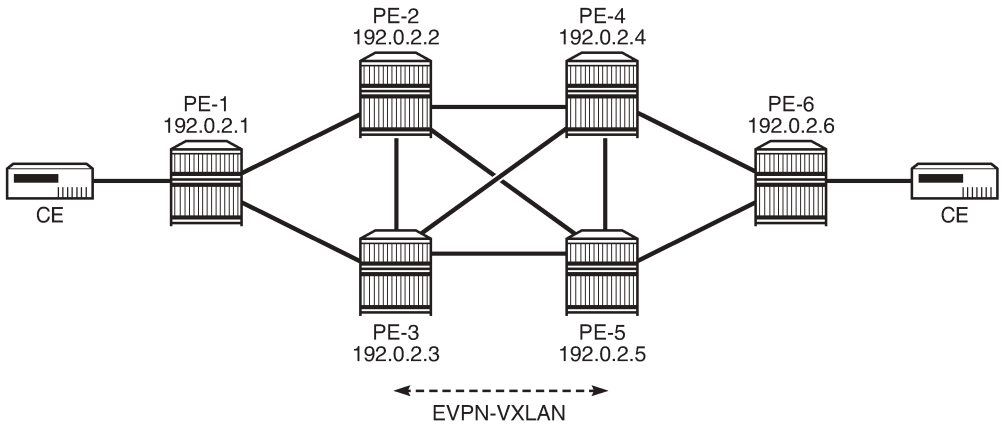

The initial configuration includes:

- Cards, MDAs, ports
- Router interfaces
- IS-IS on all router interfaces: level 2 between the core PEs and level 1 in the access networks

No MPLS protocol is configured.

BGP is configured on the core PEs for the EVPN address family with RR PE-2. The BGP configuration on RR PE-2 is as follows:

```
# on PE-2:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-update {
                evpn true
            }
            group "iBGP" {
                type internal
                split-horizon true
                family {
                    evpn true
                }
                cluster {
                    cluster-id 192.0.2.2
                }
            }
            neighbor "192.0.2.3" {
                group "iBGP"
            }
            neighbor "192.0.2.4" {
                group "iBGP"
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            neighbor "192.0.2.5" {
                group "iBGP"
            }
```

Single-homed EVPN-VXLAN Epipe using system IPv4 addresses

Figure 158: Single-homed EVPN-VXLAN Epipe 1 using system IP addresses shows the routers PE-1, PE-2, PE-4, and PE-6 configured with Epipe 1. VXLAN-EVPN is only configured on the core PEs PE-2 and PE-4.

Figure 158: Single-homed EVPN-VXLAN Epipe 1 using system IP addresses

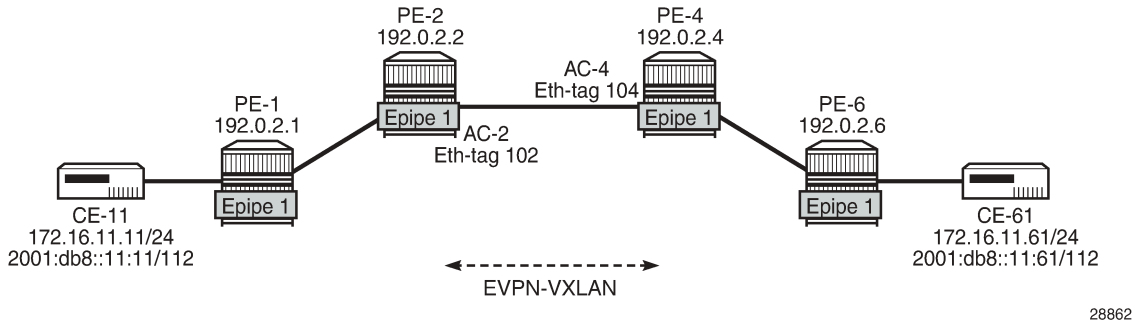

Configuration of Epipe 1

On PE-1, Epipe 1 is configured without EVPN-VXLAN, as follows.

```
# on PE-1:
configure {
    service {
        epipe "Epipe 1" {
            admin-state enable
            service-id 1
            customer "1"
            sap 1/1/1:1 {
            }
            sap 1/2/1:1 {
            }
        }
```

On PE-2, Epipe 1 is configured with EVPN-VXLAN. The local AC "AC-2" has Ethernet tag 102 and the remote AC is "AC-4" with Ethernet tag 104, as follows:

```
# on PE-2:
configure {
    service {
        epipe "Epipe-1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
            }
            sap 1/1/2:1 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
            vxlan {
                instance 1 {
                    vni 1
                }
            }
            bgp-evpn {
                evi 1
                local-attachment-circuit "AC-2" {
                    eth-tag 102
                }
                remote-attachment-circuit "AC-4" {
                    eth-tag 104
                }
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
        }
```

The Epipe configuration on PE-4 is similar, but the local AC and remote AC are swapped, as follows. Instead of a SAP, a spoke-SDP is configured toward PE-6. The SDP itself is GRE-based.

```
# on PE-4:
configure {
    service {
        epipe "Epipe-1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
            }
            spoke-sdp 46:1 {
            }
            vxlan {
                instance 1 {
                    vni 1
                }
            }
            bgp-evpn {
                evi 1
                local-attachment-circuit "AC-4" {
                    eth-tag 104
                }
                remote-attachment-circuit "AC-2" {
                    eth-tag 102
                }
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
        }
        sdp 46 {
            admin-state enable
            description "GRE SDP for single-homing"
            far-end {
                ip-address 192.0.2.6
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

On PE-6, Epipe 1 is an ordinary Epipe with spoke-SDP 64:1 toward PE-4 and SAP 1/2/1:1 toward a CE, as follows:

```
# on PE-6:
configure {
    service {
        epipe "Epipe 1" {
            admin-state enable
            service-id 1
            customer "1"
            spoke-sdp 64:1 {
            }
            sap 1/2/1:1 {
            }
        }
        sdp 64 {
            admin-state enable
            description "GRE SDP for single-homing"
            far-end {
                ip-address 192.0.2.4
            }
        }
```

Verification

VPRN 11 on PE-1 and PE-6 simulates the CEs CE-11 and CE-61. The connectivity between the CEs can be verified as follows:

```
[/]
A:admin@PE-1# ping 172.16.11.61 router-instance "VPRN 11" interval 0.1
                                                              output-format summary
PING 172.16.11.61 56 data bytes
!!!!!
---- 172.16.11.61 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 4.42ms, avg = 5.89ms, max = 10.9ms, stddev = 2.49ms
```

```
[/]
A:admin@PE-1# ping 2001:db8::11:61 router-instance "VPRN 11" interval 0.1
                                                              output-format summary
PING 2001:db8::11:61 56 data bytes
!!!!!
---- 2001:db8::11:61 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 4.14ms, avg = 5.63ms, max = 11.2ms, stddev = 2.79ms
```

On PE-2, the VXLAN destination for Epipe 1 is the system address of PE-4: 192.0.2.4, as follows. There are no VXLAN ES destinations for Epipe 1, because the service is single-homed.

```
[/]
A:admin@PE-2# show service id 1 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address                          Egress VNI          Oper      Vxlan
                                                          State     Type
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
192.0.2.4                               1                   Up        evpn
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 1
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN VXLAN ES Dest
===============================================================================
I Eth Seg Id                    TEP Address       VNI       Last Changed
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

The following BGP-EVPN information for Epipe 1 on PE-2 includes the EVI and the AC names and Ethernet tags. For Epipes, the BGP instance ID and VXLAN instance ID always equal 1.

```
[/]
A:admin@PE-2# show service id 1 bgp-evpn

===============================================================================
BGP EVPN Table
===============================================================================
EVI                : 1                  Creation Origin    : manual
Local AC Name      : AC-2
Eth Tag            : 102
Endpoint           : (Not Specified)
Ingress Label      : 0
Remote AC Name     : AC-4
Eth Tag            : 104
Endpoint           : (Not Specified)

===============================================================================
BGP EVPN VXLAN Information
===============================================================================
Admin Status       : Enabled            Bgp Instance       : 1
Vxlan Instance     : 1
Max Ecmp Routes    : 1
Default Route Tag  : none
Send EVPN Encap    : Enabled
===============================================================================
```

PE-2 has received the following BGP-EVPN AD per-EVI route with RD 192.0.2.4:1 and Ethernet tag 104 from PE-4. Epipe 1 is single-homed, so ESI=0 and there is no primary or backup node (P=B=0). Also, no control word is used, so C=0.

```
5 2021/06/30 17:14:44.605 CEST MINOR: DEBUG #2001 Base Peer 1: 192.0.2.4
"Peer 1: 192.0.2.4: UPDATE
Peer 1: 192.0.2.4 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 81
    Flag: 0x90 Type: 14 Len: 36 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.4
        Type: EVPN-AD Len: 25 RD: 192.0.2.4:1 ESI: ESI-0, tag: 104 Label: 1
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:1
        l2-attribute:MTU: 1514 C: 0 P: 0 B: 0
        bgp-tunnel-encap:VXLAN
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
"
```

The following shows the received BGP-EVPN AD per-EVI routes with RD 192.0.2.4:1 on PE-2.

```
[/]
A:admin@PE-2# show router bgp routes evpn auto-disc rd 192.0.2.4:1
===============================================================================
 BGP Router ID:192.0.2.2        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Auto-Disc Routes
===============================================================================
Flag  Route Dist.         ESI                           NextHop
      Tag                                               Label
-------------------------------------------------------------------------------
u*>i  192.0.2.4:1         ESI-0                         192.0.2.4
      104                                               VNI 1

-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

Single-homed EVPN-VXLAN Epipe using non-system IPv4 addresses

Figure 159: Single-homed EVPN-VXLAN Epipe 2 using non-system IP addresses shows the single-homed service Epipe 2 configured on PE-1, PE-3, PE-5, and PE-6. On PE-3, a loopback interface is created in the base router with IPv4 address 10.0.3.3/28. Epipe 2 uses VXLAN termination 10.0.3.1 from the same subnet.

Figure 159: Single-homed EVPN-VXLAN Epipe 2 using non-system IP addresses

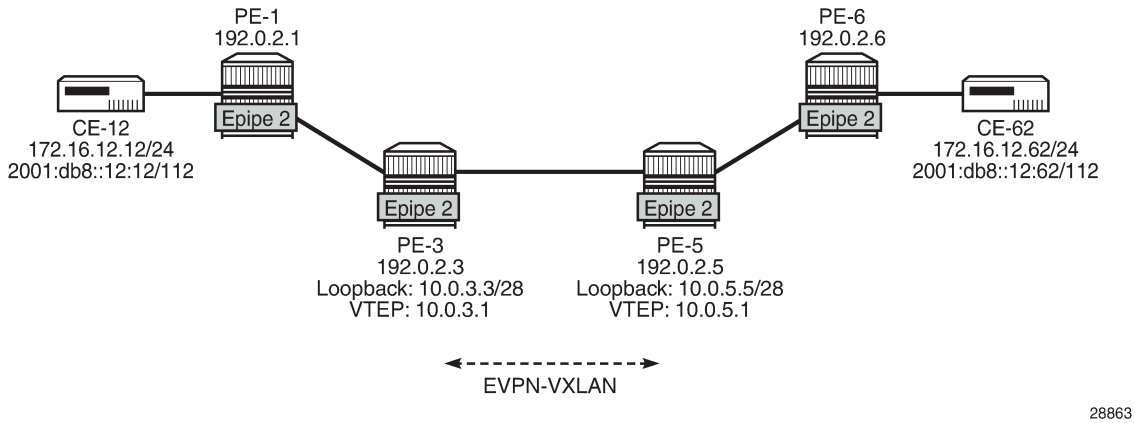

Configuration of Epipe 2

On PE-1 and PE-6, the configuration of Epipe 2 is similar to the configuration of Epipe 1.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

On PE-3, FPE needs to be configured using PXC, as described in chapter Static VXLAN Termination in Epipe Services. The following configuration is included without further explanation about FPE or PXC. The same configuration is required on PE-5.

```
# on PE-3:
configure {
    fwd-path-ext {
        sdp-id-range {
            start 10000
            end 10127
        }
        fpe 1 {
            path {
                pxc 1
            }
            application {
                vxlan-termination {
                }
            }
        }
    }
    port 1/2/5 {
        admin-state enable
        ethernet {
            mode hybrid
            encap-type dot1q
            dot1x {
                tunneling true
            }
        }
    }
    port pxc-1.a {
        admin-state enable
    }
    port pxc-1.b {
        admin-state enable
    }
    port-xc {
        pxc 1 {
            admin-state enable
            port 1/2/5
        }
    }
```

On PE-3, the following loopback interface is created and IS-IS is enabled on it. The subnet must allow multiple IP addresses; one other IP address from the subnet will be defined as VXLAN tunnel termination. The IPv6 address is only required in the next use-case, but this configuration will not be repeated in that section.

```
# on PE-3:
configure {
    router "Base" {
        interface "lo1" {
            loopback
            ipv4 {
                primary {
                    address 10.0.3.3
                    prefix-length 28
                }
            }
            ipv6 {
                address 2001::3:3 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    prefix-length 124
                }
            }
        }
        isis 0 {
            interface "lo1" {
                passive true
            }
        }
```

Up to three VXLAN tunnel terminations can be defined per system. On PE-3, the following two VXLAN tunnel terminations are configured. For Epipe 2, only the first VXLAN tunnel termination is required; the second (IPv6) VXLAN tunnel termination is used in Epipe 3. The VXLAN tunnel termination is used as VXLAN source VTEP in Epipe 2. No egress VTEP can be defined when BGP-EVPN is configured in the service; egress VTEPs are configured in static VXLAN tunnels instead.

```
# on PE-3:
configure {
    service {
        system {
            vxlan {
                tunnel-termination 10.0.3.1 {
                    fpe-id 1
                }
                tunnel-termination 2001::3:1 {
                    fpe-id 1
                }
            }
        }
        epipe "Epipe-2" {
            admin-state enable
            service-id 2
            customer "1"
            bgp 1 {
            }
            sap 1/1/1:2 {
            }
            vxlan {
                source-vtep 10.0.3.1
                instance 1 {
                    vni 2
                }
            }
            bgp-evpn {
                evi 2
                local-attachment-circuit "AC-3" {
                    eth-tag 103
                }
                remote-attachment-circuit "AC-5" {
                    eth-tag 105
                }
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
        }
```

The configuration on PE-5 is similar. The following is the service configuration on PE-5.

```
# on PE-5:
configure {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    service {
        system {
            vxlan {
                tunnel-termination 10.0.5.1 {
                    fpe-id 1
                }
                tunnel-termination 2001::5:1 {
                    fpe-id 1
                }
            }
        }
        epipe "Epipe-2" {
            admin-state enable
            service-id 2
            customer "1"
            bgp 1 {
            }
            spoke-sdp 56:2 {
            }
            vxlan {
                source-vtep 10.0.5.1
                instance 1 {
                    vni 2
                }
            }
            bgp-evpn {
                evi 2
                local-attachment-circuit "AC-5" {
                    eth-tag 105
                }
                remote-attachment-circuit "AC-3" {
                    eth-tag 103
                }
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
        }
        sdp 56 {
            admin-state enable
            description "GRE SDP for single-homing"
            far-end {
                ip-address 192.0.2.6
            }
        }
```

It is possible to use a system IPv4 address as a VXLAN tunnel termination on one of the nodes and a non-system IPv4 address on another, but that is not configured here.

Verification

The connectivity between the CEs that are emulated by VPRN 12 can be verified as follows:

```
[/]
A:admin@PE-1# ping 172.16.12.62 router-instance "VPRN 12" interval 0.1
                                                        output-format summary
PING 172.16.12.62 56 data bytes
!!!!!
---- 172.16.12.62 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
round-trip min = 4.31ms, avg = 4.67ms, max = 5.13ms, stddev = 0.268ms
```

```
[/]
A:admin@PE-1# ping 2001:db8::12:62 router-instance "VPRN 12" interval 0.1
                                                              output-format summary
PING 2001:db8::12:62 56 data bytes
!!!!!
---- 2001:db8::12:62 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 4.39ms, avg = 4.84ms, max = 5.20ms, stddev = 0.276ms
```

On PE-3, the VXLAN destination for Epipe 2 is the non-system address 10.0.5.1 on PE-5, as follows:

```
[/]
A:admin@PE-3# show service id 2 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address                          Egress VNI          Oper      Vxlan
                                                          State     Type
-------------------------------------------------------------------------------
10.0.5.1                              2                   Up        evpn
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 1
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN VXLAN ES Dest
===============================================================================
I Eth Seg Id                    TEP Address       VNI       Last Changed
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

The following BGP-EVPN information for Epipe 2 on PE-3 includes the EVI, AC names, and Ethernet tags.

```
[/]
A:admin@PE-3# show service id 2 bgp-evpn

===============================================================================
BGP EVPN Table
===============================================================================
EVI                 : 2                 Creation Origin    : manual
Local AC Name       : AC-3
Eth Tag             : 103
Endpoint            : (Not Specified)
Ingress Label       : 0
Remote AC Name      : AC-5
Eth Tag             : 105
Endpoint            : (Not Specified)

===============================================================================
BGP EVPN VXLAN Information
===============================================================================
Admin Status        : Enabled           Bgp Instance       : 1
Vxlan Instance      : 1
Max Ecmp Routes     : 1
Default Route Tag   : none
Send EVPN Encap     : Enabled
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
```

PE-3 received the following BGP-EVPN AD per-EVI route with RD 192.0.2.5:2 from PE-5. The Ethernet tag is 105 and the next-hop is the non-system address 10.0.5.1. ESI=0 for single-homed services.

```
[/]
A:admin@PE-3# show router bgp routes evpn auto-disc rd 192.0.2.5:2
===============================================================================
 BGP Router ID:192.0.2.3        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Auto-Disc Routes
===============================================================================
Flag  Route Dist.          ESI                           NextHop
      Tag                                                Label
-------------------------------------------------------------------------------
u*>i  192.0.2.5:2          ESI-0                         10.0.5.1
      105                                                VNI 2

-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

Single-homed EVPN-VXLAN Epipe using non-system IPv6 addresses

Figure 160: Single-homed EVPN-VXLAN Epipe 3 using non-system IPv6 addresses shows the example topology for single-homed EVPN-VXLAN Epipe 3 using non-system IPv6 addresses for VXLAN tunnel termination.

Figure 160: Single-homed EVPN-VXLAN Epipe 3 using non-system IPv6 addresses

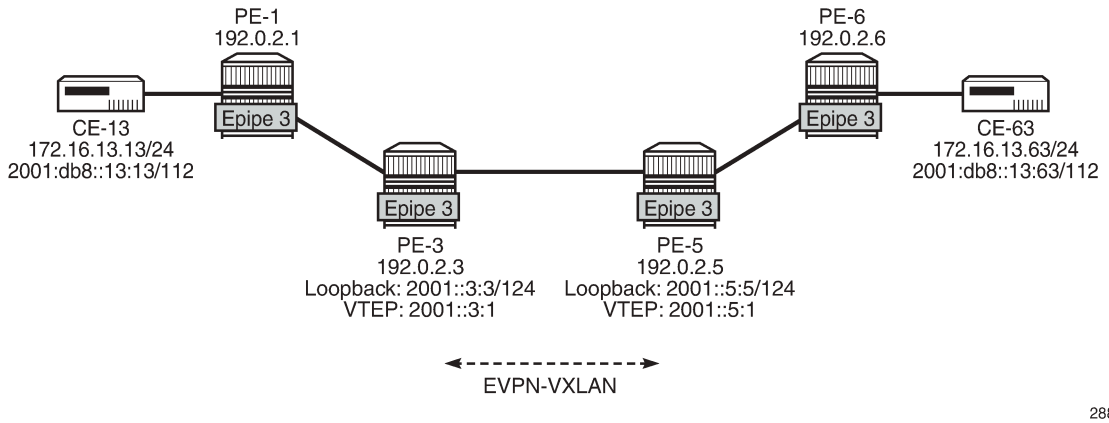

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Configuration of Epipe 3

The following single-homed Epipe 3 using non-system IPv6 addresses for the VXLAN tunnel terminations is configured on PE-3.

```
# on PE-3:
configure {
    service {
        system {
            vxlan {
                tunnel-termination 10.0.3.1 {
                    fpe-id 1
                }
                tunnel-termination 2001::3:1 {
                    fpe-id 1
                }
            }
        }
        epipe "Epipe-3" {
            admin-state enable
            service-id 3
            customer "1"
            bgp 1 {
            }
            sap 1/1/1:3 {
            }
            vxlan {
                source-vtep 2001::3:1
                instance 1 {
                    vni 3
                }
            }
            bgp-evpn {
                evi 3
                local-attachment-circuit "AC-3_v6" {
                    eth-tag 163
                }
                remote-attachment-circuit "AC-5_v6" {
                    eth-tag 165
                }
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
        }
```

The service configuration on PE-5 is similar, as follows:

```
# on PE-5:
configure {
    service {
        system {
            vxlan {
                tunnel-termination 10.0.5.1 {
                    fpe-id 1
                }
                tunnel-termination 2001::5:1 {
                    fpe-id 1
                }
            }
        }
        epipe "Epipe-3" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            admin-state enable
            service-id 3
            customer "1"
            bgp 1 {
            }
            spoke-sdp 56:3 {
            }
            vxlan {
                source-vtep 2001::5:1
                instance 1 {
                    vni 3
                }
            }
            bgp-evpn {
                evi 3
                local-attachment-circuit "AC-5_v6" {
                    eth-tag 165
                }
                remote-attachment-circuit "AC-3_v6" {
                    eth-tag 163
                }
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
        }
        sdp 56 {
            admin-state enable
            description "GRE SDP for single-homing"
            far-end {
                ip-address 192.0.2.6
            }
        }
```

Verification

The connectivity between the CEs that are emulated by VPRN 13 is verified as follows:

```
[/]
A:admin@PE-1# ping 172.16.13.63 router-instance "VPRN 13" interval 0.1
                                                              output-format summary
PING 172.16.13.63 56 data bytes
!!!!!
---- 172.16.13.63 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 4.37ms, avg = 4.68ms, max = 5.41ms, stddev = 0.377ms
```

```
[/]
A:admin@PE-1# ping 2001:db8::13:63 router-instance "VPRN 13" interval 0.1
                                                             output-format summary
PING 2001:db8::13:63 56 data bytes
!!!!!
---- 2001:db8::13:63 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 3.87ms, avg = 4.54ms, max = 5.88ms, stddev = 0.702ms
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

On PE-3, the VXLAN destination for Epipe 3 is the non-system IPv6 address 2001::5:1 on PE-5, as follows:

```
[/]
A:admin@PE-3# show service id 3 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address                          Egress VNI          Oper      Vxlan
                                                          State     Type
-------------------------------------------------------------------------------
2001::5:1                             3                   Up        evpn
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 1
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN VXLAN ES Dest
===============================================================================
I Eth Seg Id                    TEP Address       VNI       Last Changed
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

The following BGP-EVPN information for Epipe 3 on PE-3 includes the EVI and the AC names and Ethernet tags.

```
[/]
A:admin@PE-3# show service id 3 bgp-evpn

===============================================================================
BGP EVPN Table
===============================================================================
EVI                 : 3                   Creation Origin    : manual
Local AC Name       : AC-3_v6
Eth Tag             : 163
Endpoint            : (Not Specified)
Ingress Label       : 0
Remote AC Name      : AC-5_v6
Eth Tag             : 165
Endpoint            : (Not Specified)

===============================================================================
BGP EVPN VXLAN Information
===============================================================================
Admin Status        : Enabled             Bgp Instance       : 1
Vxlan Instance      : 1
Max Ecmp Routes     : 1
Default Route Tag   : none
Send EVPN Encap     : Enabled
===============================================================================
```

PE-3 received the following BGP-EVPN AD per-EVI route with RD 192.0.2.5:3 and next-hop 2001::5:1.

```
[/]
A:admin@PE-3# show router bgp routes evpn auto-disc rd 192.0.2.5:3
===============================================================================
 BGP Router ID:192.0.2.3        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Auto-Disc Routes
===============================================================================
Flag  Route Dist.         ESI                           NextHop
      Tag                                               Label
-------------------------------------------------------------------------------
u*>i  192.0.2.5:3         ESI-0                         2001::5:1
      165                                               VNI 3
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

AA and SA multi-homed EVPN-VXLAN Epipe using system IPv4 addresses

Figure 161: EVPN-VXLAN Epipe 4 with AA MH and SA MH using system IPv4 addresses shows the example topology for EVPN-VXLAN Epipe 4 with AA MH ES "vES23_4" between PE-2 and PE-3 and SA MH ES "ES45" between PE-4 and PE-5.

Figure 161: EVPN-VXLAN Epipe 4 with AA MH and SA MH using system IPv4 addresses

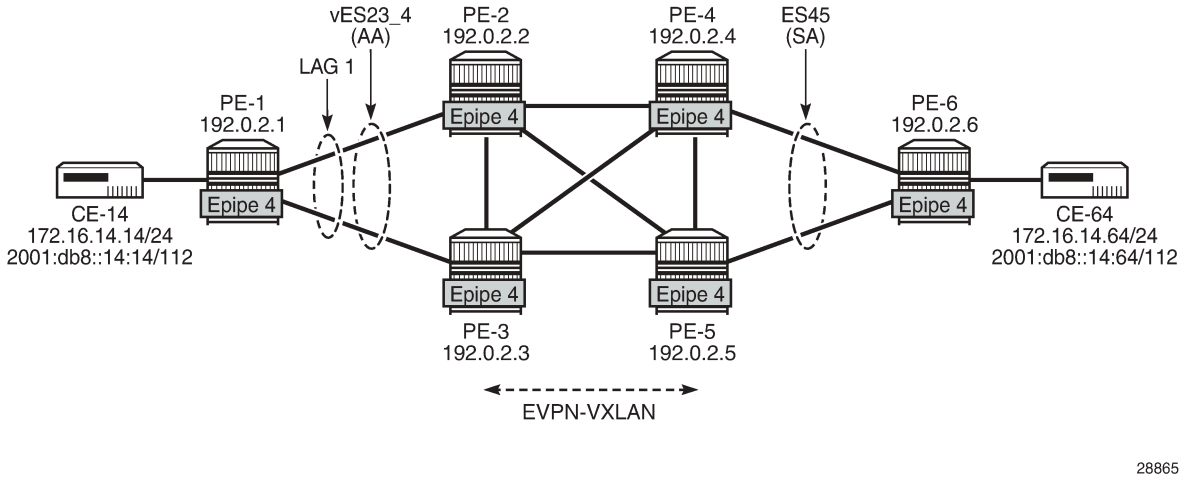

Configuration of Epipe 4

On PE-1, Epipe 4 is configured as follows:

```
# on PE-1:
configure {
    service {
        epipe "Epipe-4" {
            admin-state enable
            service-id 4
            customer "1"
            sap 1/2/1:4 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
            sap lag-1:4 {
            }
        }
```

On PE-2 and PE-3, the AA MH ES "vES23_4" is configured as a virtual ES for LAG 1 and dot1q-tag 4, so it only affects Epipe 4.

```
# on PE-2, PE-3:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "vES23_4" {
                        admin-state enable
                        type virtual
                        esi 01:00:00:00:00:23:00:00:00:04
                        multi-homing-mode all-active
                        df-election {
                            es-activation-timer 3
                        }
                        association {
                            lag "lag-1" {
                                virtual-ranges {
                                    dot1q {
                                        q-tag 4 {
                                            end 4
                                        }
                                    }
                                }
                            }
                        }
                    }
```

On PE-2 and PE-3, Epipe 4 is configured as follows. The system IPv4 address is used as VXLAN termination, the local AC Ethernet tag is 123, and the remote AC Ethernet tag is 145.

```
# on PE-2, PE-3:
configure {
    service {
        epipe "Epipe-4" {
            admin-state enable
            service-id 4
            customer "1"
            bgp 1 {
            }
            sap lag-1:4 {
            }
            vxlan {
                instance 1 {
                    vni 4
                }
            }
            bgp-evpn {
                evi 4
                local-attachment-circuit "AC-23" {
                    eth-tag 123
                }
                remote-attachment-circuit "AC-45" {
                    eth-tag 145
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                    ecmp 2
                }
            }
        }
```

On PE-4 and PE-5, the SA MH ES "ES45" is configured with a GRE SDP toward PE-6: SDP 460 on PE-4 and SDP 560 on PE-6. The following is the configuration of "ES45" on PE-4:

```
# on PE-4:
configure {
    service {
        sdp 460 {
            admin-state enable
            description "GRE SDP for SA MH"
            far-end {
                ip-address 192.0.2.6
            }
        }
        system {
            bgp {
                evpn {
                    ethernet-segment "ES45" {
                        admin-state enable
                        esi 01:00:00:00:00:45:00:00:00:04
                        multi-homing-mode single-active
                        df-election {
                            es-activation-timer 3
                        }
                        association {
                            sdp 460 {
                            }
                        }
                    }
                }
            }
        }
```

On PE-4, Epipe 4 is configured as follows. The configuration on PE-5 is similar, but with spoke-SDP 560:4 instead.

```
# on PE-4:
configure {
    service {
        epipe "Epipe-4" {
            admin-state enable
            service-id 4
            customer "1"
            bgp 1 {
            }
            spoke-sdp 460:4 {
            }
            vxlan {
                instance 1 {
                    vni 4
                }
            }
            bgp-evpn {
                evi 4
                local-attachment-circuit "AC-45" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    eth-tag 145
                }
                remote-attachment-circuit "AC-23" {
                    eth-tag 123
                }
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                    ecmp 2
                }
            }
        }
```

On PE-6, Epipe 4 is configured as follows:

```
# on PE-6:
configure {
    service {
        epipe "Epipe-4" {
            admin-state enable
            description "Epipe-4 with active/standby pseudowire"
            service-id 4
            customer "1"
            endpoint "EP" {
            }
            spoke-sdp 640:4 {
                endpoint {
                    name "EP"
                }
            }
            spoke-sdp 650:4 {
                endpoint {
                    name "EP"
                }
            }
            sap 1/2/1:4 {
            }
        }
```

Verification

The connectivity between the CEs emulated by VPRN 14 can be verified as follows:

```
[/]
A:admin@PE-1# ping 172.16.14.64 router-instance "VPRN 14" interval 0.1
                                                              output-format summary
PING 172.16.14.64 56 data bytes
!!!!!
---- 172.16.14.64 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 4.17ms, avg = 5.79ms, max = 11.5ms, stddev = 2.88ms
```

```
[/]
A:admin@PE-1# ping 2001:db8::14:64 router-instance "VPRN 14" interval 0.1
                                                              output-format summary
PING 2001:db8::14:64 56 data bytes
!!!!!
---- 2001:db8::14:64 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
round-trip min = 4.29ms, avg = 6.41ms, max = 14.4ms, stddev = 4.00ms
```

The following BGP-EVPN information for Epipe 4 includes the EVI and the AC names and Ethernet tags:

```
[/]
A:admin@PE-2# show service id 4 bgp-evpn

===============================================================================
BGP EVPN Table
===============================================================================
EVI                 : 4                  Creation Origin    : manual
Local AC Name       : AC-23
Eth Tag             : 123
Endpoint            : (Not Specified)
Ingress Label       : 0
Remote AC Name      : AC-45
Eth Tag             : 145
Endpoint            : (Not Specified)

===============================================================================
BGP EVPN VXLAN Information
===============================================================================
Admin Status        : Enabled            Bgp Instance       : 1
Vxlan Instance      : 1
Max Ecmp Routes     : 2
Default Route Tag   : none
Send EVPN Encap     : Enabled
===============================================================================
```

PE-4 received the following BGP-EVPN ES route with ESI 01:00:00:00:00:45:00:00:00:04 from PE-5:

```
[/]
A:admin@PE-4# show router bgp routes evpn eth-seg
===============================================================================
 BGP Router ID:192.0.2.4        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Eth-Seg Routes
===============================================================================
Flag  Route Dist.         ESI                           NextHop
      OrigAddr
-------------------------------------------------------------------------------
u*>i  192.0.2.5:0         01:00:00:00:00:45:00:00:00:04 192.0.2.5
      192.0.2.5

-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

Furthermore, PE-4 received the following AD per-EVI (with Ethernet tag 123 or 145) and AD per-ES (MAX-ET) routes for Epipe 4 from its three BGP peers. The ESI is non-zero for multi-homed services.

```
[/]
A:admin@PE-4# show router bgp routes evpn auto-disc
===============================================================================
 BGP Router ID:192.0.2.4        AS:64500       Local AS:64500
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Auto-Disc Routes
===============================================================================
Flag  Route Dist.         ESI                           NextHop
      Tag                                               Label
-------------------------------------------------------------------------------
---snip---

u*>i  192.0.2.2:4         01:00:00:00:00:23:00:00:00:04 192.0.2.2
      123                                               VNI 4

u*>i  192.0.2.2:4         01:00:00:00:00:23:00:00:00:04 192.0.2.2
      MAX-ET                                            LABEL 0

u*>i  192.0.2.3:4         01:00:00:00:00:23:00:00:00:04 192.0.2.3
      123                                               VNI 4

u*>i  192.0.2.3:4         01:00:00:00:00:23:00:00:00:04 192.0.2.3
      MAX-ET                                            LABEL 0

u*>i  192.0.2.5:4         01:00:00:00:00:45:00:00:00:04 192.0.2.5
      145                                               VNI 4

u*>i  192.0.2.5:4         01:00:00:00:00:45:00:00:00:04 192.0.2.5
      MAX-ET                                            LABEL 0

-------------------------------------------------------------------------------
```

In AA MH ESs, the DF for VPLS services is the forwarder for Broadcast, Unknown unicast, and Multicast (BUM) traffic. In Epipes, however, all traffic is treated as unicast. The following tools commands on PE-2 and PE-3 show that DF is not applicable for AA MH ES "vES23_4.

```
[/]
A:admin@PE-2# tools dump service system bgp-evpn ethernet-segment "vES23_4" evi 4 df

[06/30/2021 17:33:54] All Active VPWS - DF N/A
```

```
[/]
A:admin@PE-3# tools dump service system bgp-evpn ethernet-segment "vES23_4" evi 4 df

[06/30/2021 17:33:56] All Active VPWS - DF N/A
```

The following command on PE-2 shows no DF candidates for ES "vES23_4", even though PE-2 (as well as PE-3) is considered as DF (DF=yes):

```
[/]
A:admin@PE-2# show service system bgp-evpn ethernet-segment name "vES23_4" evi evi-1 4

===============================================================================
EVI DF and Candidate List
===============================================================================
EVI           SvcId         Actv Timer Rem      DF  DF Last Change
-------------------------------------------------------------------------------
4             4             0                   yes 06/30/2021 17:28:14
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
DF Candidates                           Time Added
-------------------------------------------------------------------------------
No entries found
===============================================================================
```

In the SA MH ES "ES45", PE-4 is DF out of a list of two candidates, as follows:

```
[/]
A:admin@PE-4# show service system bgp-evpn ethernet-segment name "ES45" evi evi-1 4

===============================================================================
EVI DF and Candidate List
===============================================================================
EVI           SvcId         Actv Timer Rem      DF  DF Last Change
-------------------------------------------------------------------------------
4             4             0                   yes 06/30/2021 17:30:25
===============================================================================

===============================================================================
DF Candidates                           Time Added
-------------------------------------------------------------------------------
192.0.2.4                               06/30/2021 17:30:55
192.0.2.5                               06/30/2021 17:30:55
-------------------------------------------------------------------------------
Number of entries: 2
===============================================================================
```

On NDF PE-5, the spoke-SDP is operationally down with flag StandbyForMHProtocol, as follows:

```
[/]
A:admin@PE-5# show service id 4 sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
560:4            Spok     192.0.2.6       Up      Down      524282    524281
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
```

```
[/]
A:admin@PE-5# show service id 4 sdp detail | match "Flags"
Flags              : StandbyForMHProtocol
```

The following command on PE-2 shows that the VXLAN destination for Epipe 4 is the ES "ES45" with ESI 01:00:00:00:00:45:00:00:00:04 and TEP address 192.0.2.4, which is the system IP address of the DF.

```
[/]
A:admin@PE-2# show service id 4 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address                            Egress VNI          Oper      Vxlan
                                                            State     Type
-------------------------------------------------------------------------------
No Matching Entries
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================

===============================================================================
BGP EVPN VXLAN ES Dest
===============================================================================
I Eth Seg Id                      TEP Address       VNI       Last Changed
-------------------------------------------------------------------------------
1 01:00:00:00:00:45:00:00:00:04 192.0.2.4         4         06/30/2021 17:30:58
-------------------------------------------------------------------------------
===============================================================================
```

On PE-2, the following command shows that BGP-EVPN AD per-EVI routes with Ethernet tag 145 from PE-4 (RD 192.0.2.4:4) are sent with primary flag P=1 and AD per-EVI routes with Ethernet tag 145 from PE-5 (RD 192.0.2.5:4) are sent with primary flag P=0 and backup flag B=1.

```
[/]
A:admin@PE-3# show router bgp routes evpn auto-disc tag 145 detail
                                                          | match 'C:|Route Dist'
Community      : target:64500:4 l2-attribute:MTU: 1514 C: 0 P: 1 B: 0
Route Dist.    : 192.0.2.4:4
---snip---
Community      : target:64500:4 l2-attribute:MTU: 1514 C: 0 P: 0 B: 1
Route Dist.    : 192.0.2.5:4
---snip---
```

AA and SA multi-homed EVPN-VXLAN Epipe using non-system IPv4 addresses

Figure 162: EVPN-VXLAN Epipe 5 with AA MH and SA MH using non-system IPv4 addresses shows the example topology for EVPN-VXLAN Epipe 5 with AA MH ES "vES23_5" between PE-2 and PE-3 and SA MH ES "ES45_5" between PE-4 and PE-5.

Figure 162: EVPN-VXLAN Epipe 5 with AA MH and SA MH using non-system IPv4 addresses

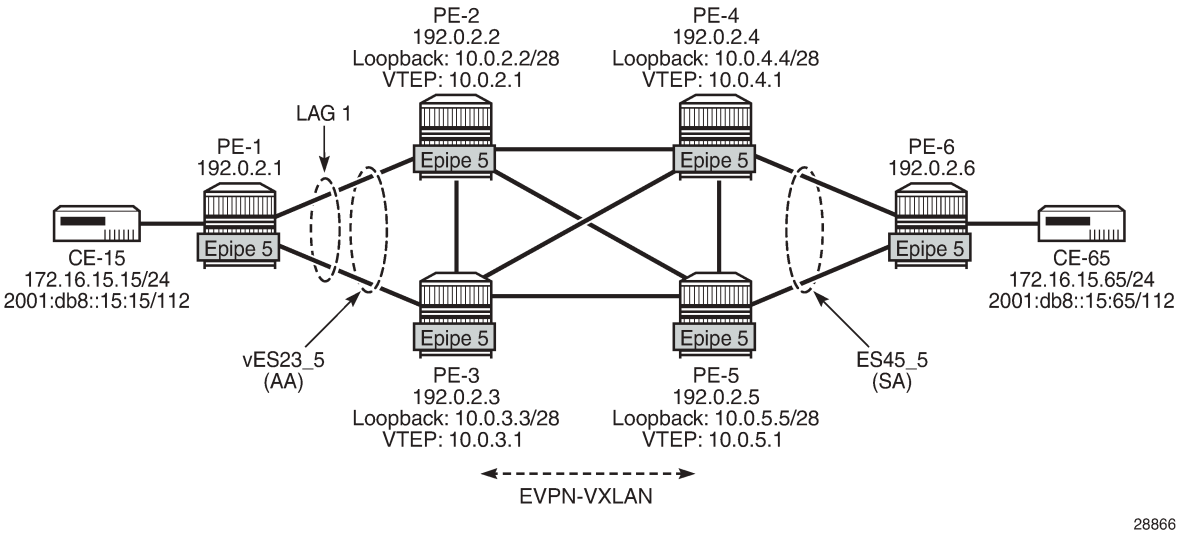

The configuration of Epipe 5 on PE-1 is similar to the configuration of Epipe 4 on PE-1, so it is not shown here. The same applies for Epipe 5 on PE-6.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

On PE-2, VTEP 10.0.2.1 is used instead of the system IP address. The ES must include two additional parameters for the DF selection: **orig-ip** and **route-next-hop**, which are both equal to the VTEP. Without these parameters, the DF selection will not work. The **orig-ip** command modifies the originator IP address of the ES route and the **route-next-hop** modifies the next-hop of the AD per-ES routes for the ES. The service configuration on PE-2 is as follows:

```
# on PE-2:
configure {
    service {
        system {
            vxlan {
                tunnel-termination 10.0.2.1 {
                    fpe-id 1
                }
                tunnel-termination 2001::2:1 {
                    fpe-id 1
                }
            }
            bgp {
                evpn {
                    ethernet-segment "vES23_5" {
                        admin-state enable
                        type virtual
                        esi 01:00:00:00:00:23:00:00:00:05
                        orig-ip 10.0.2.1
                        route-next-hop 10.0.2.1
                        multi-homing-mode all-active
                        association {
                            lag "lag-1" {
                                virtual-ranges {
                                    dot1q {
                                        q-tag 5 {
                                            end 5
                                        }
                                    }
                                }
                            }
                        }
                    }
                }
            }
        }
        epipe "Epipe-5" {
            admin-state enable
            service-id 5
            customer "1"
            bgp 1 {
            }
            sap lag-1:5 {
            }
            vxlan {
                source-vtep 10.0.2.1
                instance 1 {
                    vni 5
                }
            }
            bgp-evpn {
                evi 5
                local-attachment-circuit "AC-23_2" {
                    eth-tag 223
                }
                remote-attachment-circuit "AC-45_2" {
                    eth-tag 245
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                    ecmp 2
                }
            }
        }
```

The service configuration on PE-3 is similar.

On PE-4, the service configuration is as follows:

```
# on PE-4:
configure {
    service {
        sdp 465 {
            admin-state enable
            description "GRE SDP for SA MH_Epipe5"
            far-end {
                ip-address 192.0.2.6
            }
        }
        system {
            vxlan {
                tunnel-termination 10.0.4.1 {
                    fpe-id 1
                }
                tunnel-termination 2001::4:1 {
                    fpe-id 1
                }
            }
            bgp {
                evpn {
                    ethernet-segment "ES45_5" {
                        admin-state enable
                        esi 01:00:00:00:00:45:00:00:00:05
                        orig-ip 10.0.4.1
                        route-next-hop 10.0.4.1
                        multi-homing-mode single-active
                        association {
                            sdp 465 {
                            }
                        }
                    }
                }
            }
        }
        epipe "Epipe-5" {
            admin-state enable
            service-id 5
            customer "1"
            bgp 1 {
            }
            spoke-sdp 465:5 {
            }
            vxlan {
                source-vtep 10.0.4.1
                instance 1 {
                    vni 5
                }
            }
            bgp-evpn {
                evi 5
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                local-attachment-circuit "AC-45_2" {
                    eth-tag 245
                }
                remote-attachment-circuit "AC-23_2" {
                    eth-tag 223
                }
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                    ecmp 2
                }
            }
        }
```

In the AA MH ES, both PE-2 and PE-3 are DF. PE-4 receives BGP-EVPN autodiscovery routes with Ethernet tag 223 from PE-2 and PE-3 with the primary flag set to 1, as follows:

```
[/]
A:admin@PE-4# show router bgp routes evpn auto-disc tag 223 detail
                                                              | match 'C:|Route Dist'
Community      : target:64500:5 l2-attribute:MTU: 1514 C: 0 P: 1 B: 0
Route Dist.    : 192.0.2.2:5
Community      : target:64500:5 l2-attribute:MTU: 1514 C: 0 P: 1 B: 0
Route Dist.    : 192.0.2.2:5
Community      : target:64500:5 l2-attribute:MTU: 1514 C: 0 P: 1 B: 0
Route Dist.    : 192.0.2.3:5
Community      : target:64500:5 l2-attribute:MTU: 1514 C: 0 P: 1 B: 0
Route Dist.    : 192.0.2.3:5
```

The VXLAN destinations for Epipe 5 on PE-4 are the non-system TEP addresses 10.0.2.1 and 10.0.3.1 in ES "vES23_5" with ESI 01:00:00:00:00:23:00:00:00:05, as follows:

```
[/]
A:admin@PE-4# show service id 5 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address                          Egress VNI          Oper      Vxlan
                                                          State     Type
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================

===============================================================================
BGP EVPN VXLAN ES Dest
===============================================================================
I Eth Seg Id                    TEP Address       VNI       Last Changed
-------------------------------------------------------------------------------
1 01:00:00:00:00:23:00:00:00:05 10.0.2.1          5         06/30/2021 17:42:03
1 01:00:00:00:00:23:00:00:00:05 10.0.3.1          5         06/30/2021 17:42:03
-------------------------------------------------------------------------------
===============================================================================
```

In the SA MH ES, PE-5 is DF and PE-4 is NDF. PE-2 receives BGP-EVPN autodiscovery routes with Ethernet tag 245 from PE-4 with backup flag 1 and from PE-5 with primary flag 1, as follows:

```
[/]
A:admin@PE-2# show router bgp routes evpn auto-disc tag 245 detail
                                                              | match 'C:|Route Dist'
Community      : target:64500:5 l2-attribute:MTU: 1514 C: 0 P: 0 B: 1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Route Dist.    : 192.0.2.4:5
Community      : target:64500:5 l2-attribute:MTU: 1514 C: 0 P: 0 B: 1
Route Dist.    : 192.0.2.4:5
Community      : target:64500:5 l2-attribute:MTU: 1514 C: 0 P: 1 B: 0
Route Dist.    : 192.0.2.5:5
Community      : target:64500:5 l2-attribute:MTU: 1514 C: 0 P: 1 B: 0
Route Dist.    : 192.0.2.5:5
```

The VXLAN destination for Epipe 5 on PE-2 is the non-system TEP address 10.0.5.1 of DF PE-5 in ES "ES45_5" with ESI 01:00:00:00:00:45:00:00:00:05, as follows:

```
[/]
A:admin@PE-2# show service id 5 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address                            Egress VNI          Oper      Vxlan
                                                            State     Type
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================

===============================================================================
BGP EVPN VXLAN ES Dest
===============================================================================
I Eth Seg Id                    TEP Address       VNI       Last Changed
-------------------------------------------------------------------------------
1 01:00:00:00:00:45:00:00:00:05 10.0.5.1          5         06/30/2021 17:42:39
-------------------------------------------------------------------------------
===============================================================================
```

AA and SA multi-homed EVPN-VXLAN Epipe using non-system IPv6 addresses

Figure 163: EVPN-VXLAN Epipe 6 with AA MH and SA MH using non-system IPv6 addresses shows the example topology for EVPN-VXLAN Epipe 6 with AA MH ES "vES23_6" between PE-2 and PE-3 and SA MH ES "ES45_6" between PE-4 and PE-5.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 163: EVPN-VXLAN Epipe 6 with AA MH and SA MH using non-system IPv6 addresses

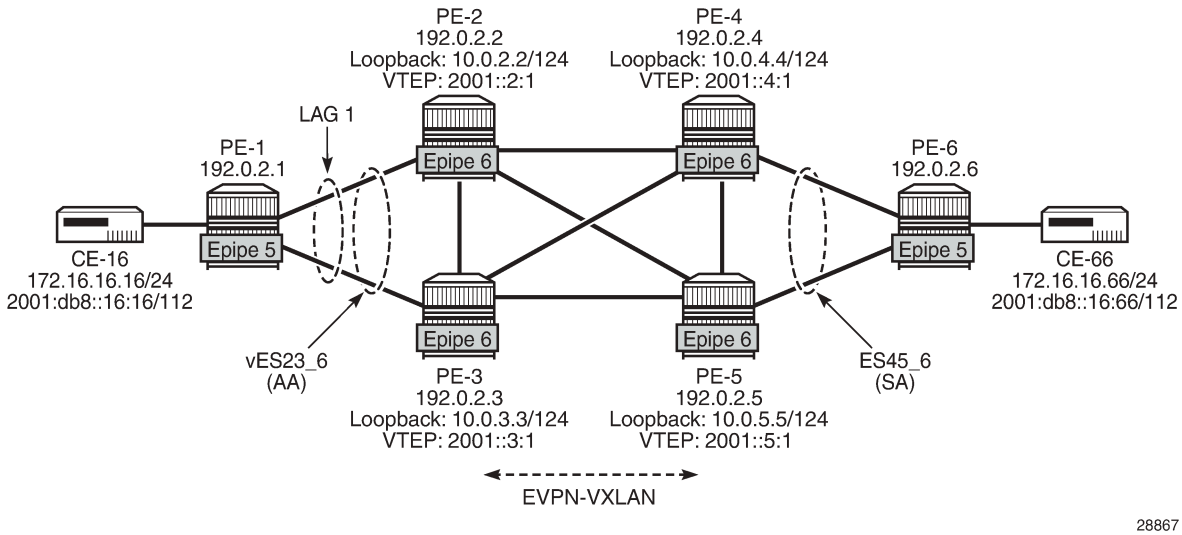

The service configuration on PE-2 is as follows:

```
# on PE-2:
configure {
    service {
        system {
            vxlan {
                tunnel-termination 2001::2:1 {
                    fpe-id 1
                }
            }
            bgp {
                evpn {
                    ethernet-segment "vES23_6" {
                        admin-state enable
                        type virtual
                        esi 01:00:00:00:00:23:00:00:00:06
                        orig-ip 2001::2:1
                        route-next-hop 2001::2:1
                        multi-homing-mode all-active
                        association {
                            lag "lag-1" {
                                virtual-ranges {
                                    dot1q {
                                        q-tag 6 {
                                            end 6
                                        }
                                    }
                                }
                            }
                        }
                    }
                }
            }
        }
        epipe "Epipe-6" {
            admin-state enable
            service-id 6
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            customer "1"
            bgp 1 {
            }
            sap lag-1:6 {
            }
            vxlan {
                source-vtep 2001::2:1
                instance 1 {
                    vni 6
                }
            }
            bgp-evpn {
                evi 6
                local-attachment-circuit "AC-23_v6" {
                    eth-tag 623
                }
                remote-attachment-circuit "AC-45_v6" {
                    eth-tag 645
                }
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                    ecmp 2
                }
            }
        }
```

The service configuration on PE-4 is as follows:

```
# on PE-4:
configure {
    service {
        system {
            vxlan {
                tunnel-termination 10.0.4.1 {
                    fpe-id 1
                }
                tunnel-termination 2001::4:1 {
                    fpe-id 1
                }
            }
            bgp {
                evpn {
                    ethernet-segment "ES45_6" {
                        admin-state enable
                        esi 01:00:00:00:00:45:00:00:00:06
                        orig-ip 2001::4:1
                        route-next-hop 2001::4:1
                        multi-homing-mode single-active
                        association {
                            sdp 466 {
                            }
                        }
                    }
                }
            }
        }
        epipe "Epipe-6" {
            admin-state enable
            service-id 6
            customer "1"
            bgp 1 {
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            spoke-sdp 466:6 {
            }
            vxlan {
                source-vtep 2001::4:1
                instance 1 {
                    vni 6
                }
            }
            bgp-evpn {
                evi 6
                local-attachment-circuit "AC-45_v6" {
                    eth-tag 645
                }
                remote-attachment-circuit "AC-23_v6" {
                    eth-tag 623
                }
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                    ecmp 2
                }
            }
        }
        sdp 466 {
            admin-state enable
            description "GRE SDP for SA MH_Epipe6"
            far-end {
                ip-address 192.0.2.6
            }
        }
```

Conclusion

EVPN-VXLAN VPWS is similar to EVPN-MPLS VPWS, and can be used in networks without MPLS.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Inter-AS Model C for VLL

This chapter describes advanced inter-AS model C for Virtual Leased Line (VLL) configurations.

Topics in this chapter include:

Applicability

This chapter was initially written for SR OS Release 8.0.R4. The MD-CLI in the current edition corresponds to SR OS Release 20.10.R2.

Overview

SR OS supports RFC 3107, *Carrying Label Information in BGP-4*, including VLL/VPLS. BGP SDPs can also be used with PBB-VPLS services.

Internet service providers are looking for mechanisms to implement the VLL and VPLS services across Autonomous Systems (ASs). Service providers may have inter-AS operation as a consequence of delivering inter-provider VLL/VPLS or because they use multiple ASs as a result of acquisitions and mergers.

The objective of this chapter is to describe the interconnection of VLL services across multiple ASs, using inter-AS model C. Inter-AS Model C involves eBGP redistribution of internal system addresses to the neighboring AS using labeled IPv4 routes.

Example topology

Figure 164: Example topology – Inter-AS model C for VLL shows the example topology used for Inter-AS Model C VLL.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 164: Example topology – Inter-AS model C for VLL

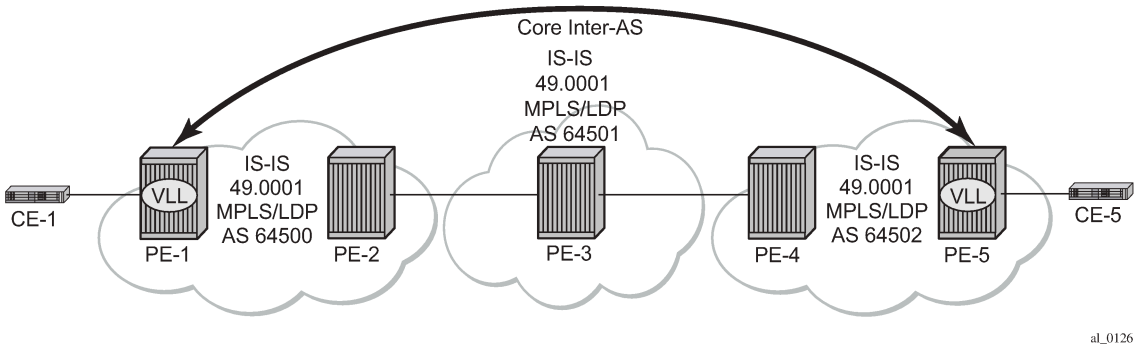

The example topology shown in Figure 164: Example topology – Inter-AS model C for VLL consists of three sites in different ASs with each site using SR OS nodes.

AS 64500 contains PE-1 and PE-2, AS 64501 contains PE-3, and AS 64502 contains PE-4 and PE-5. There is a business customer with two remote locations, Site A and Site B, with Customer Edge (CE) devices CE-1 connected to the AS 64500 via PE-1 and CE-5 connected to the AS 64502 via PE-5. A VLL Epipe service is configured between PE-1 and PE-5 to connect site A and site B.

Figure 165: Inter-AS model C for VLL

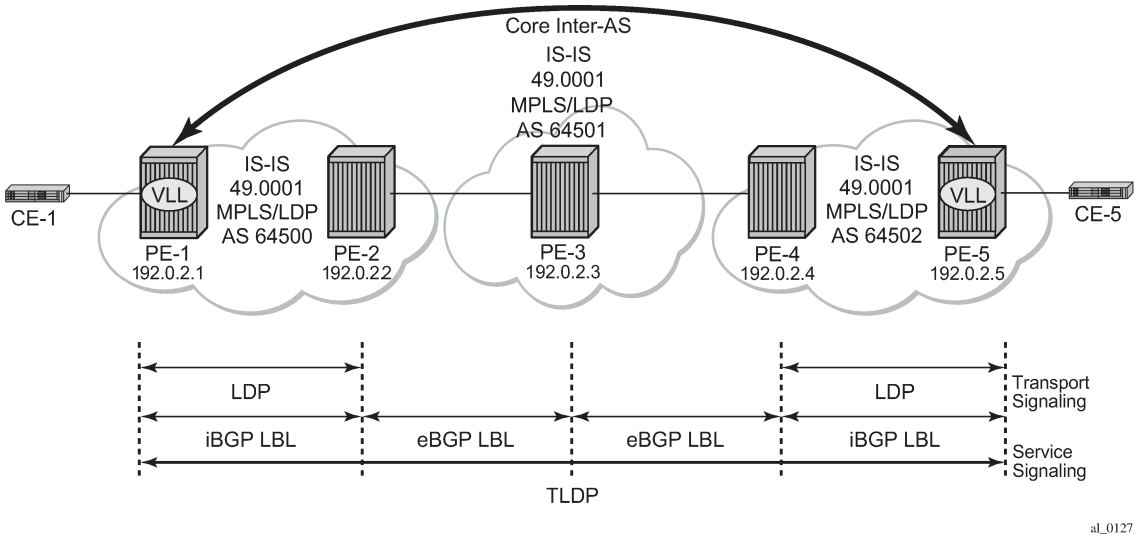

Configuration

This section describes all of the relevant configuration tasks for the detailed setup shown in Figure 166: Network setup configuration. In this particular example, the following protocols are assumed to be already configured.

- IS-IS as the IGP with all the nodes being level Level1/Level 2.
- LDP as the MPLS protocol to signal the transport tunnels within AS 64500 and AS 64502.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 166: Network setup configuration

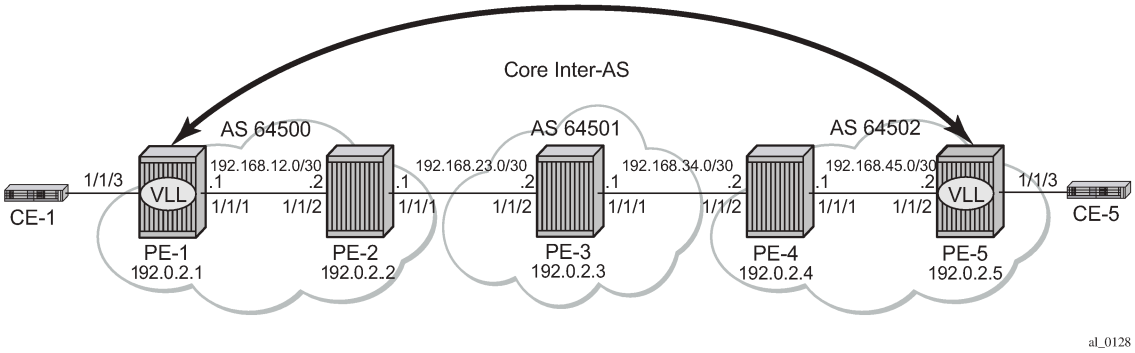

BGP configuration

A BGP tunnel must be established between PE-1 and PE-5, therefore, labeled BGP routes must be exchanged for prefixes 192.0.2.1/32 and 192.0.2.5/32 across the ASs. The following shows the BGP configuration — iBGP and eBGP — required for the PE routers to implement an Inter-AS VLL.

The BGP configuration on PE-3 in AS 64501 is as follows:

```
# on PE-3:
configure {
    router "Base" {
        autonomous-system 64501
        bgp {
            rapid-withdrawal true
            split-horizon true
            group "EBGP" {
                local-as {
                    as-number 64501
                }
            }
            neighbor "192.168.23.1" {
                group "EBGP"
                peer-as 64500
                family {
                    label-ipv4 true
                }
                import {
                    policy ["import-from64500"]
                }
                export {
                    policy ["export-from64502"]
                }
            }
            neighbor "192.168.34.2" {
                group "EBGP"
                peer-as 64502
                family {
                    label-ipv4 true
                }
                import {
                    policy ["import-from64502"]
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                export {
                    policy ["export-from64500"]
                }
            }
        }
```

The address family **label-ipv4** must be configured so that MPLS labels are carried along with MP-BGP Network Layer Reachability Information (NLRIs), see chapter "Separate BGP RIBs for Labeled Routes" in the Unicast Routing Protocols volume in the *7450 ESS, 7750 SR, and 7950 XRS MD-CLI Advanced Configuration Guide - Part I*. The setting **split-horizon** is optional and prevents that a received route is sent back to the originator, which may result in multiple routes for a certain prefix.

To export the prefixes of the nodes where the Epipe is configured (PE-1 and PE-5) to another AS, a common scenario is to advertise the prefix to be exported within the AS as labeled BGP. Therefore, an export policy is defined for prefix192.0.2.1/32 on PE-1 and this prefix will be advertised with community "64500:0" to the ASBR in AS 64500, in this case to PE-2.

On PE-2, the labeled BGP route for prefix 192.0.2.1/32 is inactive, because the IGP route for that prefix is preferred. The setting **advertise-inactive** will allow the inactive labeled BGP routes from AS 64500 to be advertised to PE-3 in AS 64501. However, for EBGP sessions on Autonomous System Border Routers (ASBRs) such as PE-2, RFC 8212 is supported, so that routes are neither imported nor exported unless specifically enabled by configuration. Export policy "export-from64500" exports all routes with community "64500:0" and import policy "import-from64502" imports all routes with community "64502:0" which is the community added by the export policy on PE-5.

On PE-5, an export policy is configured to advertise prefix 192.0.2.5/32 with community "64502:0" to ASBR PE-4 in AS 64502.

On PE-4, BGP is configured with **advertise-inactive** to advertise the labeled BGP route to its EBGP peer, PE-3. Export policy "export-64502:0" exports routes with community "64502:0" to PE-3. Import policy "import-64500:0" imports routes with community "64500:0".

On PE-3, import and export policies are configured toward the EBGP peers. PE-3 imports routes with community "64500:0" from PE-2 and exports these routes to PE-4. Simultaneously, PE-3 imports routes with community "64502:0" from PE-4 and exports these routes to PE-2.

Labeled BGP is used end-to-end between PE-1 and PE-5 and no IGP routes are to be redistributed into BGP, which would be the case if no local BGP labeled routes were advertised within AS 64500 or AS 64502 and only IGP routes were defined within these ASs.

The ASBRs PE-2, PE-3, and PE-4 will swap the BGP labels. PE-3 will advertise the labeled BGP routes learned from AS 64500 to AS 64502 and vice versa and the ASBRs will advertise these labeled routes for remote PE prefixes to their BGP peers. Eventually, PE-1 will have learned a labeled BGP route for prefix 192.0.2.5/32 and PE-5 will have learned a labeled BGP route for prefix 192.0.2.1/32 and a VLL Epipe can be established between PE-1 and PE-5.

The following policies are configured on ASBR PE-2:

```
# on PE-2:
configure {
    policy-options {
        community "64500:0" {
            member "64500:0" { }
        }
        community "64502:0" {
            member "64502:0" { }
        }
        policy-statement "export-from64500" {
            entry 10 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                from {
                    community {
                        name "64500:0"
                    }
                }
                action {
                    action-type accept
                }
            }
        }
        policy-statement "import-from64502" {
            entry 10 {
                from {
                    community {
                        name "64502:0"
                    }
                }
                action {
                    action-type accept
                }
            }
        }
```

The BGP configuration of PE-2 in AS 64500 is as follows:

```
# on PE-2:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            rapid-withdrawal true
            split-horizon true
            group "EBGP" {
                local-as {
                    as-number 64500
                }
            }
            group "IBGP" {
            }
            neighbor "192.0.2.1" {
                group "IBGP"
                next-hop-self true
                peer-as 64500
                family {
                    label-ipv4 true
                }
            }
            neighbor "192.168.23.2" {
                advertise-inactive true
                group "EBGP"
                peer-as 64501
                family {
                    label-ipv4 true
                }
                import {
                    policy ["import-from64502"]
                }
                export {
                    policy ["export-from64500"]
                }
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The BGP configuration of ASBR PE-4 in AS 64502 is as follows.The import and export policies are similar to the policies on PE-2.

```
# on PE-4:
configure {
    router "Base" {
        autonomous-system 64502
        bgp {
            rapid-withdrawal true
            split-horizon true
            group "EBGP" {
                local-as {
                    as-number 64502
                }
            }
            group "IBGP" {
            }
            neighbor "192.0.2.5" {
                group "IBGP"
                next-hop-self true
                peer-as 64502
                family {
                    label-ipv4 true
                }
            }
            neighbor "192.168.34.1" {
                advertise-inactive true
                group "EBGP"
                peer-as 64501
                family {
                    label-ipv4 true
                }
                import {
                    policy ["import-from64500"]
                }
                export {
                    policy ["export-from64502"]
                }
            }
        }
```

PE-1 and PE-5 are the PEs to which the CEs are connected in AS 64500 and AS 64502. PE-1 and PE-5 advertise their system prefixes as labeled BGP routes to their BGP peers within the AS.

The BGP configuration of PE-1 is as follows:

```
# on PE-1:
configure {
    router "Base"{
        autonomous-system 64500
        bgp {
            rapid-withdrawal true
            split-horizon true
            group "IBGP" {
                export {
                    policy ["export-PE-1"]
                }
            }
            neighbor "192.0.2.2" {
                group "IBGP"
                next-hop-self true
                peer-as 64500
                family {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    label-ipv4 true
                }
            }
        }
```

The BGP configuration of PE-5 in AS 64502 is as follows:

```
# on PE-5:
configure {
    router "Base" {
        autonomous-system 64502
        bgp {
            rapid-withdrawal true
            split-horizon true
            group "IBGP" {
                export {
                    policy ["export-PEsys"]
                }
            }
            neighbor "192.0.2.4" {
                group "IBGP"
                next-hop-self true
                peer-as 64502
                family {
                    label-ipv4 true
                }
            }
        }
```

Policy configuration

The export policies on PE-1 and PE-5 advertise the system addresses to the remote AS. The added communities are used by the export and import policies on PE-2, PE-3, and PE-4. The export policy on PE-1 has a prefix list that only contains prefix 192.0.2.1/32 as follows:

```
# on PE-1:
configure {
    policy-options {
        community "64500:0" {
            member "64500:0" { }
        }
        prefix-list "PE-1" {
            prefix 192.0.2.1/32 type exact {
            }
        }
        policy-statement "export-PE-1" {
            entry 10 {
                from {
                    prefix-list ["PE-1"]
                }
                action {
                    action-type accept
                    origin igp
                    community {
                        add ["64500:0"]
                    }
                }
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

A similar export policy can be configured for prefix 192.0.2.5/32 on PE-5. However, the export policy on PE-5 is slightly different: the policy has a prefix list that can be applied for prefixes on multiple PEs, but in this case, only prefix 192.0.2.5/32 will be exported:

```
# on PE-5:
configure {
    policy-options {
        community "64502:0" {
            member "64502:0" { }
        }
        prefix-list "PEsys" {
            prefix 192.0.2.0/29 type longer {
            }
        }
        policy-statement "export-PEsys" {
            entry 10 {
                from {
                    prefix-list ["PEsys"]
                    protocol {
                        name [direct]
                    }
                }
                action {
                    action-type accept
                    origin igp
                    community {
                        add ["64502:0"]
                    }
                }
            }
        }
```

The same policy could have been applied on PE-1.

Service configuration

Once BGP is configured, the configuration requires the service to be defined (Epipe 1). The focus here is a VLL service, however, it is also possible to have a similar configuration with VPLS services.

The following shows the service level configuration on PE-1:

```
# on PE-1:
configure {
     service {
        epipe "Epipe 1" {
            admin-state enable
            description "Tunnel-PE-1-PE-5"
            service-id 1
            customer "1"
            spoke-sdp 15:1 {
            }
            sap 1/1/3:1 {
            }
        }
        sdp 15 {
            admin-state enable
            delivery-type mpls
            bgp-tunnel true
            far-end {
                ip-address 192.0.2.5
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
        }
```

The following CLI shows the service level configuration on PE-5:

```
# on PE-5:
configure {
    service {
        epipe "Epipe 1" {
            admin-state enable
            description "Tunnel-PE-5-PE-1"
            service-id 1
            customer "1"
            spoke-sdp 51:1 {
            }
            sap 1/1/3:1 {
            }
        }
        sdp 51 {
            admin-state enable
            delivery-type mpls
            bgp-tunnel true
            far-end {
                ip-address 192.0.2.1
            }
        }
```

Show commands and troubleshooting

On PE-5, BGP tunnels exist to the remote AS system addresses that are using LDP as a transport mechanism and the configuration of end-to-end SDPs over which T-LDP service labels are exchanged.

In the following sections, the same commands are launched on the nodes in the following order: first on PE-1 and PE-5; then on PE-3, and finally, on PE-2 and PE-4.

Show commands and troubleshooting on PE-1

The following shows information about SDP 15 on PE-1:

```
[]
A:admin@PE-1# show service sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId  AdmMTU  OprMTU  Far End          Adm  Opr         Del     LSP   Sig
-------------------------------------------------------------------------------
15     0       1552    192.0.2.5        Up   Up          MPLS    B     TLDP
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
Legend: R = RSVP, L = LDP, B = BGP, M = MPLS-TP, n/a = Not Applicable
        I = SR-ISIS, O = SR-OSPF, T = SR-TE, F = FPE
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

On PE-1, the VLL Epipe service is up, as follows:

```
[]
A:admin@PE-1# show service service-using

===============================================================================
Services
===============================================================================
ServiceId    Type      Adm  Opr  CustomerId Service Name
-------------------------------------------------------------------------------
1            Epipe     Up   Up   1          Epipe 1
2147483648   IES       Up   Down 1          _tmnx_InternalIesService
2147483649   intVpls   Up   Down 1          _tmnx_InternalVplsService
-------------------------------------------------------------------------------
Matching Services : 3
-------------------------------------------------------------------------------
===============================================================================
```

Two LDP sessions have been established from PE-1: a link LDP session with neighbor PE-2 in AS 64500 and a targeted LDP session with PE-5 in AS 64502, as follows:

```
[]
A:admin@PE-1# show router ldp session ipv4

===============================================================================
LDP IPv4 Sessions
===============================================================================
Peer LDP Id         Adj Type  State         Msg Sent  Msg Recv  Up Time
-------------------------------------------------------------------------------
192.0.2.2:0         Link      Established   109       111       0d 00:04:31
192.0.2.5:0         Targeted  Established   21        23        0d 00:01:20
-------------------------------------------------------------------------------
No. of IPv4 Sessions: 2
===============================================================================
```

The route table on PE-1 shows that the system IP address of PE-5 is reachable using a BGP tunnel:

```
[]
A:admin@PE-1# show router route-table

===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
192.0.2.1/32                                  Local   Local     00h06m03s  0
       system                                                        0
192.0.2.2/32                                  Remote  ISIS      00h04m48s  15
       192.168.12.2                                                  10
192.0.2.5/32                                  Remote  BGP_LABEL 00h01m46s  170
       192.0.2.2 (tunneled)                                          10
192.168.12.0/30                               Local   Local     00h06m03s  0
       int-PE-1-PE-2                                                 0
-------------------------------------------------------------------------------
No. of Routes: 4
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The following tunnel-table on PE-1 shows the details of the LDP, SDP, and BGP tunnels.

```
[]
A:admin@PE-1# show router tunnel-table

===============================================================================
IPv4 Tunnel Table (Router: Base)
===============================================================================
Destination           Owner     Encap TunnelId  Pref   Nexthop        Metric
   Color
-------------------------------------------------------------------------------
192.0.2.2/32          ldp       MPLS  65537     9      192.168.12.2   10
192.0.2.5/32          sdp       MPLS  15        5      192.0.2.5      0
192.0.2.5/32          bgp       MPLS  262145    12     192.0.2.2      1000
-------------------------------------------------------------------------------
Flags: B = BGP or MPLS backup hop available
       L = Loop-Free Alternate (LFA) hop available
       E = Inactive best-external BGP route
       k = RIB-API or Forwarding Policy backup hop
===============================================================================
```

The service details for Epipe 1 on PE-1 are as follows:

```
[]
A:admin@PE-1# show service id 1 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 1                   Vpn Id            : 0
Service Type      : Epipe
MACSec enabled    : no
Name              : Epipe 1
Description       : Tunnel-PE-1-PE-5
Customer Id       : 1                   Creation Origin   : manual
Last Status Change: 01/21/2021 16:01:07
Last Mgmt Change  : 01/21/2021 16:00:53
Test Service      : No
Admin State       : Up                  Oper State        : Up
MTU               : 1514
Vc Switching      : False
SAP Count         : 1                   SDP Bind Count    : 1
Per Svc Hashing   : Disabled
Vxlan Src Tep Ip  : N/A
Force QTag Fwd    : Disabled
Oper Group        : <none>

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                             Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/3:1                            q-tag        1518    1518    Up   Up
sdp:15:1 S(192.0.2.5)                  Spok         0       1552    Up   Up
===============================================================================
```

ICMP is used to verify the IP connectivity from PE-1 to the system IP address of PE-5:

```
[]
A:admin@PE-1# ping 192.0.2.5
PING 192.0.2.5 56 data bytes
64 bytes from 192.0.2.5: icmp_seq=1 ttl=64 time=1.91ms.
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
64 bytes from 192.0.2.5: icmp_seq=2 ttl=64 time=2.06ms.
64 bytes from 192.0.2.5: icmp_seq=3 ttl=64 time=2.02ms.
64 bytes from 192.0.2.5: icmp_seq=4 ttl=64 time=2.01ms.
64 bytes from 192.0.2.5: icmp_seq=5 ttl=64 time=2.02ms.

---- 192.0.2.5 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 1.91ms, avg = 2.01ms, max = 2.06ms, stddev = 0.050ms
```

Show commands and troubleshooting on PE-5

The same commands on PE-5 result in the following output:

```
[]
A:admin@PE-5# show service sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId  AdmMTU  OprMTU  Far End          Adm  Opr         Del     LSP   Sig
-------------------------------------------------------------------------------
51     0       1552    192.0.2.1        Up   Up          MPLS    B     TLDP
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
Legend: R = RSVP, L = LDP, B = BGP, M = MPLS-TP, n/a = Not Applicable
        I = SR-ISIS, O = SR-OSPF, T = SR-TE, F = FPE
===============================================================================
```

```
[]
A:admin@PE-5# show service service-using

===============================================================================
Services
===============================================================================
ServiceId    Type      Adm  Opr  CustomerId Service Name
-------------------------------------------------------------------------------
1            Epipe     Up   Up   1          Epipe 1
2147483648   IES       Up   Down 1          _tmnx_InternalIesService
2147483649   intVpls   Up   Down 1          _tmnx_InternalVplsService
-------------------------------------------------------------------------------
Matching Services : 3
-------------------------------------------------------------------------------
===============================================================================
```

```
[]
A:admin@PE-5# show router ldp session ipv4

===============================================================================
LDP IPv4 Sessions
===============================================================================
Peer LDP Id         Adj Type  State         Msg Sent  Msg Recv  Up Time
-------------------------------------------------------------------------------
192.0.2.1:0         Targeted  Established   52        53        0d 00:04:07
192.0.2.4:0         Link      Established   185       188       0d 00:07:57
-------------------------------------------------------------------------------
No. of IPv4 Sessions: 2
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
```

```
[]
A:admin@PE-5# show router route-table

===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
192.0.2.1/32                                    Remote  BGP_LABEL 00h05m47s  170
       192.0.2.4 (tunneled)                                         10
192.0.2.4/32                                    Remote  ISIS      00h08m13s  15
       192.168.45.1                                                 10
192.0.2.5/32                                    Local   Local     00h08m19s  0
       system                                                       0
192.168.45.0/30                                 Local   Local     00h08m19s  0
       int-PE-5-PE-4                                                0
-------------------------------------------------------------------------------
No. of Routes: 4
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

```
[]
A:admin@PE-5# show router tunnel-table

===============================================================================
IPv4 Tunnel Table (Router: Base)
===============================================================================
Destination           Owner     Encap TunnelId  Pref   Nexthop        Metric
   Color
-------------------------------------------------------------------------------
192.0.2.1/32          sdp       MPLS  51        5      192.0.2.1      0
192.0.2.1/32          bgp       MPLS  262145    12     192.0.2.4      1000
192.0.2.4/32          ldp       MPLS  65537     9      192.168.45.1   10
-------------------------------------------------------------------------------
Flags: B = BGP or MPLS backup hop available
       L = Loop-Free Alternate (LFA) hop available
       E = Inactive best-external BGP route
       k = RIB-API or Forwarding Policy backup hop
===============================================================================
```

```
[]
A:admin@PE-5# show service id 1 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 1                   Vpn Id            : 0
Service Type      : Epipe
MACSec enabled    : no
Name              : Epipe 1
Description       : Tunnel-PE-5-PE-1
Customer Id       : 1                   Creation Origin   : manual
Last Status Change: 01/21/2021 16:01:07
Last Mgmt Change  : 01/21/2021 16:00:49
Test Service      : No
Admin State       : Up                  Oper State        : Up
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
MTU               : 1514
Vc Switching      : False
SAP Count         : 1                   SDP Bind Count    : 1
Per Svc Hashing   : Disabled
Vxlan Src Tep Ip  : N/A
Force QTag Fwd    : Disabled
Oper Group        : <none>

------------------------------------------------------------------------------
Service Access & Destination Points
------------------------------------------------------------------------------
Identifier                              Type         AdmMTU  OprMTU  Adm  Opr
------------------------------------------------------------------------------
sap:1/1/3:1                             q-tag        1518    1518    Up   Up
sdp:51:1 S(192.0.2.1)                   Spok         0       1552    Up   Up
==============================================================================
```

```
[]
A:admin@PE-5# ping 192.0.2.1
PING 192.0.2.1 56 data bytes
64 bytes from 192.0.2.1: icmp_seq=1 ttl=64 time=1.83ms.
64 bytes from 192.0.2.1: icmp_seq=2 ttl=64 time=2.06ms.
64 bytes from 192.0.2.1: icmp_seq=3 ttl=64 time=2.01ms.
64 bytes from 192.0.2.1: icmp_seq=4 ttl=64 time=2.08ms.
64 bytes from 192.0.2.1: icmp_seq=5 ttl=64 time=2.15ms.

---- 192.0.2.1 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 1.83ms, avg = 2.03ms, max = 2.15ms, stddev = 0.107ms
```

On PE-5, the BGP route to the system IP address of PE-1 can been seen with PE-4 as the next hop:

```
[]
A:admin@PE-5# show router bgp routes label-ipv4
===============================================================================
 BGP Router ID:192.0.2.5        AS:64502       Local AS:64502
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop (Router)                                 Path-Id     IGP Cost
      As-Path                                                      Label
-------------------------------------------------------------------------------
u*>i  192.0.2.1/32                                     100         None
      192.0.2.4                                        None        10
      64501 64500                                                  524285
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

On PE-5, the FIB on slot 1 shows that the system IP address of PE-1 is reachable using BGP over an LDP transport to PE-4:

```
[]
A:admin@PE-5# show router fib 1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
FIB Display
===============================================================================
Prefix [Flags]                                                  Protocol
  NextHop
-------------------------------------------------------------------------------
192.0.2.1/32                                                    BGP_LABEL
  192.0.2.4 (Transport:LDP)
192.0.2.4/32                                                    ISIS
  192.168.45.1 (int-PE-5-PE-4)
192.0.2.5/32                                                    LOCAL
  192.0.2.5 (system)
192.168.45.0/30                                                 LOCAL
  192.168.45.0 (int-PE-5-PE-4)
-------------------------------------------------------------------------------
Total Entries : 4
-------------------------------------------------------------------------------
===============================================================================
```

Show commands on PE-3

The **show** commands on router PE-3 in AS 64501 are as follows:

```
[]
A:admin@PE-3# show router bgp summary all

===============================================================================
BGP Summary
===============================================================================
Legend : D - Dynamic Neighbor
===============================================================================
Neighbor
Description
ServiceId          AS PktRcvd InQ  Up/Down   State|Rcv/Act/Sent (Addr Family)
                      PktSent OutQ
-------------------------------------------------------------------------------
192.168.23.1
Def. Instance  64500       22    0 00h09m12s 1/1/1 (Lbl-IPv4)
                           22    0
192.168.34.2
Def. Instance  64502       23    0 00h09m04s 1/1/1 (Lbl-IPv4)
                           25    0

-------------------------------------------------------------------------------
```

```
[]
A:admin@PE-3# show router bgp routes label-ipv4
===============================================================================
 BGP Router ID:192.0.2.3        AS:64501       Local AS:64501
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP Routes
===============================================================================
Flag  Network                                          LocalPref   MED
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
      Nexthop (Router)                                      Path-Id     IGP Cost
      As-Path                                                           Label
------------------------------------------------------------------------------
u*>i  192.0.2.1/32                                          None        None
      192.168.23.1                                          None        0
      64500                                                             524285
u*>i  192.0.2.5/32                                          None        None
      192.168.34.2                                          None        0
      64502                                                             524284
------------------------------------------------------------------------------
Routes : 2
==============================================================================
```

The BGP labels are swapped at PE-3, as follows:

```
[]
A:admin@PE-3# show router bgp inter-as-label

==============================================================================
BGP Inter-AS labels
Flags: B - entry has backup, P - entry is promoted
==============================================================================
NextHop                      Received       Advertised     Label
                             Label          Label          Origin
------------------------------------------------------------------------------
192.168.23.1                 524285         524287         External
192.168.34.2                 524284         524286         External
------------------------------------------------------------------------------
Total Labels allocated:   2
==============================================================================
```

The routing table on PE-3 includes BGP labeled routes to PE-1 and PE-5, as follows:

```
[]
A:admin@PE-3# show router route-table

==============================================================================
Route Table (Router: Base)
==============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
------------------------------------------------------------------------------
192.0.2.1/32                                  Remote  BGP_LABEL 00h10m02s  170
       192.168.23.1                                                 0
192.0.2.3/32                                  Local   Local     00h12m42s  0
       system                                                       0
192.0.2.5/32                                  Remote  BGP_LABEL 00h09m23s  170
       192.168.34.2                                                 0
192.168.23.0/30                               Local   Local     00h12m42s  0
       int-PE-3-PE-2                                                0
192.168.34.0/30                               Local   Local     00h12m42s  0
       int-PE-3-PE-4                                                0
------------------------------------------------------------------------------
No. of Routes: 5
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
==============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Show commands on PE-2

The commands on PE-2 are as follows:

```
[]
A:admin@PE-2# show router bgp summary all

===============================================================================
BGP Summary
===============================================================================
Legend : D - Dynamic Neighbor
===============================================================================
Neighbor
Description
ServiceId          AS PktRcvd InQ  Up/Down   State|Rcv/Act/Sent (Addr Family)
                      PktSent OutQ
-------------------------------------------------------------------------------
192.0.2.1
Def. Instance  64500       36    0 00h16m24s 1/0/1 (Lbl-IPv4)
                           36    0
192.168.23.2
Def. Instance  64501       36    0 00h16m19s 1/1/1 (Lbl-IPv4)
                           36    0

-------------------------------------------------------------------------------
```

The BGP labels are swapped by PE-2 as follows:

```
[]
A:admin@PE-2# show router bgp inter-as-label

===============================================================================
BGP Inter-AS labels
Flags: B - entry has backup, P - entry is promoted
===============================================================================
NextHop                      Received       Advertised     Label
                             Label          Label          Origin
-------------------------------------------------------------------------------
192.0.2.1                    524285         524285         Internal
192.168.23.2                 524286         524284         External
-------------------------------------------------------------------------------
Total Labels allocated:   2
===============================================================================
```

```
[]
A:admin@PE-2# show router route-table

===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
192.0.2.1/32                                  Remote  ISIS      00h13m43s  15
       192.168.12.1                                               10
192.0.2.2/32                                  Local   Local     00h13m50s  0
       system                                                     0
192.0.2.5/32                                  Remote  BGP_LABEL 00h11m01s  170
       192.168.23.2                                               0
192.168.12.0/30                               Local   Local     00h13m50s  0
       int-PE-2-PE-1                                              0
192.168.23.0/30                               Local   Local     00h13m50s  0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
       int-PE-2-PE-3                                                    0
-------------------------------------------------------------------------------
No. of Routes: 5
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

Show commands on PE-4

The **show** commands on PE-4 are the following:

```
[]
A:admin@PE-4# show router bgp summary all

===============================================================================
BGP Summary
===============================================================================
Legend : D - Dynamic Neighbor
===============================================================================
Neighbor
Description
ServiceId          AS PktRcvd InQ  Up/Down   State|Rcv/Act/Sent (Addr Family)
                      PktSent OutQ
-------------------------------------------------------------------------------
192.0.2.5
Def. Instance  64502       29    0 00h12m43s 1/0/1 (Lbl-IPv4)
                           29    0
192.168.34.1
Def. Instance  64501       29    0 00h12m53s 1/1/1 (Lbl-IPv4)
                           30    0

-------------------------------------------------------------------------------
```

```
[]
A:admin@PE-4# show router bgp routes label-ipv4
===============================================================================
 BGP Router ID:192.0.2.4        AS:64502       Local AS:64502
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop (Router)                                 Path-Id     IGP Cost
      As-Path                                                      Label
-------------------------------------------------------------------------------
u*>i  192.0.2.1/32                                     None        None
      192.168.34.1                                     None        0
      64501 64500                                                  524287
*i    192.0.2.5/32                                     100         None
      192.0.2.5                                        None        10
      No As-Path                                                   524285
-------------------------------------------------------------------------------
Routes : 2
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
```

```
[]
A:admin@PE-4# show router bgp inter-as-label

===============================================================================
BGP Inter-AS labels
Flags: B - entry has backup, P - entry is promoted
===============================================================================
NextHop                       Received       Advertised     Label
                              Label          Label          Origin
-------------------------------------------------------------------------------
192.0.2.5                     524285         524284         Internal
192.168.34.1                  524287         524285         External
-------------------------------------------------------------------------------
Total Labels allocated:   2
===============================================================================
```

Conclusion

The BGP tunnel-based SDP binding is allowed for VLL and VPLS services, including PBB-VPLS. Using RFC 3107, it is possible to implement inter-AS Model C VLLs.

The example used in this chapter illustrates the configuration of an Inter-AS VLL providing access to CE sites. Troubleshooting commands also have been shown to verify all the procedures.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

L2 Multicast in EVPN-MPLS VPRN R-VPLS with All-Active Multi-Homing

This chapter provides information about L2 Multicast in EVPN-MPLS VPRN R-VPLS with All-Active Multi-Homing.

Topics in this chapter include:

Applicability

The information and configuration in this chapter are based on SR OS Release 23.7.R1.

Overview

IPv4 multicast traffic can be forwarded from an EVPN-MPLS service into an attached R-VPLS service in which the receiving devices are using EVPN all-active multi-homing.

The routed service to which the R-VPLS service attaches can be an IES or a VPRN service. In this way, IPv4 multicast traffic can be transported using native IP for the IES case or NG-MVPN technologies for the VPRN case.

This feature requires:

- IGMP support on the R-VPLS IP interface
- Forwarding IPv4 multicast traffic from the IP interface of a VPRN or IES to its EVPN-MPLS R-VPLS service
- IGMP snooping within the VPLS of the R-VPLS service
- IGMP snooping state synchronization based on the ESI label to synchronize the IGMP snooping state between the all-active (R-)VPLS LAG SAPs

The configuration used in this chapter is the NG-MVPN scenario as shown in Figure 167: Multicast From an EVPN-MPLS Service Into an R-VPLS With All-Active EVPN Multi-Homing.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 167: Multicast From an EVPN-MPLS Service Into an R-VPLS With All-Active EVPN Multi-Homing

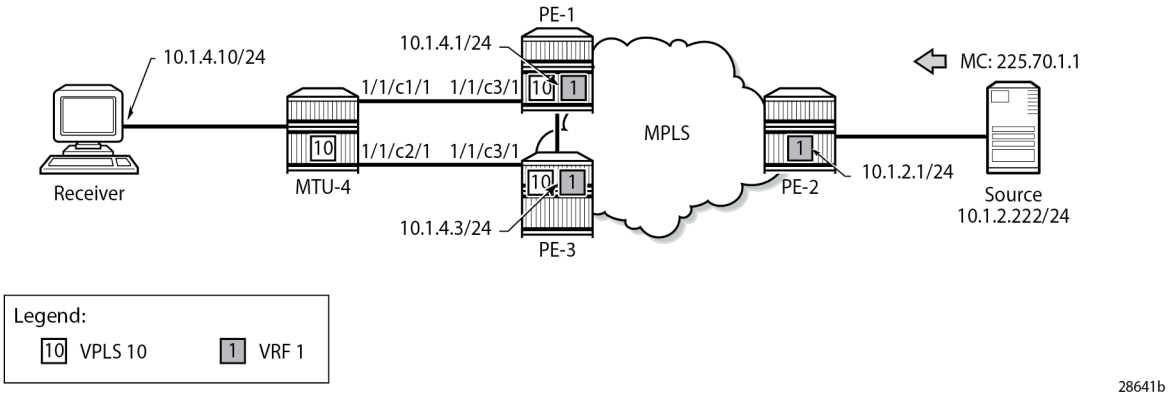

A multicast stream is emitted by the source connected to PE-2 with group address 225.70.1.1. A multicast receiver connected to MTU-4 joins group 225.70.1.1. MTU-4 is connected to PE-1 and PE-3 through an all-active multi-homing EVPN Ethernet segment comprising LAG 1. On MTU-4, LAG 1 comprises port 1/1/c1/1 and 1/1/c2/1, and this LAG is used in VPLS 10. On PE-1 and PE-3, VPLS 10 is interconnected with VPRN 1 through an Integrated Routing and Bridging (IRB) interface. VPRN 1 is defined in PE-1, PE-2, and PE-3, and uses NG-MVPN for transporting the multicast traffic through the core of the network. See the EVPN for MPLS Tunnels and EVPN for MPLS Tunnels in Routed VPLS chapters for more information about EVPN. See the "NG-MVPN Configuration with MPLS" and "NG-MVPN Configuration with PIM" chapters in the Layer 3 Services volume of the *7450 ESS, 7750 SR, and 7950 XRS Advanced Configuration Guide - Part II* for more information about NG-MVPN.

Configuration

The initial configuration on the PE nodes includes the following:

- Cards, MDAs, ports
- Router interfaces
- IS-IS (alternatively, OSPF can be used)
- MPLS tunnels between the PEs: LDP- or RSVP-based

BGP is required at the core of the network, using the VPN IPv4 and MVPN IPv4 address families between all PEs, for supporting unicast and multicast traffic on VPRN services, and additionally using the EVPN address family between PE-1 and PE-3 to support EVPN services. The BGP configurations for PE-1, PE-2, and PE-3 are as follows:

```
# on PE-1:
configure {
    router {
        autonomous-system 64496
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            family {
                ipv4 false
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                vpn-ipv4 true
                mvpn-ipv4 true
                evpn true
            }
            ebgp-default-reject-policy {
                import false
                export false
            }
            rapid-update {
                evpn true
            }
            group "iBGP" { }
            neighbor "192.0.2.2" {
                group "iBGP"
                peer-as 64496
            }
            neighbor "192.0.2.3" {
                group "iBGP"
                peer-as 64496
            }
        }
    }
}
```

```
# on PE-2:
configure {
    router {
        autonomous-system 64496
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            family {
                ipv4 false
                vpn-ipv4 true
                mvpn-ipv4 true
            }
            ebgp-default-reject-policy {
                import false
                export false
            }
            rapid-update {
                evpn true
            }
            group "iBGP" { }
            neighbor "192.0.2.1" {
                group "iBGP"
                peer-as 64496
            }
            neighbor "192.0.2.3" {
                group "iBGP"
                peer-as 64496
            }
        }
    }
}
```

```
# on PE-3:
configure {
    router {
        autonomous-system 64496
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            family {
                ipv4 false
                vpn-ipv4 true
                mvpn-ipv4 true
                evpn true
            }
            ebgp-default-reject-policy {
                import false
                export false
            }
            rapid-update {
                evpn true
            }
            group "iBGP" { }
            neighbor "192.0.2.1" {
                group "iBGP"
                peer-as 64496
            }
            neighbor "192.0.2.2" {
                group "iBGP"
                peer-as 64496
            }
        }
    }
}
```

The receiver connected to MTU-4 joins group 225.70.1.1, and the corresponding multicast stream is emitted by the source that is connected to PE-2. MTU-4 is connected to PE-1 and PE-3 through an all-active multi-homing EVPN Ethernet segment comprising LAG 1. The VPLS and the LAG on MTU-4 are defined as follows:

```
# on MTU-4:
configure {
    service {
        vpls "mcast-vpls" {
            admin-state enable
            service-id 10
            customer "1"
            igmp-snooping {
                admin-state enable
            }
            sap 1/2/c1/1 { }
            sap lag-1:10 { }
        }
    }
}

configure {
    lag "lag-1" {
        admin-state enable
        encap-type dot1q
        mode access
        port 1/1/c1/1 { }
        port 1/1/c2/1 { }
    }
}
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The all-active multi-homing Ethernet segment esi-13 is configured identically on PE-1 and PE-3, as follows. See the EVPN for MPLS Tunnels and EVPN for MPLS Tunnels in Routed VPLS chapters for more information.

```
# on PE-1 and PE-3:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "esi-13" {
                        admin-state enable
                        esi 0x01000000001300000001
                        multi-homing-mode all-active
                        df-election {
                            es-activation-timer 3
                            service-carving-mode manual
                            manual {
                                preference {
                                    mode non-revertive
                                    value 30
                                }
                            }
                        }
                        association {
                            lag "lag-1" { }
                        }
                    }
                }
            }
        }
    }
}
```

The multi-homed access circuits of esi-13 are located on port 1/1/c3/1 for PE-1 and PE-3, so the LAG is configured identically, as follows:

```
# on PE-1 and PE-3:
configure {
    lag "lag-1" {
        admin-state enable
        encap-type dot1q
        mode access
        port 1/1/c3/1 { }
    }
}
```

Also, the EVPN VPLS service with ID 10 is configured identically on PE-1 and PE-3, as follows. The *mcast-vpls* name is needed to link VPLS 10 to VPRN 1 at a later stage, without requiring a physical loop or hairpin. The **routed-vpls** command enables the VPLS to become an R-VPLS. The **igmp-snooping** and **mrouter-port true** commands are required for multicast to work correctly in an all-active multi-homed scenario.

```
# on PE-1 and PE-3:
configure {
    service {
        vpls "mcast-vpls" {
            admin-state enable
            service-id 10
            customer "1"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            routed-vpls {
                multicast {
                    ipv4 {
                        igmp-snooping {
                            mrouter-port true
                        }
                    }
                }
            }
            bgp 1 { }
            igmp-snooping {
                admin-state enable
            }
            bgp-evpn {
                evi 111
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            sap lag-1:10 { }
        }
    }
}
```

The VPRN service with ID 1 provides the connection toward MTU-4 via VPLS 10, through the *int-MCAST-VPLS* interface with address 10.1.4.1/24 on PE-1, and with address 10.1.4.3/24 on PE-3. This L3 interface is linked to VPLS 10 with the **vpls "mcast-vpls" { }** command. The *int-MCAST-VPLS* interface is also included in the IGMP and PIM configurations of VPRN 1. The full configuration of VPRN 1 on PE-1 is as follows. The configuration of VPRN 1 on PE-3 is similar.

```
# on PE-1:
configure {
    service {
        vprn "VPRN 1" {
            admin-state enable
            service-id 1
            customer "1"
            igmp {
                ssm-translate {
                    group-range start 225.70.1.1 end 225.70.255.255 {
                        source 10.1.2.222 { }
                    }
                }
                interface "int-MCAST-VPLS" { }
                interface "int-PE-1-CE-1" { }
            }
            pim {
                interface "int-MCAST-VPLS" { }
                interface "system" { }
            }
            mvpn {
                c-mcast-signaling bgp
                mdt-type receiver-only
                auto-discovery {
                    type bgp
                }
                vrf-target {
                    unicast true
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
                provider-tunnel {
                    inclusive {
                        mldp {
                            admin-state enable
                        }
                    }
                    selective {
                        data-threshold {
                            group-prefix 224.0.0.0/4 {
                                threshold 1
                            }
                        }
                        mldp {
                            admin-state enable
                        }
                    }
                }
            }
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:1"
                    vrf-target {
                        community "target:64496:1"
                    }
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            interface "int-MCAST-VPLS" {
                ipv4 {
                    primary {
                        address 10.1.4.1
                        prefix-length 24
                    }
                }
                vpls "mcast-vpls" { }
            }
            interface "int-PE-1-CE-1" {
                ipv4 {
                    primary {
                        address 10.1.1.1
                        prefix-length 24
                    }
                }
                sap 1/2/c1/1 { }
            }
            interface "system" {
                loopback true
                ipv4 {
                    primary {
                        address 192.0.2.101
                        prefix-length 32
                    }
                }
            }
        }
    }
}
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The full configuration of VPRN 1 on PE-2 is as follows. The *int-PE-2-CE-2-source* interface provides the connection to the multicast source.

```
# on PE-2:
configure {
    service {
        vprn "VPRN 1" {
            admin-state enable
            service-id 1
            customer "1"
            pim {
                interface "int-PE-2-CE-2-source" { }
                interface "system" { }
            }
            mvpn {
                c-mcast-signaling bgp
                mdt-type sender-only
                auto-discovery {
                    type bgp
                }
                vrf-target {
                    unicast true
                }
                provider-tunnel {
                    inclusive {
                        mldp {
                            admin-state enable
                        }
                    }
                    selective {
                        data-threshold {
                            group-prefix 224.0.0.0/4 {
                                threshold 1
                            }
                        }
                        mldp {
                            admin-state enable
                        }
                    }
                }
            }
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:1"
                    vrf-target {
                        community "target:64496:1"
                    }
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            interface "int-PE-2-CE-2-source" {
                ipv4 {
                    primary {
                        address 10.1.2.1
                        prefix-length 24
                    }
                }
                sap 1/2/c1/1 { }
            }
            interface "system" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                loopback true
                ipv4 {
                    primary {
                        address 192.0.2.102
                        prefix-length 32
                    }
                }
            }
        }
    }
}
```

Verification

The following command shows that *esi-13* is an all-active multi-homed Ethernet segment, on PE-1. The same command can be executed on PE-3.

```
[/]
A:admin@PE-1# show service system bgp-evpn ethernet-segment name "esi-13"

===============================================================================
Service Ethernet Segment
===============================================================================
Name                    : esi-13
Eth Seg Type            : None
Admin State             : Enabled            Oper State         : Up
ESI                     : 01:00:00:00:00:13:00:00:00:01
Oper ESI                : 01:00:00:00:00:13:00:00:00:01
Auto-ESI Type           : None
AC DF Capability        : Include
Multi-homing            : allActive          Oper Multi-homing  : allActive
ES SHG Label            : 524282
Source BMAC LSB         : None
Lag                     : lag-1
ES Activation Timer     : 3 secs
Oper Group              : (Not Specified)
Svc Carving             : manual             Oper Svc Carving   : manual
Cfg Range Type          : lowest-pref

-------------------------------------------------------------------------------
DF Pref Election Information
-------------------------------------------------------------------------------
Preference     Preference     Last Admin Change        Oper Pref      Do No
Mode           Value                                   Value          Preempt
-------------------------------------------------------------------------------
non-revertive  30             07/20/2023 11:24:10      30             Disabled
-------------------------------------------------------------------------------
EVI Ranges: <none>
ISID Ranges: <none>
Vprn NextHop EVI Ranges : <none>
===============================================================================
```

The output from the following commands on PE-1 and PE-3 shows that for esi-13, PE-1 is Non-Designated Forwarder (NDF), whereas PE-3 is Designated Forwarder (DF).

```
[/]
A:admin@PE-1# show service id 10 ethernet-segment "esi-13"

===============================================================================
SAP Ethernet-Segment Information
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
SAP                   Eth-Seg                         Status
-------------------------------------------------------------------------------
lag-1:10              esi-13                          NDF
===============================================================================
No sdp entries
No vxlan instance entries
```

```
[/]
A:admin@PE-3# show service id 10 ethernet-segment "esi-13"

===============================================================================
SAP Ethernet-Segment Information
===============================================================================
SAP                   Eth-Seg                         Status
-------------------------------------------------------------------------------
lag-1:10              esi-13                          DF
===============================================================================
No sdp entries
No vxlan instance entries
```

A stream with group address 225.70.1.1 is started by the multicast source and joined by the multicast receiver connected to MTU-4. This stream is forwarded from PE-2 to PE-3; PE-1 is not involved in the forwarding.

PE-1 maintains IGMP state for group 225.70.1.1 in VPRN 1, and so does PE-3. PE-1 and PE-3 synchronize IGMP state using a data-driven mechanism. The forwarding list includes the *int-MCAST-VPLS* interface, as follows:

```
[/]
A:admin@PE-1# show router 1 igmp group 225.70.1.1 interfaces
===============================================================================
IGMP Interface Groups
===============================================================================

(*,225.70.1.1)                                          UpTime: 0d 00:01:41
    Fwd List  : int-MCAST-VPLS
-------------------------------------------------------------------------------
Entries : 1
===============================================================================
```

PE-1 maintains PIM state for group 225.70.1.1, as follows. The outgoing interfaces list is empty and the forwarding rate is zero; both are indications that PE-1 is not forwarding any multicast traffic.

```
[/]
A:admin@PE-1# show router 1 pim group 225.70.1.1 detail

===============================================================================
PIM Source Group ipv4
===============================================================================
Group Address      : 225.70.1.1
Source Address     : 10.1.2.222
RP Address         : 0
Advt Router        : 192.0.2.2
Flags              :                    Type              : (S,G)
Mode               : sparse
MRIB Next Hop      : 192.0.2.2
MRIB Src Flags     : remote
Keepalive Timer Exp: 0d 00:01:50
Up Time            : 0d 00:02:34        Resolved By       : rtable-u
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Up JP State        : Not Joined         Up JP Expiry       : 0d 00:00:00
Up JP Rpt          : Not Joined StarG   Up JP Rpt Override : 0d 00:00:00

Register State     : No Info
Reg From Anycast RP: No

Rpf Neighbor       : 192.0.2.2
Incoming Intf      : mpls-if-73728
Outgoing Intf List :

Curr Fwding Rate   : 0.000 kbps
Forwarded Packets  : 0                  Discarded Packets  : 0
Forwarded Octets   : 0                  RPF Mismatches     : 0
Spt threshold      : 0 kbps             ECMP opt threshold : 7
Admin bandwidth    : 1 kbps
-------------------------------------------------------------------------------
Groups : 1
===============================================================================
```

PE-2 and PE-3 are forwarding the stream as indicated by the PIM state for this group, as follows:

```
[/]
A:admin@PE-2# show router 1 pim group 225.70.1.1 detail

===============================================================================
PIM Source Group ipv4
===============================================================================
Group Address      : 225.70.1.1
Source Address     : 10.1.2.222
RP Address         : 0
Advt Router        : 192.0.2.2
Flags              :                    Type               : (S,G)
Mode               : sparse
MRIB Next Hop      : 10.1.2.222
MRIB Src Flags     : direct
Keepalive Timer    : Not Running
Up Time            : 0d 00:02:04        Resolved By        : rtable-u

Up JP State        : Joined             Up JP Expiry       : 0d 00:00:00
Up JP Rpt          : Not Joined StarG   Up JP Rpt Override : 0d 00:00:00

Register State     : No Info
Reg From Anycast RP: No

Rpf Neighbor       : 10.1.2.222
Incoming Intf      : int-PE-2-CE-2-source
Outgoing Intf List : mpls-if-73728 (mpls-if-73730)

Curr Fwding Rate   : 9745.632 kbps
Forwarded Packets  : 60341              Discarded Packets  : 0
Forwarded Octets   : 89425362           RPF Mismatches     : 0
Spt threshold      : 0 kbps             ECMP opt threshold : 7
Admin bandwidth    : 1 kbps
-------------------------------------------------------------------------------
Groups : 1
===============================================================================
```

```
[/]
A:admin@PE-3# show router 1 pim group 225.70.1.1 detail

===============================================================================
PIM Source Group ipv4
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Group Address      : 225.70.1.1
Source Address     : 10.1.2.222
RP Address         : 0
Advt Router        : 192.0.2.2
Flags              :                    Type              : (S,G)
Mode               : sparse
MRIB Next Hop      : 192.0.2.2
MRIB Src Flags     : remote
Keepalive Timer Exp: 0d 00:01:48
Up Time            : 0d 00:02:35        Resolved By       : rtable-u

Up JP State        : Joined             Up JP Expiry      : 0d 00:00:24
Up JP Rpt          : Not Joined StarG   Up JP Rpt Override : 0d 00:00:00

Register State     : No Info
Reg From Anycast RP: No

Rpf Neighbor       : 192.0.2.2
Incoming Intf      : mpls-if-73728
Incoming SPMSI Intf: mpls-if-73730
Outgoing Intf List : int-MCAST-VPLS

Curr Fwding Rate   : 9745.632 kbps
Forwarded Packets  : 86065              Discarded Packets  : 0
Forwarded Octets   : 127548330          RPF Mismatches     : 0
Spt threshold      : 0 kbps             ECMP opt threshold : 7
Admin bandwidth    : 1 kbps
-------------------------------------------------------------------------------
Groups : 1
===============================================================================
```

The outgoing interfaces on PE-2 and PE-3 are the *mpls-if-73728* PMSI interface and the *int-MCAST-VPLS* interfaces, respectively. The properties of the S-PMSI interface are as follows:

```
[/]
A:admin@PE-2# show router 1 pim tunnel-interface "mpls-if-73728" detail

===============================================================================
PIM Interface ipv4 mpls-if-73728
===============================================================================
Admin Status       : Up                 Oper Status       : Up
IPv4 Admin Status  : Up                 IPv4 Oper Status  : Up
DR                 : 192.0.2.2
Auto-created       : No
Transport Type     : MVPN-Pmsi

-------------------------------------------------------------------------------
PIM Group Source
-------------------------------------------------------------------------------
Group Address      : 225.70.1.1
Source Address     : 10.1.2.222
Interface          : mpls-if-73728      Type              : (S,G)
RP Address         : 0.0.0.0
Up Time            : 0d 00:02:12

Join Prune State   : Join               Expires           : Never
Prune Pend Expires : N/A

Assert State       : No Info
-------------------------------------------------------------------------------
Interfaces : 1
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The stream is received on the incoming PMSI interface *mpls-if-73728* on PE-3. The properties of this PMSI interface are as follows:

```
[/]
A:admin@PE-3# show router 1 pim tunnel-interface "mpls-if-73728" detail

===============================================================================
PIM Interface ipv4 mpls-if-73728
===============================================================================
Admin Status       : Up                 Oper Status        : Up
IPv4 Admin Status  : Up                 IPv4 Oper Status   : Up
DR                 : 192.0.2.2
Auto-created       : No
Transport Type     : MVPN-Pmsi
-------------------------------------------------------------------------------
Interfaces : 1
===============================================================================
```

PE-3 sends this multicast stream to MTU-4, which in turn sends it to the receiver that sent the join, so the path taken by the multicast stream runs via PE-2, PE-3, and MTU-4.

In the example from Figure 167: Multicast From an EVPN-MPLS Service Into an R-VPLS With All-Active EVPN Multi-Homing, and the commands and traces that follow, PE-1 is the active IGMP querier using address 10.1.4.1, sending out the queries across the L2 domain. The group queries are sent by PE-1 to PE-3 across the EVPN-MPLS tunnel because PE-3 is DF for *esi-13*, then forwarded onto MTU-4 to reach the (potential) receiver. MTU-4 relays the IGMP responses from the receiver to one of the links; in this example, the link between MTU-4 and PE-1. When the IGMP response for joining the 225.1.70.1 stream is received on PE-1, this event is signaled across the EVPN-MPLS tunnel because it is received over esi-13. This way, the IGMP state is synchronized between PE-3 and PE-1 in a data-driven way.

The basic IGMP snooping state for VPLS 10 on PE-1 and PE-3 is as follows. The output shows that IGMP snooping is enabled on ports sap:lag-1:10, rvpls, and evpn-mpls.

```
[/]
A:admin@PE-1# show service id 10 igmp-snooping base

===============================================================================
IGMP Snooping Base info for service 10
===============================================================================
Admin State : Up
Querier     : 10.1.4.1 on rvpls int-MCAST-VPLS
SBD service : N/A
Evpn-proxy  : Disabled
-------------------------------------------------------------------------------
Port                      Oper MRtr Pim  Send Max  Max  Max   MVR       Num
Id                        Stat Port Port Qrys Grps Srcs Grp   From-VPLS Grps
                                                  Srcs
-------------------------------------------------------------------------------
sap:lag-1:10              Up   No   No   No   None None None  Local     1
rvpls                     Up   Yes  No   N/A  N/A  N/A  N/A   N/A       N/A
evpn-mpls
                          Up   Yes  No   N/A  N/A  N/A  N/A   N/A       N/A
===============================================================================
```

```
[/]
A:admin@PE-3# show service id 10 igmp-snooping base

===============================================================================
IGMP Snooping Base info for service 10
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Admin State : Up
Querier     : 10.1.4.1 on evpn-mpls
SBD service : N/A
Evpn-proxy  : Disabled
-------------------------------------------------------------------------------
Port                      Oper MRtr Pim  Send Max  Max  Max   MVR       Num
Id                        Stat Port Port Qrys Grps Srcs Grp   From-VPLS Grps
                                                      Srcs
-------------------------------------------------------------------------------
sap:lag-1:10              Up   No   No   No   None None None  Local     1
rvpls                     Up   Yes  No   N/A  N/A  N/A  N/A   N/A       N/A
evpn-mpls
                          Up   Yes  No   N/A  N/A  N/A  N/A   N/A       N/A
===============================================================================
```

PE-1 sends the IGMP queries on VPRN 1 via the *int-MCAST-VPLS* interface, so the VPLS that is referenced in the *int-MCAST-VPLS* interface registers the ports on which the IGMP queries are received as multicast router ports. EVPN-MPLS tunnels are always multicast router ports. The following output displays the source addresses of the multicast routers:

```
[/]
A:admin@PE-1# show service id 10 igmp-snooping mrouters

===============================================================================
IGMP Snooping Multicast Routers for service 10
===============================================================================
MRouter          Port Id                   Up Time        Expires   Version
-------------------------------------------------------------------------------
10.1.4.1         rvpls                     0d 00:21:54    231s      3
-------------------------------------------------------------------------------
Number of mrouters: 1
===============================================================================
```

```
[/]
A:admin@PE-3# show service id 10 igmp-snooping mrouters

===============================================================================
IGMP Snooping Multicast Routers for service 10
===============================================================================
MRouter          Port Id                   Up Time        Expires   Version
-------------------------------------------------------------------------------
10.1.4.1         evpn-mpls                 0d 00:21:25    229s      3
-------------------------------------------------------------------------------
Number of mrouters: 1
===============================================================================
```

The IGMP snooping querier properties for VPLS 10 on PE-1 and PE-3 are as follows:

```
[/]
A:admin@PE-1# show service id 10 igmp-snooping querier

===============================================================================
IGMP Snooping Querier info for service 10
===============================================================================
Port Id                : r-vpls int-MCAST-VPLS
IP Address             : 10.1.4.1
Expires                : 130s
Up Time                : 0d 00:21:30
Version                : 3

General Query Interval  : 125s
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Query Response Interval : 10.0s
Robust Count            : 2
===============================================================================

[/]
A:admin@PE-3# show service id 10 igmp-snooping querier

===============================================================================
IGMP Snooping Querier info for service 10
===============================================================================
Port Id                 : evpn-mpls
IP Address              : 10.1.4.1
Expires                 : 253s
Up Time                 : 0d 00:21:01
Version                 : 3

General Query Interval  : 125s
Query Response Interval : 10.0s
Robust Count            : 2
===============================================================================
```

IGMP snooping in VPLS 10 registers the reports in the IGMP snooper port database (port-db). The port-db can be displayed with a show command, and specifying a SAP limits the output generated by this command, as follows:

```
[/]
A:admin@PE-1# show service id 10 igmp-snooping port-db sap lag-1:10

===============================================================================
IGMP Snooping SAP lag-1:10 Port-DB for service 10
===============================================================================
Group Address   Mode    Type    From-VPLS  Up Time        Expires  Num   MC
                                                                    Src   Stdby
-------------------------------------------------------------------------------
225.70.1.1      exclude dynamic local      0d 00:04:10    never    0
-------------------------------------------------------------------------------
Number of groups: 1
===============================================================================

[/]
A:admin@PE-3# show service id 10 igmp-snooping port-db sap lag-1:10

===============================================================================
IGMP Snooping SAP lag-1:10 Port-DB for service 10
===============================================================================
Group Address   Mode    Type    From-VPLS  Up Time        Expires  Num   MC
                                                                    Src   Stdby
-------------------------------------------------------------------------------
225.70.1.1      exclude dynamic local      0d 00:04:12    230s     0
-------------------------------------------------------------------------------
Number of groups: 1
===============================================================================
```

IGMP snooping statistics show the number of received, transmitted, and forwarded IGMP messages per type, and also provide drop counts per error type, as follows:

```
[/]
A:admin@PE-1# show service id 10 igmp-snooping statistics

===============================================================================
IGMP Snooping Statistics for service 10
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Message Type            Received      Transmitted   Forwarded
-------------------------------------------------------------------------------
General Queries         0             0             17
Group Queries           0             2             2
Group-Source Queries    0             0             0
V1 Reports              0             0             0
V2 Reports              0             0             0
V3 Reports              12            6             6
V2 Leaves               0             0             0
Unknown Type            0             N/A           0
EVPN SMET Routes        0             0             N/A
-------------------------------------------------------------------------------
Drop Statistics
-------------------------------------------------------------------------------
Bad Length               : 0
Bad IP Checksum          : 0
Bad IGMP Checksum        : 0
Bad Encoding             : 0
No Router Alert          : 0
Zero Source IP           : 0
Wrong Version            : 0
Lcl-Scope Packets        : 0
Rsvd-Scope Packets       : 0

Send Query Cfg Drops     : 0
Import Policy Drops      : 0
Exceeded Max Num Groups  : 0
Exceeded Max Num Sources : 0
Exceeded Max Num Grp Srcs: 0
MCAC Policy Drops        : 0
MCS Failures             : 0

MVR From VPLS Cfg Drops  : 0
MVR To SAP Cfg Drops     : 0
===============================================================================

[/]
A:admin@PE-3# show service id 10 igmp-snooping statistics

===============================================================================
IGMP Snooping Statistics for service 10
===============================================================================
Message Type            Received      Transmitted   Forwarded
-------------------------------------------------------------------------------
General Queries         12            0             9
Group Queries           2             2             0
Group-Source Queries    0             0             0
V1 Reports              0             0             0
V2 Reports              0             0             0
V3 Reports              12            6             0
V2 Leaves               0             0             0
Unknown Type            0             N/A           0
EVPN SMET Routes        0             0             N/A
-------------------------------------------------------------------------------
Drop Statistics
-------------------------------------------------------------------------------
Bad Length               : 0
Bad IP Checksum          : 0
Bad IGMP Checksum        : 0
Bad Encoding             : 0
No Router Alert          : 0
Zero Source IP           : 0
Wrong Version            : 0
Lcl-Scope Packets        : 0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Rsvd-Scope Packets       : 0

Send Query Cfg Drops     : 0
Import Policy Drops      : 0
Exceeded Max Num Groups  : 0
Exceeded Max Num Sources : 0
Exceeded Max Num Grp Srcs: 0
MCAC Policy Drops        : 0
MCS Failures             : 0

MVR From VPLS Cfg Drops  : 0
MVR To SAP Cfg Drops     : 0
===============================================================================
```

Debug

Debugging is useful for troubleshooting purposes, and the debug configuration used on PE-1 and PE-3 for checking IGMP and IGMP snooping functionalities is as follows:

```
debug {
    router "VPRN 1" {
        igmp {
            packet {
                dropped true
                ingress true
                egress true
            }
        }
    }
    service {
        vpls "mcast-vpls" {
            igmp-snooping {
                packet {
                    dropped true
                    ingress true
                    egress true
                    evpn-mpls true
                    sap lag-1:10 { }
                }
            }
        }
    }
}
```

When group 225.70.1.1 is joined, the trace on PE-1 is as follows. Event 4 is the IGMPv3 join message for group 225.70.1.1 received on SAP lag-1:10 in VPLS 10 from the receiver. The reception of this message is synchronized across the EVPN-MPLS tunnel for VPLS 10, as indicated by event 5. Event 6 is the IGMPv3 join message as received on interface *int-MCAST-VPLS* by VPRN 1.

```
4 2023/07/20 11:42:52.720 CEST MINOR: DEBUG #2001 Base IGMP
"IGMP: RX packet on svc 10
  from chaddr 04:0f:ff:00:01:41
  Port  : sap lag-1:10
  SrcIp : 0.0.0.0
  DstIp : 224.0.0.22
  Type  : V3 REPORT
    Num Group Records: 1
        Group Record Type: CHG_TO_EXCL (4), AuxDataLen 0, Num Sources 0
          Group Addr: 225.70.1.1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
"

5 2023/07/20 11:42:52.720 CEST MINOR: DEBUG #2001 Base IGMP
"IGMP: TX packet on svc 10
  from chaddr 5e:00:00:16:04:0f
  send towards ES : esi-13
  Port  : evpn-mpls
  SrcIp : 0.0.0.0
  DstIp : 224.0.0.22
  Type  : V3 REPORT
    Num Group Records: 1
        Group Record Type: CHG_TO_EXCL (4), AuxDataLen 0, Num Sources 0
          Group Addr: 225.70.1.1

"

---snip---

6 2023/07/20 11:42:52.720 CEST MINOR: DEBUG #2001 vprn1 IGMP[2]
"IGMP[2]: RX-PKT
[000 00:28:13.480] IGMP interface int-MCAST-VPLS [ifIndex 4] V3 PDU: 0.0.0.0 -> 224.0.0.22 pdu
Len 16
    Type: V3 REPORT maxrespCode 0x0 checkSum 0xf7b6
    Num Group Records: 1
        Group Record 0
        Type: CHG_TO_EXCL, AuxDataLen 0, Num Sources 0
        Mcast Addr: 225.70.1.1
        Source Address List

"
```

The trace on PE-3 is as follows. Event 5 is the reception of the snooping state synchronization across the EVPN-MPSL tunnel, and event 8 is the IGMPv3 join as received on interface *int-MCAST-VPLS* by VPRN 1.

```
5 2023/07/20 11:42:52.726 CEST MINOR: DEBUG #2001 Base IGMP
"IGMP: RX packet on svc 10
  from chaddr 04:0f:ff:00:01:41
  received via evpn-mpls on ES : esi-13
  Port  : sap lag-1:10
  SrcIp : 0.0.0.0
  DstIp : 224.0.0.22
  Type  : V3 REPORT
    Num Group Records: 1
        Group Record Type: CHG_TO_EXCL (4), AuxDataLen 0, Num Sources 0
          Group Addr: 225.70.1.1

"
---snip---

8 2023/07/20 11:42:52.726 CEST MINOR: DEBUG #2001 vprn1 IGMP[2]
"IGMP[2]: RX-PKT
[000 00:28:07.620] IGMP interface int-MCAST-VPLS [ifIndex 4] V3 PDU: 0.0.0.0 -> 224.0.0.22 pdu
Len 16
    Type: V3 REPORT maxrespCode 0x0 checkSum 0xf7b6
    Num Group Records: 1
        Group Record 0
        Type: CHG_TO_EXCL, AuxDataLen 0, Num Sources 0
        Mcast Addr: 225.70.1.1
        Source Address List
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
"
```

Similar events are logged when the multicast receiver leaves the 225.70.1.1 group.

Conclusion

By connecting customers to EVPN-MPLS VPRN/IES routed services via an R-VPLS, service providers can offer IPv4 multicast services to customers in an all-active multi-homing scenario.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

L2 Services with Auto-GRE Spoke-SDPs

This chapter provides information about L2 Services with Auto-GRE Spoke-SDPs.

Topics in this chapter include:

Applicability

This chapter was initially written for SR OS Release 16.0.R4, but the MD-CLI in the current edition is based on SR OS Release 21.5.R1. Auto-GRE spoke-SDPs are supported in L2 services in SR OS Release 16.0.R1, and later.

Overview

When the connectivity between nodes is IP-based (not MPLS), VPWS and VPLS services can use manually provisioned or auto-generated GRE transport tunnels. For auto-GRE transport tunnels, the signaling can be BGP or Targeted LDP (T-LDP). BGP signaling is more scalable than T-LDP, because T-LDP requires point-to-point sessions between communicating peers.

Auto-GRE spoke-SDPs can be used in the following services:

- BGP-VPLS with BGP signaling
- LDP VPLS using BGP-AD with T-LDP signaling
- BGP-VPWS with BGP signaling
- Dynamic Multi-segment Pseudowire (MS-PW) spoke-SDP Forwarding Equivalence Class (FEC) 129 with T-LDP signaling (*not* supported in MD-CLI for SR OS Release 21.5.R1)

PW templates for auto-GRE spoke-SDPs are configured with **auto-gre-sdp true**.

```
*[ex:/configure service pw-template "PW3"]
A:admin@PE-1# auto-gre-sdp ?

 auto-gre-sdp <boolean>
 <boolean>  - ([true]|false)
 Default    - false

 'auto-gre-sdp' is: immutable

    Use a GRE tunnel to automatically create an SDP

    Warning: Modifying this element recreates 'configure service pw-template "PW3"'
    automatically for the new value to take effect.

 Immutable fields      - pw-template-id, provisioned-sdp, auto-gre-sdp
---snip---
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The **auto-gre-sdp** parameter can be combined with the parameter **provisioned-sdp prefer**, but not with **provisioned-sdp use** (because that might contradict the use of auto-GRE spoke-SDPs), as follows:

```
*[ex:/configure service pw-template "PW3"]
A:admin@PE-1# commit
MINOR: SVCMGR #5626: configure service pw-template "PW3" auto-gre-sdp - not compatible with
 auto-gre-sdp - auto-gre-sdp is not allowed with used-provisioned-sdp
```

The auto-GRE SDP and SDP binding are created after a matching BGP route has been received. Subsequent requests for an auto-GRE SDP of the same type and to the same destination as an existing auto-GRE SDP will use the existing auto-GRE SDP.

Downstream fragmentation is allowed for auto-GRE SDPs by clearing the Don't Fragment (DF) bit in the GRE IP header. The following command controls fragmentation for a PW template:

```
configure {
    service {
        pw-template "PW40" {
            pw-template-id 40
            allow-fragmentation true
            auto-gre-sdp true
        }
```

The following PW template parameters are not supported with GRE tunnels and will be ignored when a GRE SDP is auto-created:

- Hash label
- Entropy label
- SDP include/exclude (there is no mechanism to configure an SDP admin group for auto-GRE SDPs)

However, these parameters are relevant for provisioned MPLS SDPs when the PW template is configured with **provisioned-sdp prefer**.

The **pw-template-binding** parameter in the **bgp <..>** context of the L2 service allows to configure the PW template to be used. It is possible to define multiple PW template bindings within a service. The mechanism for selecting the PW template is as follows:

- In BGP-VPWS, BGP-VPLS, and BGP-AD services, the PW template binding selection is based on matching the configured import Route Targets (RTs) for a PW template binding with the RTs in the received routes.
- The binding with the first matching RT is chosen. If no import RTs are configured, the lowest PW template binding ID is used.
- It is not possible to add RTs to BGP-VPWS BGP updates using import or export policies, because they are ignored. However, the RT exported to select the destination service can be used on the receiving PE with PW template binding statement to influence the PW template to be selected; see the first use case in the Configuration section.
- If the selected PW template is configured with **provisioned-sdp prefer** and an SDP with a matching far-end address exists, the system chooses the SDP with the lowest metric from the tunnel table. If multiple matching SDPs with the same metric occur, the highest SDP ID that is operationally up is chosen.

The following **tools** command allows for PW template bindings to change:

```
[/]
A:admin@PE-1# tools perform service id 1 eval-pw-template
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
[policy-id] <number>
<number>  - <1..2147483647>

   <number> - <1..2147483647>
```

The policy ID refers to the PW template currently in use. With the **allow-service-impact** option, the current binding will be torn down and re-signaled.

Configuration

Figure 168: Example topology shows the example topology with three PEs in AS 64500. Services will be configured on PE-1 and PE-2, and PE-3 is the route reflector (RR).

Figure 168: Example topology

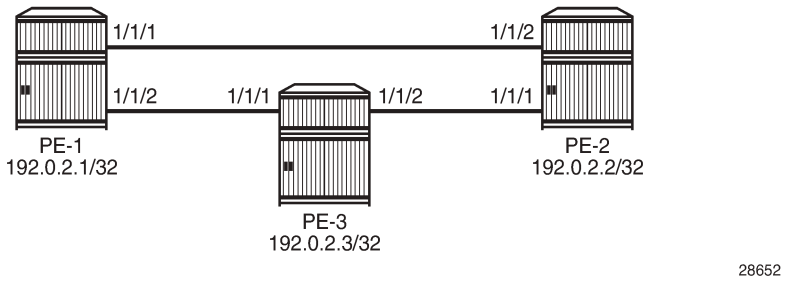

The initial configuration on the three PEs includes:

- Cards, MDAs, ports
- Router interfaces
- IS-IS as IGP (alternatively, OSPF can be used)

Auto-GRE spoke-SDPs are configured in the following use cases:

1. BGP-VPLS with BGP signaling
2. BGP-AD in VPLS with T-LDP signaling
3. BGP-VPWS with BGP signaling

In these three use cases (BGP-VPLS, BGP-AD, BGP-VPWS), BGP is configured for the L2-VPN address family. In each of the use cases, two L2 services will be configured using different PW templates with **auto-gre-sdp**: one with **provisioned-sdp prefer** and one without.

Auto-GRE spoke-SDPs in BGP-VPLS

Figure 169: BGP-VPLS with auto-GRE spoke-SDPs shows the example topology with BGP-VPLSs 1 and 2 configured on PE-1 and PE-2. BGP is configured for the L2-VPN address family with PE-3 as Route Reflector (RR). The CEs are emulated through VPRNs configured on the PEs and connected to the VPLSs via Port Cross-connect (PXC).

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 169: BGP-VPLS with auto-GRE spoke-SDPs

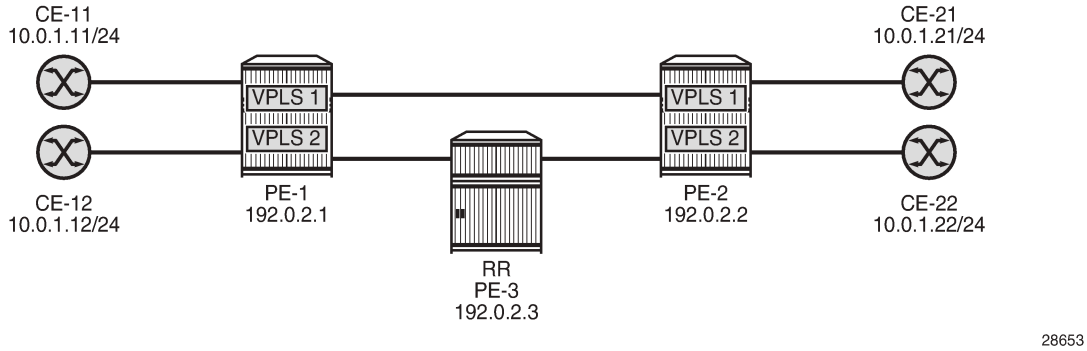

BGP configuration

For the BGP-VPLS, BGP-AD, and BGP-VPWS use cases, BGP is configured with the L2-VPN address family. The BGP configuration on PE-1 and PE-2 is identical, as follows:

```
# on PE-1, PE-2::
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            rapid-withdrawal true
            split-horizon true
            group "WAN" {
                type internal
                family {
                    l2-vpn true
                }
            }
            neighbor "192.0.2.3" {
                group "WAN"
            }
```

On RR PE-3, BGP is configured as follows:

```
# on RR PE-3:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            rapid-withdrawal true
            split-horizon true
            group "WAN" {
                type internal
                family {
                    l2-vpn true
                }
                cluster {
                    cluster-id 192.0.2.3
                }
            }
            neighbor "192.0.2.1" {
                group "WAN"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
            neighbor "192.0.2.2" {
                group "WAN"
            }
```

Service configuration

The configuration of BGP-VPLS services is described in the BGP VPLS chapter.

PW template 10 is configured with **auto-gre-sdp**; PW template 20 is configured with **provisioned-sdp prefer** and **auto-gre-sdp**. Because only IP connectivity is present between the nodes (no MPLS), the provisioned SDP is GRE-based using BGP signaling (no T-LDP). VPLS 1 has PW template bindings with IDs 10 and 20; VPLS 2 is configured with PW template binding 20. The service configuration on PE-1 is as follows:

```
# on PE-1:
configure {
    service {
        pw-template "PW10-auto-GRE" {
            pw-template-id 10
            auto-gre-sdp true
        }
        pw-template "PW20-auto-GRE_prefer-prov" {
            pw-template-id 20
            provisioned-sdp prefer
            auto-gre-sdp true
        }
        sdp 12 {
            admin-state enable
            signaling bgp
            far-end {
                ip-address 192.0.2.2
            }
        }
        vpls "BGP-VPLS-1" {
            admin-state enable
            description "BGP-VPLS with auto-GRE spoke-SDP"
            service-id 1
            customer "1"
            bgp 1 {
                route-distinguisher "64500:1"
                route-target {
                    export "target:64500:1"
                    import "target:64500:1"
                }
                pw-template-binding "PW10-auto-GRE" {
                }
                pw-template-binding "PW20-auto-GRE_prefer-prov" {
                }
            }
            bgp-vpls {
                admin-state enable
                maximum-ve-id 100
                ve {
                    name "PE-1"
                    id 1
                }
            }
            sap pxc-10.a:1 {        # SAP to connect to CE-11
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
        vpls "BGP-VPLS-2" {
            admin-state enable
            description "BGP-VPLS with auto-GRE spoke-SDP_prefer provisioned SDP"
            service-id 2
            customer "1"
            bgp 1 {
                route-distinguisher "64500:2"
                route-target {
                    export "target:64500:2"
                    import "target:64500:2"
                }
                pw-template-binding "PW20-auto-GRE_prefer-prov" {
                }
            }
            bgp-vpls {
                admin-state enable
                maximum-ve-id 100
                ve {
                    name "PE-1"
                    id 1
                }
            }
            sap pxc-10.a:2 {        # SAP to connect to CE-12
            }
        }
```

The service configuration on PE-2 is similar, but the VE name is "PE-2" and the VE ID equals 2 instead, as follows:

```
# on PE-2:
configure {
    service {
        pw-template "PW10-auto-GRE" {
            pw-template-id 10
            auto-gre-sdp true
        }
        pw-template "PW20-auto-GRE_prefer-prov" {
            pw-template-id 20
            provisioned-sdp prefer
            auto-gre-sdp true
        }
        sdp 21 {
            admin-state enable
            signaling bgp
            far-end {
                ip-address 192.0.2.1
            }
        }
        vpls "BGP-VPLS-1" {
            admin-state enable
            description "BGP-VPLS with auto-GRE spoke-SDP"
            service-id 1
            customer "1"
            bgp 1 {
                route-distinguisher "64500:1"
                route-target {
                    export "target:64500:1"
                    import "target:64500:1"
                }
                pw-template-binding "PW10-auto-GRE" {
                }
                pw-template-binding "PW20-auto-GRE_prefer-prov" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
            }
            bgp-vpls {
                admin-state enable
                maximum-ve-id 100
                ve {
                    name "PE-2"
                    id 2
                }
            }
            sap pxc-10.a:1 {        # SAP to connect to CE-21
            }
        }
        vpls "BGP-VPLS-2" {
            admin-state enable
            description "BGP-VPLS with auto-GRE spoke-SDP_prefer provisioned SDP"
            service-id 2
            customer "1"
            bgp 1 {
                route-distinguisher "64500:2"
                route-target {
                    export "target:64500:2"
                    import "target:64500:2"
                }
                pw-template-binding "PW20-auto-GRE_prefer-prov" {
                }
            }
            bgp-vpls {
                admin-state enable
                maximum-ve-id 100
                ve {
                    name "PE-2"
                    id 2
                }
            }
            sap pxc-10.a:2 {        # SAP to connect to CE-22
            }
        }
```

The following L2-VPN routes are received on PE-1: one for VPLS 1 with RD 64500:1 and another for VPLS 2 with RD 64500:2.

```
[/]
A:admin@PE-1# show router bgp routes l2-vpn
===============================================================================
 BGP Router ID:192.0.2.1        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP L2VPN Routes
===============================================================================
Flag  RouteType                  Prefix                            MED
      RD                         SiteId                            Label
      Nexthop                    VeId                  BlockSize   LocalPref
      As-Path                    BaseOffset            vplsLabelBa
                                                       se
-------------------------------------------------------------------------------
u*>i  VPLS                       -                     -           0
      64500:1                    -                                 -
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
      192.0.2.2                 2                     8          100
      No As-Path                1                     524280
u*>i  VPLS                      -                     -          0
      64500:2                   -                     -
      192.0.2.2                 2                     8          100
      No As-Path                1                     524272
-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

VPLS 1 is configured with two PW template bindings without import RT. Because the PW template binding with the lowest ID is preferred, PW template 10 is used and therefore, the following GRE SDP 32767 is auto-created:

```
[/]
A:admin@PE-1# show service id 1 sdp detail

===============================================================================
Services: Service Destination Points Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 32767:4294967295  -(192.0.2.2)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id             : 32767:4294967295         Type              : BgpVpls
PW-Template Id     : 10
Split Horiz Grp    : (Not Specified)
Etree Root Leaf Tag: Disabled                 Etree Leaf AC     : Disabled
VC Type            : Ether                    VC Tag            : n/a
Admin Path MTU     : 0                        Oper Path MTU     : 8954
Delivery           : GRE
Far End            : 192.0.2.2                Tunnel Far End    : n/a
Oper Tunnel Far End: 192.0.2.2
---snip---

Admin State        : Up                       Oper State        : Up
MinReqd SdpOperMTU : 1514
Acct. Pol          : None                     Collect Stats     : Disabled
Ingress Label      : 524281                   Egress Label      : 524280
---snip---

Last Status Change : 06/23/2021 14:24:54      Signaling         : BGP
---snip---
```

VPLS 2 is configured with PW template binding 20, which prefers provisioned SDPs, so the provisioned SDP 12 is used, as follows:

```
[/]
A:admin@PE-1# show service id 2 sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
12:4294967294    BgpVpls  192.0.2.2       Up      Up        524273    524272
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

In VPLS 1, the PW template binding selection can be changed by configuring a non-matching import RT to PW template 10, as follows:

```
# on PE-1:
configure {
    service {
        vpls "BGP-VPLS-1" {
            bgp 1 {
                pw-template-binding "PW10-auto-GRE" {
                    import-rt ["target:64500:999"]
                }
            }
```

This does not change the selected PW template during service operation and PW template 10 remains in use, as follows:

```
[/]
A:admin@PE-1# show service id 1 sdp detail | match "PW-Template"
PW-Template Id     : 10
```

The following **tools** command forces the system to re-evaluate the PW template binding:

```
[/]
A:admin@PE-1# tools perform service id 1 eval-pw-template 10 allow-service-impact
eval-pw-template succeeded for Svc 1 32767:4294967295 Policy 10
```

When the PW template binding is re-evaluated, PW template binding 20 is selected and the provisioned SDP 12 is used, as follows:

```
[/]
A:admin@PE-1# show service id 1 sdp detail | match "PW-Template"
PW-Template Id     : 20
```

```
[/]
A:admin@PE-1# show service id 1 sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
12:4294967293    BgpVpls  192.0.2.2       Up      Up        524281    524280
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
```

Auto-GRE spoke-SDPs in LDP-VPLS using BGP-AD

Figure 170: LDP-VPLS using BGP-AD with auto-GRE Spoke-SDPs shows the example topology with VPLSs 3 and 4 configured with BGP-AD on PE-1 and PE-2. The BGP configuration is identical to the one for BGP-VPLS.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 170: LDP-VPLS using BGP-AD with auto-GRE Spoke-SDPs

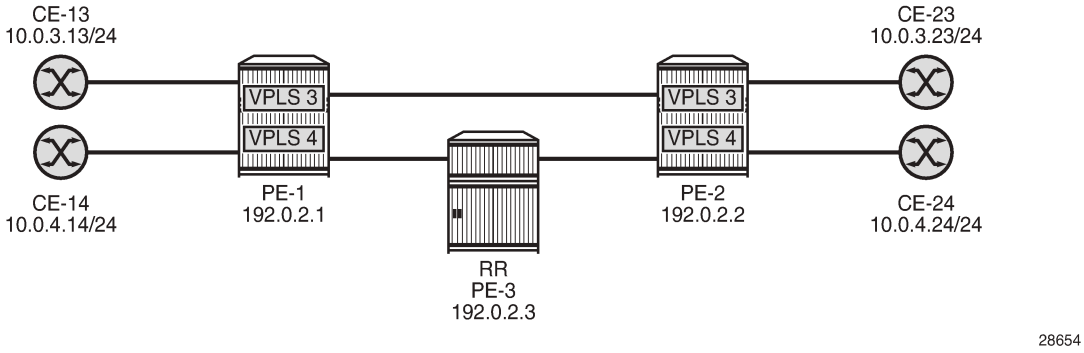

The following T-LDP session is configured between PE-1 and PE-2:

```
# on PE-1:
configure {
    router "Base" {
        ldp {
            targeted-session {
                peer 192.0.2.2 {
                }
            }
```

```
# on PE-2:
configure {
    router "Base" {
        ldp {
            targeted-session {
                peer 192.0.2.1 {
                }
            }
```

The following T-LDP signaled SDP is configured on PE-1 and PE-2:

```
# on PE-1:
configure {
    service {
        sdp 120 {
            admin-state enable
            far-end {
                ip-address 192.0.2.2
            }
        }
```

```
# on PE-2:
configure {
    service {
        sdp 210 {
            admin-state enable
            far-end {
                ip-address 192.0.2.1
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The service configuration on PE-1 and PE-2 is as follows; see chapter LDP VPLS Using BGP Auto-Discovery for a description of BGP-AD in LDP VPLS. PW templates 10 and 20 are the same as in the preceding example.

```
# on PE-1, PE-2:
configure {
    service {
        pw-template "PW10-auto-GRE" {
            pw-template-id 10
            auto-gre-sdp true
        }
        pw-template "PW20-auto-GRE_prefer-prov" {
            pw-template-id 20
            provisioned-sdp prefer
            auto-gre-sdp true
        }
        vpls "BGP-AD VPLS-3" {
            admin-state enable
            description "BGP-AD for LDP VPLS with auto-GRE spoke-SDP"
            service-id 3
            customer "1"
            bgp 1 {
                route-distinguisher "64500:3"
                route-target {
                    export "target:64500:3"
                    import "target:64500:3"
                }
                pw-template-binding "PW10-auto-GRE" {
                }
                pw-template-binding "PW20-auto-GRE_prefer-prov" {
                }
            }
            bgp-ad {
                admin-state enable
                vpls-id "64500:3"
            }
            sap pxc-10.a:3 {          # SAP to connect to CE-13 (PE-1) or CE-23 (PE-2)
            }
        }
        vpls "BGP-AD VPLS-4" {
            admin-state enable
            description "BGP-AD for LDP VPLS with auto-GRE spoke-SDP pref-prov-SDP"
            service-id 4
            customer "1"
            bgp 1 {
                route-distinguisher "64500:4"
                route-target {
                    export "target:64500:4"
                    import "target:64500:4"
                }
                pw-template-binding "PW20-auto-GRE_prefer-prov" {
                }
            }
            bgp-ad {
                admin-state enable
                vpls-id "64500:4"
            }
            sap pxc-10.a:4 {         # SAP to connect to CE-14 (PE-1) or CE-24 (PE-2)
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

PE-1 has received the following L2-VPN BGP-AD routes:

```
[/]
A:admin@PE-1# show router bgp routes l2-vpn bgp-ad
===============================================================================
 BGP Router ID:192.0.2.1        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP L2VPN-AD Routes
===============================================================================
Flag  RouteType                  Prefix                            MED
      RD                         SiteId                            Label
      Nexthop                    VeId                  BlockSize   LocalPref
      As-Path                    BaseOffset            vplsLabelBa
                                                       se
-------------------------------------------------------------------------------
u*>i  AutoDiscovery              192.0.2.2             -           0
      64500:3                    -                                 -
      192.0.2.2                  -                     -           100
      No As-Path                 -                     -
u*>i  AutoDiscovery              192.0.2.2             -           0
      64500:4                    -                                 -
      192.0.2.2                  -                     -           100
      No As-Path                 -                     -
-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

The following shows the used SDPs on PE-1: BGP-signaled SDP 12 (used by VPLS 1 and 2) and T-LDP-signaled SDPs 120 and 32767.

```
[/]
A:admin@PE-1# show service sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId  AdmMTU  OprMTU  Far End          Adm  Opr         Del     LSP   Sig
-------------------------------------------------------------------------------
12     0       8954    192.0.2.2        Up   Up          GRE     n/a   BGP
120    0       8954    192.0.2.2        Up   Up          GRE     n/a   TLDP
32767  0       8954    192.0.2.2        Up   Up          GRE     n/a   TLDP
-------------------------------------------------------------------------------
Number of SDPs : 3
-------------------------------------------------------------------------------
Legend: R = RSVP, L = LDP, B = BGP, M = MPLS-TP, n/a = Not Applicable
        I = SR-ISIS, O = SR-OSPF, T = SR-TE, F = FPE
===============================================================================
```

The following shows that PW template 10 is used in VPLS 3 and that auto-GRE SDP 32767 is used, with T-LDP signaling:

```
[/]
A:admin@PE-1# show service id 3 sdp detail

===============================================================================
Services: Service Destination Points Details
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 32767:4294967292  -(192.0.2.2)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id            : 32767:4294967292         Type              : BgpAd
PW-Template Id    : 10
AGI               : 64500:3                  SDP Bind Source   : bgp-l2vpn
Local AII         : 192.0.2.1
Remote AII        : 192.0.2.2
Split Horiz Grp   : (Not Specified)
Etree Root Leaf Tag: Disabled                Etree Leaf AC     : Disabled
VC Type           : Ether                    VC Tag            : n/a
Admin Path MTU    : 0                        Oper Path MTU     : 8954
Delivery          : GRE
Far End           : 192.0.2.2                Tunnel Far End    : n/a
Oper Tunnel Far End: 192.0.2.2
---snip---

Admin State       : Up                       Oper State        : Up
---snip---

Last Status Change : 06/23/2021 14:30:31     Signaling         : TLDP
---snip---
```

The following shows that the T-LDP signaled GRE SDP 120 is used in VPLS 4, not the BGP-signaled GRE SDP 12:

```
[/]
A:admin@PE-1# show service id 4 sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
120:4294967291   BgpAd    192.0.2.2       Up      Up        524269    524269
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
```

Auto-GRE spoke-SDPs in BGP-VPWS

Figure 171: BGP-VPWS with auto-GRE spoke-SDPs shows the example topology with BGP-VPWS Epipes 5 and 6 on PE-1 and PE-2. The BGP configuration is identical to the one for BGP-VPLS.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 171: BGP-VPWS with auto-GRE spoke-SDPs

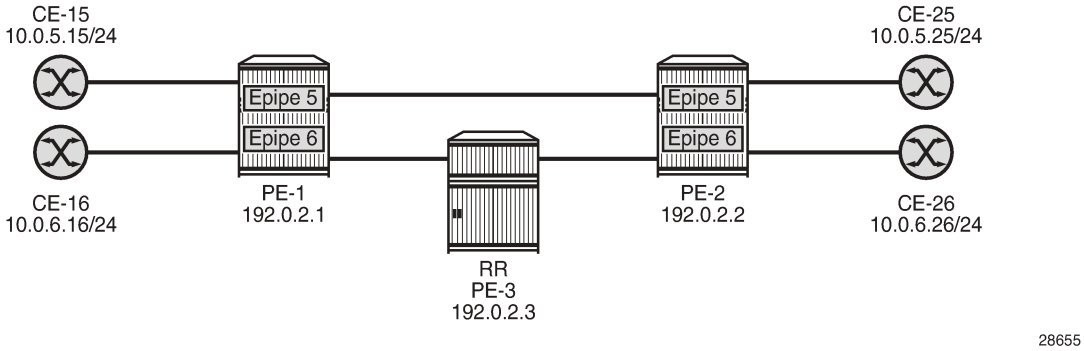

Chapter BGP Virtual Private Wire Services describes the configuration of BGP VPWS. The configuration of Epipes 5 and 6 on PE-1 is as follows:

```
# on PE-1:
configure {
    service {
        pw-template "PW10-auto-GRE" {
            pw-template-id 10
            auto-gre-sdp true
        }
        pw-template "PW20-auto-GRE_prefer-prov" {
            pw-template-id 20
            provisioned-sdp prefer
            auto-gre-sdp true
        }
        epipe "BGP-VPWS-5" {
            admin-state enable
            description "BGP-VPWS with auto-GRE spoke-SDP"
            service-id 5
            customer "1"
            bgp 1 {
                route-distinguisher "64500:5"
                route-target {
                    export "target:64500:5"
                    import "target:64500:5"
                }
                pw-template-binding "PW10-auto-GRE" {
                }
                pw-template-binding "PW20-auto-GRE_prefer-prov" {
                }
            }
            bgp-vpws {
                admin-state enable
                local-ve {
                    name "PE-1"
                    id 1
                }
                remote-ve "PE-2" {
                    id 2
                }
            }
            sap pxc-10.a:5 {        # SAP to connect to CE-15
            }
        }
        epipe "BGP-VPWS-6" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            admin-state enable
            description "BGP-VPWS with auto-GRE spoke-SDP_prefer provisioned SDP"
            service-id 6
            customer "1"
            bgp 1 {
                route-distinguisher "64500:6"
                route-target {
                    export "target:64500:6"
                    import "target:64500:6"
                }
                pw-template-binding "PW20-auto-GRE_prefer-prov" {
                }
            }
            bgp-vpws {
                admin-state enable
                local-ve {
                    name "PE-1"
                    id 1
                }
                remote-ve "PE-2" {
                    id 2
                }
            }
            sap pxc-10.a:6 {        # SAP to connect to CE-16
            }
        }
```

The configuration of the Epipes is similar on PE-2, but the VE names and VE IDs are different, as follows:

```
# on PE-2:
configure {
    service {
        pw-template "PW10-auto-GRE" {
            pw-template-id 10
            auto-gre-sdp true
        }
        pw-template "PW20-auto-GRE_prefer-prov" {
            pw-template-id 20
            provisioned-sdp prefer
            auto-gre-sdp true
        }
        epipe "BGP-VPWS-5" {
            admin-state enable
            description "BGP-VPWS with auto-GRE spoke-SDP"
            service-id 5
            customer "1"
            bgp 1 {
                route-distinguisher "64500:5"
                route-target {
                    export "target:64500:5"
                    import "target:64500:5"
                }
                pw-template-binding "PW10-auto-GRE" {
                }
                pw-template-binding "PW20-auto-GRE_prefer-prov" {
                }
            }
            bgp-vpws {
                admin-state enable
                local-ve {
                    name "PE-2"
                    id 2
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                remote-ve "PE-1" {
                    id 1
                }
            }
            sap pxc-10.a:5 {        # SAP to connect to CE-25
            }
        }
        epipe "BGP-VPWS-6" {
            admin-state enable
            description "BGP-VPWS with auto-GRE spoke-SDP_prefer provisioned SDP"
            service-id 6
            customer "1"
            bgp 1 {
                route-distinguisher "64500:6"
                route-target {
                    export "target:64500:6"
                    import "target:64500:6"
                }
                pw-template-binding "PW20-auto-GRE_prefer-prov" {
                }
            }
            bgp-vpws {
                admin-state enable
                local-ve {
                    name "PE-2"
                    id 2
                }
                remote-ve "PE-1" {
                    id 1
                }
            }
            sap pxc-10.a:6 {        # SAP to connect to CE-26
            }
        }
```

PE-1 receives the following BGP-VPWS routes from PE-2:

```
[/]
A:admin@PE-1# show router bgp routes l2-vpn bgp-vpws
===============================================================================
 BGP Router ID:192.0.2.1        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP L2VPN-VPWS Routes
===============================================================================
Flag  RouteType                  Prefix                           MED
      RD                         SiteId                           Label
      Nexthop                    VeId                  BlockSize  LocalPref
      As-Path                    BaseOffset            vplsLabelBa
                                                       se
-------------------------------------------------------------------------------
u*>i  VPWS                       -                     -          0
      64500:5                    -                                -
      192.0.2.2                  2                     1          100
      No As-Path                 1                     524268
u*>i  VPWS                       -                     -          0
      64500:6                    -                                -
      192.0.2.2                  2                     1          100
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
      No As-Path                  1                       524267
-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

The following SDP bindings are used on PE-1: the first two are used by BGP-VPLS services VPLS 1 and 2, the third and fourth are used by BGP-AD in LDP VPLS 3 and 4, and the last two are used by BGP-VPWS services Epipe 5 and 6. For the last two, SDP 32766 is auto-created, whereas SDP 12 is provisioned with BGP signaling.

```
[/]
A:admin@PE-1# show service sdp-using

===============================================================================
SDP Using
===============================================================================
SvcId      SdpId              Type   Far End              Opr   I.Label E.Label
                                                         State
-------------------------------------------------------------------------------
1          12:4294967293      BgpVp* 192.0.2.2            Up    524281  524280
2          12:4294967294      BgpVp* 192.0.2.2            Up    524273  524272
3          32767:4294967292   BgpAd  192.0.2.2            Up    524270  524270
4          120:4294967291     BgpAd  192.0.2.2            Up    524269  524269
5          32766:4294967290   BgpVp* 192.0.2.2            Up    524268  524268
6          12:4294967289      BgpVp* 192.0.2.2            Up    524267  524267
-------------------------------------------------------------------------------
Number of SDPs : 6
-------------------------------------------------------------------------------
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

Epipe 5 uses the following auto-GRE SDP 32766 with BGP signaling:

```
[/]
A:admin@PE-1# show service id 5 sdp detail

===============================================================================
Services: Service Destination Points Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 32766:4294967290  -(192.0.2.2)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id             : 32766:4294967290         Type             : BgpVpws
PW-Template Id     : 10
VC Type            : Ether                    VC Tag           : n/a
Admin Path MTU     : 0                        Oper Path MTU    : 8954
Delivery           : GRE
Far End            : 192.0.2.2                Tunnel Far End   : n/a
Oper Tunnel Far End: 192.0.2.2
---snip---

Admin State        : Up                       Oper State       : Up
---snip---

Last Status Change : 06/23/2021 14:36:00      Signaling        : BGP
---snip---
```

PW template 20 is used in Epipe 6, so the BGP-signaled GRE SDP 12 is used, as follows:

```
[/]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
A:admin@PE-1# show service id 6 sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
12:4294967289    BgpVpws  192.0.2.2       Up      Up        524267    524267
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
```

Conclusion

In IP-based networks, auto-GRE spoke-SDPs can be used in VPWS and VPLS services. Manually configured GRE tunnels are not an option in networks — such as LTE networks — where it is common to assign IP addresses dynamically from a pool of addresses, but auto-GRE spoke-SDPs can be applied instead.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Layer 2 Multicast Optimization for EVPN-VXLAN — Assisted Replication

This chapter provides information about Layer 2 Multicast Optimization for EVPN-VXLAN — Assisted Replication.

Topics in this chapter include:

Applicability

This chapter was initially written for SR OS Release 14.0.R4, but the CLI in the current edition is based on SR OS Release 23.3.R3. Layer 2 multicast optimization for EVPN-VXLAN - Assisted Replication (AR) is supported in SR OS Release 14.0.R4, and later.

Overview

Typically, EVPN-VXLAN can use either Ingress Replication (IR) or Protocol Independent Multicast (PIM) for Broadcast, Unknown unicast, and Multicast (BUM) traffic (although SR OS does not support PIM along with EVPN-VXLAN). PIM requires keeping multicast state awareness per subnet per tenant in the core routers, which may not scale. Not all core routers support PIM.

IR inefficiency is usually tolerable in EVPN networks for broadcast and unknown unicast traffic; however, it is not tolerable for multicast traffic:

- Broadcast traffic can be reduced by the proxy-ARP and proxy-ND capabilities supported by EVPN.
- Unknown unicast traffic is greatly reduced in virtualized Data Center (DC) networks where all MAC and IP addresses are learned in the control or management planes. In such cases, unknown MAC addresses are always outside the DC. An **unknown-mac-route** can be enabled to ensure that the unknown unicast traffic is sent only to the DC gateway, which minimizes flooding within the DC.
- Multicast traffic may be an issue for the hypervisors holding the multicast sources, because the hypervisors need to replicate the multicast traffic to the remote VXLAN Tunnel Endpoints (VTEPs). The multicast replication at the hypervisors is a software process and the throughput can be heavily impacted. This is also true when VPLS services are used in the Virtual Service Router (VSR) and many replicas must be done from the VSR. Using a dedicated service node to replicate the multicast traffic on behalf of the hypervisors can help, but the replication capabilities of such service nodes are limited too.

SR OS supports the Assisted Replication (AR) feature for IPv4 VXLAN tunnels (both replicator and leaf functions) in compliance with the non-selective mode described in *draft-ietf-bess-evpn-optimized-ir*. AR is a Layer 2 multicast optimization feature that helps software-based PEs and Network Virtualization Edge (NVE) devices with low-performance replication capabilities to deliver Broadcast and Multicast (BM) Layer 2 traffic to remote VTEPs in the VPLS.

SR OS nodes support the AR-Replicator (AR-R) and AR-Leaf (AR-L) functions, although not simultaneously on the same service. Nodes configured as AR-L select an AR-R within a service and send

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

all BM packets to this AR-R. AR-Rs replicate traffic to all the VTEPs in the VPLS on behalf of the AR-Ls, so BM traffic is delivered to all VPLS participants without any packet loss caused by performance issues. Unknown unicast packets follow the same path as known unicast packets to avoid packet reordering. Therefore, no AR-R is used for unknown unicast traffic.

When multiple AR-Rs exist in a service, the AR-L performs per-service load-balancing of the BM traffic. The AR-L lists the candidate AR-Rs, ordered by IP address and VXLAN Network Identifier (VNI); candidate 0 having the lowest IP address and VNI. The replicator is selected using a modulo function of the service ID and the number of candidate AR-Rs. For example, assume that VPLS 1 has two candidate AR-Rs: because 1 modulo 2 equals 1, the second AR-R in the list is selected. In case of failure, a new AR-R is selected. If there are no more AR-Rs, the system falls back to IR.

Figure 172: PMSI Tunnel Attribute - Flags shows an EVPN route-type 3, an Inclusive Multicast Ethernet Tag (IMET) route containing a PMSI tunnel attribute with a flags octet. Flag L was already defined in RFC 6514. *Draft-ietf-bess-evpn-optimized-ir* defines additional flags: type, BM, and U. The BM and U flags are used for Pruned Flood Lists (PFL) signaling and they are not supported.

Figure 172: PMSI Tunnel Attribute - Flags

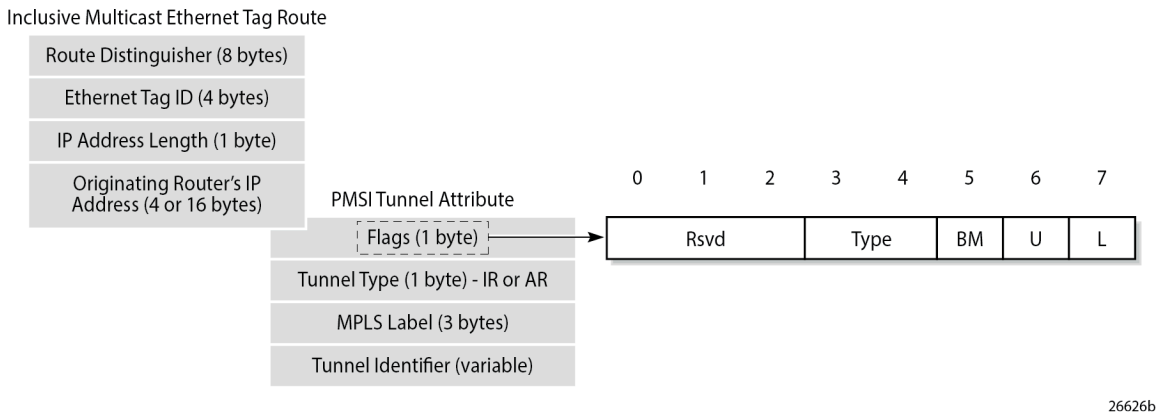

The type field has two bits that define the AR role of the advertising router, as follows:

- Type 00 = Regular Network Virtualization Edge (RNVE) - indicates that AR is not supported and IR is applied instead (for backward compatibility)
- Type 01= AR-R
- Type 10 = AR-L
- Type 11 = reserved

The tunnel type in the PMSI tunnel attribute can be configured with the following options for IR and AR:

- Tunnel type 0x06 = (non-optimized) IR, sent by AR-R and AR-L if **ingress-repl-inc-mcast-advertisement** is enabled, which is the default option
- Tunnel type 0x0A = type AR, originated by AR-R

For regular IR routes, the originating router's IP address equals the system IP address. The MPLS label and tunnel identifier must be used as described in RFC 7432. The tunnel identifier is set to a routable address of the PE.

For AR routes, the originating router's IP address and the tunnel identifier are both set to the AR IP address (AR-IP) configured in the **service system vxlan** context. The AR-IP must be previously defined as a loopback interface address in the base router and must be different from the IR IP address (IR-IP).

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Note:
If the AR-IP loopback interface is down, the router does not withdraw the AR route. However, the remote AR-Ls is not able to resolve the AR route's BGP next-hop if the AR-IP is no longer propagated in the IGP.

Figure 173: EVPN Assisted Replication for VXLAN shows the example topology with the multicast source connected to a hypervisor PE-3 that acts as AR-L, which sends an IR route containing the system address of PE-3. The AR-R PE-1 sends an AR route that uses AR-IPs instead of IR-IPs; for example, PE-1 has AR-IP 1.1.1.1 and IR-IP 192.0.2.1.

Figure 173: EVPN Assisted Replication for VXLAN

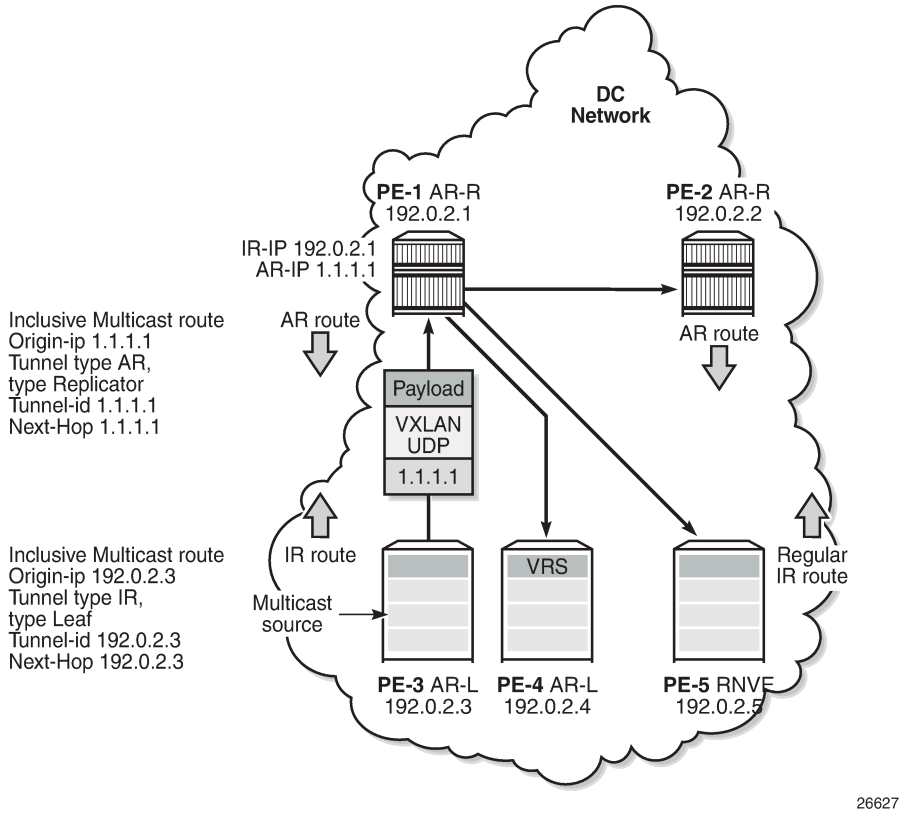

Hypervisor PE-3 sends the BM traffic to the AR-R, which replicates it to all the VTEPs in the VPLS, except to PE-3.

Table 10: Inclusive multicast route information sent by different AR roles shows the inclusive multicast route information sent by each role in an AR-capable service.

Table 10: Inclusive multicast route information sent by different AR roles

AR role	function	inclusive multicast route advertised
AR-R	assists AR-Ls	IR inclusive multicast route (tunnel = 0x06 = IR, IR-IP, type = 0 = none)

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

AR role	function	inclusive multicast route advertised
		AR inclusive multicast route (tunnel = 0x0A = AR, AR-IP, type = 1 = AR-R)
AR-L	sends BM only to AR-R	IR inclusive multicast route (tunnel = 0x06 = IR, IR- IP, type = 2 = AR-L)
RNVE	non-AR support	IR inclusive multicast route (tunnel = 0x06 = IR, IR- IP, type = 0 = none)

Unicast traffic (known or unknown) is processed as normal. For BM traffic, the AR-R uses AR or IR based on the IP destination address (DA):

- If IP DA equals the AR-IP, the AR-R replicates to the VTEPs in the VXLAN service, except for the VTEP over which the BM traffic was received.
- If IP DA equals the IR-IP, normal IR forwarding is done.

Non-optimized-IR nodes are unaware of the PMSI tunnel attribute flag definition with the additional flags for AR, so they ignore the information in the flags field.

The *draft-ietf-bess-evpn-optimized-ir* describes the following three types of IR optimizations:

- Non-selective AR - the chosen AR-R replicates the BM traffic to all NVEs in the Ethernet VPN Instance (EVI) except for the source NVE.
- Selective AR - AR-Rs replicate BM traffic to only their AR-L set and the rest of the AR-Rs. Selective AR allows a "multi-stage" AR replication, as opposed to a "single-stage" AR replication.
- Pruned Flood Lists - AR-Ls can signal PFL flags to be pruned from the flood lists for BM or for unknown unicast traffic. PFL may be used in combination with AR.

This chapter only describes non-selective AR.

Configure AR-R and AR-L

The AR-IP is configured on the AR-R, as follows:

```
configure {
    service {
        system {
            vxlan {
                assisted-replication {
                    ip-address ?

 ip-address <unicast-ipv4-address>
 <unicast-ipv4-address> - <d.d.d.d>

    IP address for assisted replication in the router
```

The AR-IP is the IPv4 address of a loopback interface in the base router instance. When attempting to configure an AR-IP and the loopback address does not exist, the following error message is raised:

```
configure {
    service {
        system {
            vxlan {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                assisted-replication {
                    ip-address 1.1.1.1

MINOR: MGMT_CORE #4001: configure service system vxlan assisted-replication ip-address
 - loopback interface with address (max prefix) needed for assisted replication
 - configure router "Base"
```

The AR types replicator and leaf are configured in a VPLS with the following command:

```
configure {
    service {
        vpls "VPLS 10" {
            vxlan {
                instance 1 {
                    vni 1
                    assisted-replication ?

 assisted-replication

 Choice: role
 leaf                :+ Enter the leaf context
 replicator          :- AR role as replicator
```

When attempting to configure an AR-R before the AR-IP is set, the following error is raised:

```
configure {
    service {
        vpls "VPLS 10" {
            customer "1"
            service-id 10
            vxlan {
                instance 1 {
                    vni 1
                    assisted-replication {
                        replicator
MINOR: MGMT_CORE #4001: configure service vpls "VPLS 10" vxlan instance 1 assisted-replication
 replicator
 - assisted-replication ip address needed for replication role
 - configure service system vxlan assisted-replication ip-address
```

The assisted-replication-time can only be configured on leaf nodes. The following error is raised after an attempt to configure the assisted-replication-time on an AR-R:

```
configure {
    service {
        vpls "VPLS 10" {
            vxlan {
                instance 1 {
                    assisted-replication {
                        replicator {
                            acttime 5
MINOR: MGMT_CORE #2201: Unknown element - 'acttime'
```

The **acttime** can optionally be activated, and works as follows. When the router creates an AR-R destination for the first time, the assisted replication time must expire before this AR-R destination is eligible as candidate AR-R to forward BM traffic. Upon expiration, the router runs the AR-R selection (service ID modulo the number of AR-Rs provides the selected AR-R in the ordered list of candidate AR-Rs). The AR-R EVPN destination is created as "BM" and the destinations to the remaining nodes is shown as "U".

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The **acttime** allows the AR-R some time to program the leaf VTEPs in the following cases:

- Configuration of a new AR-R
- AR-R rebooting
- AR-R going operationally down and up again

If the timer is zero (default value), the AR-R may receive packets from a VTEP that has not been programmed yet, in which case the AR-R drops the packets.

With the AR-Rs and AR-Ls configured, IMET AR routes can be exchanged. IR can be enabled or disabled independently of the AR configuration. The following command is required to enable IR inclusive multicast routes, and is enabled by default:

```
configure {
    service {
        vpls "VPLS 10" {
            bgp-evpn {
                routes {
                    incl-mcast {
                        advertise-ingress-replication true
```

BGP-EVPN routes

By default, IR is enabled in BGP-EVPN. The following IMET IR route is sent from PE-5 (RNVE) to Route Reflector (RR) PE-1. The flags in the PMSI Tunnel Attribute (PTA) indicate that regular IR is used to forward BUM traffic (tunnel type: 0x06). The AR type is "None", because AR is disabled on PE-5. The IR-IP 192.0.2.5 is used as next-hop, originator IP address, and tunnel endpoint. The MPLS label corresponds to the VNI.

```
A:admin@PE-5# //
A:PE-5# show debug
debug
    router "Base"
        bgp
            update
        exit
    exit
exit
A:PE-5# //
```

```
On PE-5:
14 2023/07/12 10:59:55.416 CEST MINOR: DEBUG #2001 Base Peer 1: 192.0.2.1
"Peer 1: 192.0.2.1: UPDATE
Peer 1: 192.0.2.1 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 77
    Flag: 0x90 Type: 14 Len: 28 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.5
        Type: EVPN-INCL-MCAST Len: 17 RD: 192.0.2.5:1, tag: 0, orig_addr len: 32, orig_addr:
 192.0.2.5
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:1
        bgp-tunnel-encap:VXLAN
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    Flag: 0xc0 Type: 22 Len: 9 PMSI:
        Tunnel-type Ingress Replication (6)
        Flags: (0x0)[Type: None BM: 0 U: 0 Leaf: not required]
        MPLS Label 1
        Tunnel-Endpoint 192.0.2.5
"
```

A similar IMET IR route is sent from AR-L PE-3 toward RR PE-1, as follows. The difference is that the flags indicate that PE-3 is configured as an AR-L for the VPLS. The IR-IP 192.0.2.3 is used as next-hop, originator address, and tunnel endpoint.

```
On PE-3:
10 2023/07/12 10:58:29.634 CEST MINOR: DEBUG #2001 Base Peer 1: 192.0.2.1
"Peer 1: 192.0.2.1: UPDATE
Peer 1: 192.0.2.1 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 77
    Flag: 0x90 Type: 14 Len: 28 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.3
        Type: EVPN-INCL-MCAST Len: 17 RD: 192.0.2.3:1, tag: 0, orig_addr len: 32, orig_addr:
 192.0.2.3
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:1
        bgp-tunnel-encap:VXLAN
    Flag: 0xc0 Type: 22 Len: 9 PMSI:
        Tunnel-type Ingress Replication (6)
        Flags: (0x10)[Type: AR Leaf BM: 0 U: 0 Leaf: not required]
        MPLS Label 1
        Tunnel-Endpoint 192.0.2.3
"
```

The IMET IR routes contain the system IP addresses of the nodes, not the AR-IPs.

The following AR route is advertised from AR-R PE-1. The tunnel type is AR and the flags indicate that PE-1 is configured as AR-R. The AR-IP 1.1.1.1 is the next-hop address, the originator address, and the tunnel endpoint.

```
On PE-1:
4 2023/07/12 10:55:15.069 CEST MINOR: DEBUG #2001 Base Peer 1: 192.0.2.4
"Peer 1: 192.0.2.4: UPDATE
Peer 1: 192.0.2.4 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 77
    Flag: 0x90 Type: 14 Len: 28 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 1.1.1.1
        Type: EVPN-INCL-MCAST Len: 17 RD: 192.0.2.1:1, tag: 0, orig_addr len: 32, orig_addr:
 1.1.1.1
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:1
        bgp-tunnel-encap:VXLAN
    Flag: 0xc0 Type: 22 Len: 9 PMSI:
        Tunnel-type Assisted Replication (10)
        Flags: (0x8)[Type: AR Replicator BM: 0 U: 0 Leaf: not required]
        MPLS Label 1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        Tunnel-Endpoint 1.1.1.1
"
```

Besides IMET AR routes, PE-1 may also advertise IMET IR routes to the other nodes using IR-IP 192.0.2.1 (system IP address). By default, BGP-EVPN has IR enabled. For example, the following IMET IR route is advertised to PE-4:

```
On PE-1:
3 2023/07/12 10:55:15.069 CEST MINOR: DEBUG #2001 Base Peer 1: 192.0.2.4
"Peer 1: 192.0.2.4: UPDATE
Peer 1: 192.0.2.4 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 77
    Flag: 0x90 Type: 14 Len: 28 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.1
        Type: EVPN-INCL-MCAST Len: 17 RD: 192.0.2.1:1, tag: 0, orig_addr len: 32, orig_addr:
 192.0.2.1
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:1
        bgp-tunnel-encap:VXLAN
    Flag: 0xc0 Type: 22 Len: 9 PMSI:
        Tunnel-type Ingress Replication (6)
        Flags: (0x0)[Type: None BM: 0 U: 0 Leaf: not required]
        MPLS Label 1
        Tunnel-Endpoint 192.0.2.1
"
```

The following IMET routes have been received by PE-4:

```
[/]
A:admin@PE-4# show router bgp routes evpn incl-mcast
===============================================================================
 BGP Router ID:192.0.2.4        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Inclusive-Mcast Routes
===============================================================================
Flag  Route Dist.         OrigAddr
      Tag                 NextHop
-------------------------------------------------------------------------------
u*>i  192.0.2.1:1         1.1.1.1
      0                   1.1.1.1

u*>i  192.0.2.1:1         192.0.2.1
      0                   192.0.2.1

u*>i  192.0.2.2:1         2.2.2.2
      0                   2.2.2.2

u*>i  192.0.2.2:1         192.0.2.2
      0                   192.0.2.2

u*>i  192.0.2.3:1         192.0.2.3
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
      0                   192.0.2.3

u*>i  192.0.2.5:1         192.0.2.5
      0                   192.0.2.5

-------------------------------------------------------------------------------
Routes : 6
===============================================================================
```

Configuration

Figure 174: Example topology shows the example topology with PE-1 and PE-2 as AR-R nodes, PE-3 and PE-4 as AR-L nodes, and PE-5 as RNVE node. The multicast source is connected to PE-3, which is a low-performance node. PE-1 acts as an RR for all nodes.

Figure 174: Example topology

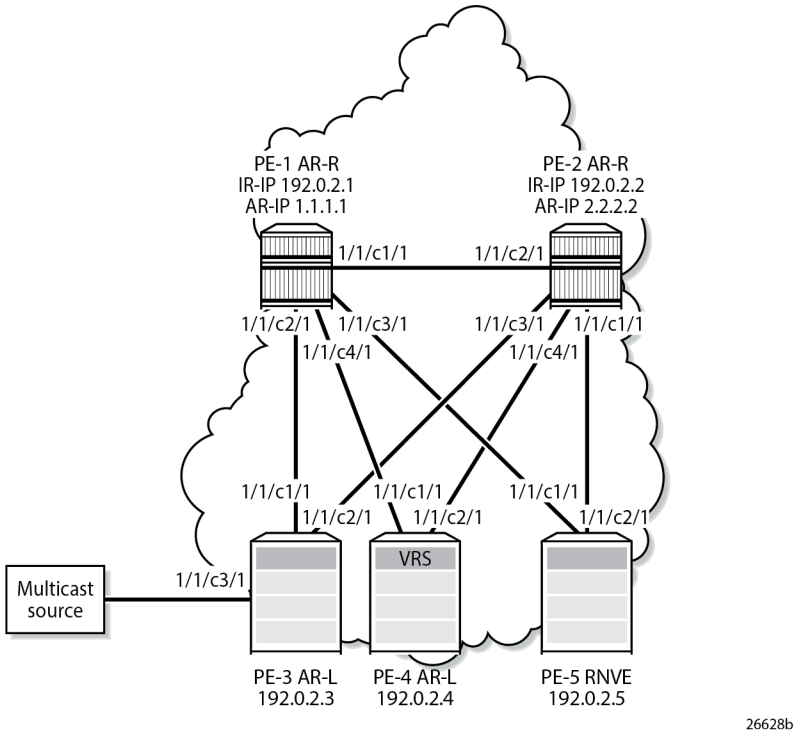

The initial configuration on the nodes includes:

- Cards, MDAs, ports
- Router interfaces between the nodes
- IS-IS as IGP (alternatively, OSPF can be used)

BGP is configured for address family EVPN with RR PE-1. The BGP configuration on PE-1 is as follows:

```
On PE-1:
configure {
    router "Base" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            split-horizon true
            ebgp-default-reject-policy {
                import false
                export false
            }
            rapid-update {
                evpn true
            }
            group "DC" {
                peer-as 64500
                family {
                    evpn true
                }
                cluster {
                    cluster-id 192.0.2.1
                }
            }
            neighbor "192.0.2.2" {
                group "DC"
            }
            neighbor "192.0.2.3" {
                group "DC"
            }
            neighbor "192.0.2.4" {
                group "DC"
            }
            neighbor "192.0.2.5" {
                group "DC"
            }
        }
```

The BGP configuration on the other nodes is as follows:

```
On the other PEs:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            split-horizon true
            ebgp-default-reject-policy {
                import false
                export false
            }
            rapid-update {
                evpn true
            }
            group "DC" {
                peer-as 64500
                family {
                    evpn true
                }
            }
            neighbor "192.0.2.1" {
                group "DC"
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
```

VPLS 10 is configured on all nodes. PE-1 is configured as AR-R with AR-IP 1.1.1.1, which must be configured as loopback IPv4 address in the base router and as AR-IP that can be shared between services. When attempting to configure an AR-IP with an IP address that does not exist in the base router, the following error is raised:

```
configure {
    service {
        system {
            vxlan {
                assisted-replication {
                    ip-address 1.1.1.1
MINOR: MGMT_CORE #4001: configure service system vxlan assisted-replication ip-address
 - loopback interface with address (max prefix) needed for assisted replication
 - configure router "Base"
```

First, a loopback interface is configured in the base router. The IP address needs to be routable and, in this example, an export policy exporting this IP address is configured in IS-IS. Alternatively, a static route can be configured or an additional IS-IS passive interface can be configured for the loopback interface. The IP address is then configured as AR-IP in the **service system vxlan** context. PE-1 is configured as AR-R for VPLS 10, as follows:

```
On PE-1:
configure {
    policy-options {
        prefix-list "AR-IP" {
            prefix 1.1.1.1/32 type exact {
            }
        }
        policy-statement "export_AR-IP" {
            entry 10 {
                from {
                    prefix-list ["AR-IP"]
                }
                action {
                    action-type accept
                }
            }
        }
    }
    router "Base" {
        interface "AR-IP" {
            loopback
            ipv4 {
                primary {
                    address 1.1.1.1
                    prefix-length 32
                }
            }
        }
        isis 0 {
            export-policy ["export_AR-IP"]
        }
    }
    service {
        system {
            vxlan {
                assisted-replication {
                    ip-address 1.1.1.1
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
        }
        vpls "VPLS 10" {
            customer "1"
            service-id 10
            admin-state enable
            vxlan {
                instance 1 {
                    vni 1
                    assisted-replication {
                        replicator
                    }
                }
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 1
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
        }
    }
```

The configuration is similar on PE-2, but with AR-IP 2.2.2.2 instead of 1.1.1.1.

PE-3 and PE-4 are configured as AR-L nodes for VPLS 10. No AR-IP needs to be configured. The configuration of VPLS 10 on PE-3 is as follows:

```
On PE-3:
configure {
    service {
        vpls "VPLS 10" {
            admin-state enable
            service-id 10
            customer "1"
            vxlan {
                instance 1 {
                    vni 1
                    assisted-replication {
                        leaf { }
                    }
                }
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 1
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
            sap 1/1/c3/1 {    # sap for ingress traffic from STC
            }
            sap 1/2/c1/1:1 {    # sap for egress traffic to VPLS 10
            }
```

Multicast traffic enters SAP 1/1/c3/1, whereas receiving hosts can be connected to other SAPs, such as SAP 1/2/c1/1:1. The configuration of VPLS 10 on PE-4 is similar, but no multicast source is connected.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

When a node is configured as AR-L, optionally the **acttime** can be configured to define the waiting time before the leaf can begin sending multicast traffic to a new replicator or a replicator that was rebooted. The default is zero seconds, in which case the AR-L starts sending packets to the AR-R without delay. Nokia recommends configuring a **acttime** value different from zero.

```
configure {
    service {
        vpls "VPLS 10" {
            vxlan {
                instance 1 {
                    vni 1
                    assisted-replication {
                        leaf {
                            acttime ?

 acttime <number>
 <number> - <1..255> - seconds

    Time for the leaf to wait before sending traffic to a new replicator
```

PE-5 is configured as an RNVE node for VPLS 10, as follows:

```
On PE-5:
configure {
    service {
        vpls "VPLS 10" {
            admin-state enable
            service-id 10
            customer "1"
            vxlan {
                instance 1 {
                    vni 1
                }
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 1
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
            sap 1/2/c1/1:1 {    # sap for egress traffic to VPLS 10
            }
```

BGP-EVPN IMET routes are exchanged between the nodes. The following IMET routes are used on AR-L PE-3, with two routes from each AR-R: one IR route with BGP next-hop 192.0.2.x and one AR route with BGP next-hop x.x.x.x (with x equal to 1 or 2).

```
[/]
A:admin@PE-3# show router bgp routes evpn incl-mcast
===============================================================================
 BGP Router ID:192.0.2.3        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
BGP EVPN Inclusive-Mcast Routes
===============================================================================
Flag  Route Dist.         OrigAddr
      Tag                 NextHop
-------------------------------------------------------------------------------
u*>i  192.0.2.1:1         1.1.1.1
      0                   1.1.1.1

u*>i  192.0.2.1:1         192.0.2.1
      0                   192.0.2.1

u*>i  192.0.2.2:1         2.2.2.2
      0                   2.2.2.2

u*>i  192.0.2.2:1         192.0.2.2
      0                   192.0.2.2

u*>i  192.0.2.4:1         192.0.2.4
      0                   192.0.2.4

u*>i  192.0.2.5:1         192.0.2.5
      0                   192.0.2.5

-------------------------------------------------------------------------------
Routes : 6
===============================================================================
```

When the AR-R has no local attachment circuits, such as SAPs or SDP-bindings, it should not generate regular IR routes. This can be controlled by disabling **advertise-ingress-replication** on PE-1 and PE-2, as follows:

```
On PE-1 and PE-2:
configure {
    service {
        vpls "VPLS 10" {
            bgp-evpn {
                routes {
                    incl-mcast {
                        advertise-ingress-replication false
```

When IR is disabled on the AR-Rs, no IR routes are sent to the other nodes and PE-3 only sees the AR routes from PE-1 and PE-2, as follows:

```
[/]
A:admin@PE-3# show router bgp routes evpn incl-mcast
===============================================================================
 BGP Router ID:192.0.2.3        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Inclusive-Mcast Routes
===============================================================================
Flag  Route Dist.         OrigAddr
      Tag                 NextHop
-------------------------------------------------------------------------------
u*>i  192.0.2.1:1         1.1.1.1
      0                   1.1.1.1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
u*>i  192.0.2.2:1         2.2.2.2
      0                   2.2.2.2

u*>i  192.0.2.4:1         192.0.2.4
      0                   192.0.2.4

u*>i  192.0.2.5:1         192.0.2.5
      0                   192.0.2.5

-------------------------------------------------------------------------------
Routes : 4
===============================================================================
```

The detailed information about the AR route sent by AR-R PE-1 can be shown with the following command. The AR tunnel has endpoint 1.1.1.1.

```
[/]
A:admin@PE-3# show router bgp routes evpn incl-mcast rd 192.0.2.1:1 hunt
===============================================================================
 BGP Router ID:192.0.2.3        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Inclusive-Mcast Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : n/a
Nexthop        : 1.1.1.1
Path Id        : None
From           : 192.0.2.1
---snip---
Community      : target:64500:1 bgp-tunnel-encap:VXLAN
Cluster        : No Cluster Members
Originator Id  : None                    Peer Router Id : 192.0.2.1
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
EVPN type      : INCL-MCAST
Tag            : 0
Originator IP  : 1.1.1.1
Route Dist.    : 192.0.2.1:1
Route Tag      : 0
---snip---
-------------------------------------------------------------------------------
PMSI Tunnel Attributes :
Tunnel-type    : Assisted Replication
Flags          : Type: AR-Replicator(1) BM: 0 U: 0 Leaf: not required
MPLS Label     : VNI 1
Tunnel-Endpoint: 1.1.1.1
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The following command shows the VXLAN destinations for VPLS 10 on PE-3:

```
[/]
A:admin@PE-3# show service id 10 vxlan destinations

===============================================================================
Egress VTEP, VNI (Instance 1)
===============================================================================
VTEP Address                                    Egress VNI Oper  Mcast Num
                                                           State       MACs
-------------------------------------------------------------------------------
1.1.1.1                                         1          Up    BM    0
2.2.2.2                                         1          Up    -     0
192.0.2.4                                       1          Up    U     0
192.0.2.5                                       1          Up    U     0
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 4
-------------------------------------------------------------------------------
---snip---
===============================================================================
```

PE-3 is configured as AR-L and no **acttime** is defined (default). Four egress VTEPs are listed: the system IP addresses are used for IR routes and the AR-IPs are used for AR routes. All BM traffic is forwarded to AR-IP 1.1.1.1 on PE-1. The AR-R in use is selected by the modulo operation on the service ID (10). In this example, two AR-Rs are available, and the service ID modulo 2 equals zero: 10 mod 2 = 0. This is the lowest possible outcome, so the first AR-R in the ordered candidate list is used. The AR-Rs are ordered by IP and VNI, with candidate 0 the lowest IP and VNI.

```
[/]
A:admin@PE-3# show service id 10 vxlan assisted-replication replicator

===============================================================================
Vxlan AR Replicator Candidates
===============================================================================
Inst  VTEP Address           Egr VNI  In Use  In Candidate List Pending Time
-------------------------------------------------------------------------------
1     1.1.1.1                1        yes     yes               0
1     2.2.2.2                1        no      yes               0
-------------------------------------------------------------------------------
Number of entries : 2
-------------------------------------------------------------------------------
===============================================================================
```

Within a service, no load-sharing is done between the AR-Rs. However, different AR-Rs can be used for different services.

- If PE-3 were configured as AR-L in VPLS 11, the calculation would be as follows: 11 mod 2 = 1; therefore, the second AR-R in the list would be selected.
- When three AR-Rs were available for VPLS 11, the calculation would be: 11 mod 3 = 2, so the third AR-R in the list would be used.

In case different VNIs are configured for the AR-Rs, the lowest IP address is always higher in the list, even when the VNI is higher. This can be shown when the VPLS VXLAN configuration on PE-1 is modified with VNI 99 instead of VNI 1, as follows:

```
On PE-1:
configure {
    service {
        vpls "VPLS 10" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            bgp-evpn {
                delete vxlan 1
            }
            delete vxlan

configure {
    service {
        vpls "VPLS 10" {
            vxlan {
                instance 1 {
                    vni 99
                    assisted-replication {
                        replicator

configure {
    service {
        vpls "VPLS 10" {
            bgp-evpn {
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
```

The list of AR-Rs on PE-3 shows that the first entry is the VTEP with the lowest IP address (1.1.1.1), even though the VNI 99 is higher than 1:

```
[/]
A:admin@PE-3# show service id 10 vxlan assisted-replication replicator

===============================================================================
Vxlan AR Replicator Candidates
===============================================================================
Inst  VTEP Address           Egr VNI  In Use  In Candidate List Pending Time
-------------------------------------------------------------------------------
1     1.1.1.1                99       yes     yes               0
1     2.2.2.2                1        no      yes               0
-------------------------------------------------------------------------------
Number of entries : 2
-------------------------------------------------------------------------------
===============================================================================
```

Note:
If the AR-IP loopback interface is down, BGP does not withdraw the AR route. When the route to the AR-IP is signaled using IGP, the route is removed from the routing table and the AR-L selects another AR-R. However, when a static route is defined for the AR-IP, a black-hole exists when the AR-IP interface is down.

PE-5 is configured as an RNVE node that signals regular IMET IR routes and is unaware of the AR-R and AR-L roles in the EVI. RNVE nodes ignore IMET AR routes. In the example, only PE-3, PE-4, and PE-5 send IMET IR updates, so the list of VTEP addresses on PE-5 only contains PE-3 and PE-4, as follows:

```
[/]
A:admin@PE-5# show service id 10 vxlan destinations

===============================================================================
Egress VTEP, VNI (Instance 1)
===============================================================================
VTEP Address                                    Egress VNI Oper  Mcast Num
                                                           State       MACs
-------------------------------------------------------------------------------
192.0.2.3                                       1          Up    BUM   0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
192.0.2.4                                             1          Up    BUM   0
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 2
-------------------------------------------------------------------------------
---snip---
===============================================================================
```

The RNVE is unaware of AR-Rs; therefore, the list of AR-Rs is empty on PE-5:

```
[/]
A:admin@PE-5# show service id 10 vxlan assisted-replication replicator

===============================================================================
Vxlan AR Replicator Candidates
===============================================================================
Inst  VTEP Address           Egr VNI  In Use  In Candidate List Pending Time
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

Verification of multicast traffic

The multicast source connected to PE-3 generates multicast traffic. PE-3 acts as AR-L and forwards the multicast packets to AR-R PE-1. In this example topology, multicast traffic enters port 1/1/c3/1 on PE-3 and is forwarded to egress port 1/1/c1/1 toward PE-1. Port statistics are cleared and traffic is generated, then the port statistics are verified.

```
[/]
A:admin@PE-3# show port 1/1/c1/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c1/1                                  82                           8878
                                       48890                       75662990
===============================================================================

[/]
A:admin@PE-3# show port 1/1/c2/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c2/1                                  67                           7654
                                          68                           7912
===============================================================================

[/]
A:admin@PE-3# show port 1/1/c3/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Id                          Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c3/1                             48809                       73213500
                                         0                              0
===============================================================================
```

Besides the multicast traffic, IGP signaling is sent and received on the network interfaces. This explains why the counters on the network interface 1/1/c1/1 toward PE-1 show a slightly higher value than on the interface 1/1/c3/1 toward the multicast source. No multicast traffic is forwarded to PE-2, which is an AR-R candidate, but not used. AR-L PE-3 selected PE-1 for VPLS 10.

When the AR-R PE-1 receives the multicast traffic from PE-3, it forwards the traffic to PE-4 and PE-5 within the VXLAN service. The VXLAN information for VPLS 10 on PE-1 shows that PE-2 is not in the list of egress VTEPs. The reason is that PE-2 does not have any SAPs or SDP-bindings and no IMET IR route is sent by PE-2 because **advertise-ingress-replication** is disabled.

```
[/]
A:admin@PE-1# show service id 10 vxlan destinations

===============================================================================
Egress VTEP, VNI (Instance 1)
===============================================================================
VTEP Address                                    Egress VNI Oper  Mcast Num
                                                           State       MACs
-------------------------------------------------------------------------------
192.0.2.3                                       1          Up    BUM   0
192.0.2.4                                       1          Up    BUM   0
192.0.2.5                                       1          Up    BUM   0
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 3
-------------------------------------------------------------------------------
---snip---
===============================================================================
```

AR-R PE-1 receives the multicast traffic from PE-3 on port 1/1/c2/1 and forwards it to the egress ports 1/1/c3/1 toward PE-5 and 1/1/c4/1 toward PE-4, as follows. No multicast traffic needs to be forwarded to egress port 1/1/c1/1 toward PE-2. Source squelching ensures that the traffic is not sent back to the originator AR-L PE-3. PE-1 has no local SAPs or SDP-bindings.

```
[/]
A:admin@PE-1# show port 1/1/c1/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c1/1                                45                           4959
                                        45                           5077
===============================================================================

[/]
A:admin@PE-1# show port 1/1/c2/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
1/1/c2/1                              48855                          75659086
                                        44                             4823
===============================================================================

[/]
A:admin@PE-1# show port 1/1/c3/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c3/1                              47                             5055
                                   48857                          75659322
===============================================================================

[/]
A:admin@PE-1# show port 1/1/c4/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c4/1                              47                             5118
                                   48855                          75659050
===============================================================================
```

An egress AR-L or RNVE node performs regular egress BUM forwarding procedures. Packets are replicated to local SAPs or SDP-bindings, but not to VXLAN-bindings.

AR-R failure scenarios

When the AR-IP interface on the used AR-R is down for any kind of reason, the route to this AR-IP is removed from the routing table on AR-L PE-3, and PE-3 selects AR-R PE-2. To simulate an AR-R failure, the AR-IP interface on PE-1 is disabled, as follows:

```
On PE-1:
configure {
    router "Base" {
        interface "AR-IP" {
            admin-state disable
```

After a while, the routing table on PE-3 does not contain an entry for prefix 1.1.1.1/32 anymore, as follows:

```
[/]
A:admin@PE-3# show router route-table 1.1.1.1/32

===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
No. of Routes: 0
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

AR-R PE-1 is not eligible anymore when the AR-IP is not reachable. PE-2 is now selected as AR-R, so BM traffic is forwarded to PE-2. Log 99 on PE-3 shows the change in AR-R from PE-1 to PE-2, as follows:

```
On PE-3:
136 2023/07/12 11:34:57.482 CEST MINOR: SVCMGR #2090 Base
"Assisted replicator in service 10 changed to VTEP 2.2.2.2, Egress VNI 1 vxlan-instance 1."
```

The VXLAN destinations for VPLS 10 on PE-3 do not include VTEP 1.1.1.1 anymore, as follows:

```
[/]
A:admin@PE-3# show service id 10 vxlan destinations

===============================================================================
Egress VTEP, VNI (Instance 1)
===============================================================================
VTEP Address                                        Egress VNI Oper  Mcast Num
                                                               State       MACs
-------------------------------------------------------------------------------
2.2.2.2                                             1          Up    BM    0
192.0.2.4                                           1          Up    U     0
192.0.2.5                                           1          Up    U     0
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 3
-------------------------------------------------------------------------------
---snip---
===============================================================================
```

Only PE-2 is listed as AR-R for VPLS 10 on PE-3, and PE-2 is the selected AR-R for VPLS 10, as follows:

```
[/]
A:admin@PE-3# show service id 10 vxlan assisted-replication replicator

===============================================================================
Vxlan AR Replicator Candidates
===============================================================================
Inst  VTEP Address           Egr VNI  In Use  In Candidate List Pending Time
-------------------------------------------------------------------------------
1     2.2.2.2                1        yes     yes               0
-------------------------------------------------------------------------------
Number of entries : 1
-------------------------------------------------------------------------------
===============================================================================
```

Incoming multicast traffic on port 1/1/c3/1 on PE-3 is now forwarded to port 1/1/c2/1 toward PE-2, as follows:

```
[/]
A:admin@PE-3# show port 1/1/c1/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c1/1                                  61                           6793
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                                        60                              6694
===============================================================================

[/]
A:admin@PE-3# show port 1/1/c2/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c2/1                                45                              5497
                                     48855                          75660441
===============================================================================

[/]
A:admin@PE-3# show port 1/1/c3/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c3/1                             48810                          73215000
                                         0                                 0
===============================================================================
```

When the AR-IP interface on AR-R PE-2 is also disabled, no AR-R is available anymore and PE-3 reverts to IR instead.

```
On PE-2:
configure {
    router "Base" {
        interface "AR-IP" {
            admin-state disable
```

The following log 99 message on AR-L PE-3 indicates that there is no AR-R anymore (VTEP 0.0.0.0, Egress VNI 0).

```
On PE-3:
2 2023/07/12 11:38:34.902 CEST MINOR: SVCMGR #2090 Base
"Assisted replicator in service 10 changed to VTEP 0.0.0.0, Egress VNI 0 vxlan-instance 1."
```

The list of VXLAN destinations for VPLS 10 on PE-3 does not include any AR-R (VTEP 1.1.1.1 or 2.2.2.2) anymore, as follows:

```
[/]
A:admin@PE-3# show service id 10 vxlan destinations

===============================================================================
Egress VTEP, VNI (Instance 1)
===============================================================================
VTEP Address                                    Egress VNI Oper  Mcast Num
                                                           State       MACs
-------------------------------------------------------------------------------
192.0.2.4                                       1          Up    BUM   0
192.0.2.5                                       1          Up    BUM   0
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 2
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
---snip---
===============================================================================
```

```
[/]
A:admin@PE-3# show service id 10 vxlan assisted-replication replicator

===============================================================================
Vxlan AR Replicator Candidates
===============================================================================
Inst  VTEP Address           Egr VNI  In Use  In Candidate List Pending Time
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

In this case, IR is done for all BUM traffic toward PE-4 and PE-5.

Conclusion

AR uses replicators to forward broadcast and multicast traffic on behalf of less-performing nodes that are configured as AR-Ls. AR is primarily used for L2 multicast optimization in data centers, but may also be used in any network using overlay EVPN-VXLAN tunnels.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

LDP VPLS Using BGP Auto-Discovery

This chapter provides information about LDP VPLS using BGP Auto-Discovery.

Topics in this chapter include:

Applicability

This chapter was initially written for SR OS Release 9.0.R3. The MD-CLI in this edition is based on SR OS Release 20.10.R2. There are no prerequisites for this configuration.

Knowledge of BGP-auto-discovery RFC 6074 architecture and functionality, RFC 4447 Pseudo-wire set-up using label distribution protocol is assumed throughout this chapter, as well as knowledge of Multi-Protocol BGP (MP-BGP).

Overview

MPLS-based Virtual Private LAN Services (VPLS) may have many different provisioning models to allow the signaling of pseudowires between Provider Edge (PE) routers containing VPLS instances.

Network Management System (NMS) provisioning using Label Distribution Protocol (LDP) signaling is a well understood method of provisioning of Layer 2 VPLS services as described in RFC 4762. This relies on the provisioning of pseudowires between VPLS instances using LDP signaling with a common Virtual Circuit (VC) identifier within the label mapping message to instantiate pseudowires.

Border Gateway Protocol Auto-Discovery (BGP-AD), as specified in RFC 6074, is an alternative method of provisioning of Layer 2 PE routers containing VPLS service instances to those described above where PEs in a common VPLS instance are automatically discovered using BGP-AD techniques.

Each PE router advertises the presence of VPLS instances to other PE routers using defined parameters within a BGP update message.

LDP is used as the pseudowire signaling protocol and relies on the auto-discovery of VPLS endpoints to instantiate pseudowires instead of manually provisioning virtual circuits. Locally configured parameters, along with BGP learned parameters, are used to determine local and remote VPLS endpoints, which are used by LDP to signal service labels to peer routers.

Figure 175: Example topology shows the example topology with six SR OS nodes located in the same autonomous system (AS).There are three PEs and RR-6 acts as a route reflector for the AS. The PE routers are all VPLS-aware. The provider (P) routers are VPLS-unaware and do not take part in the BGP process. A full mesh VPLS between PE-1, PE-2, and PE-3 is described.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 175: Example topology

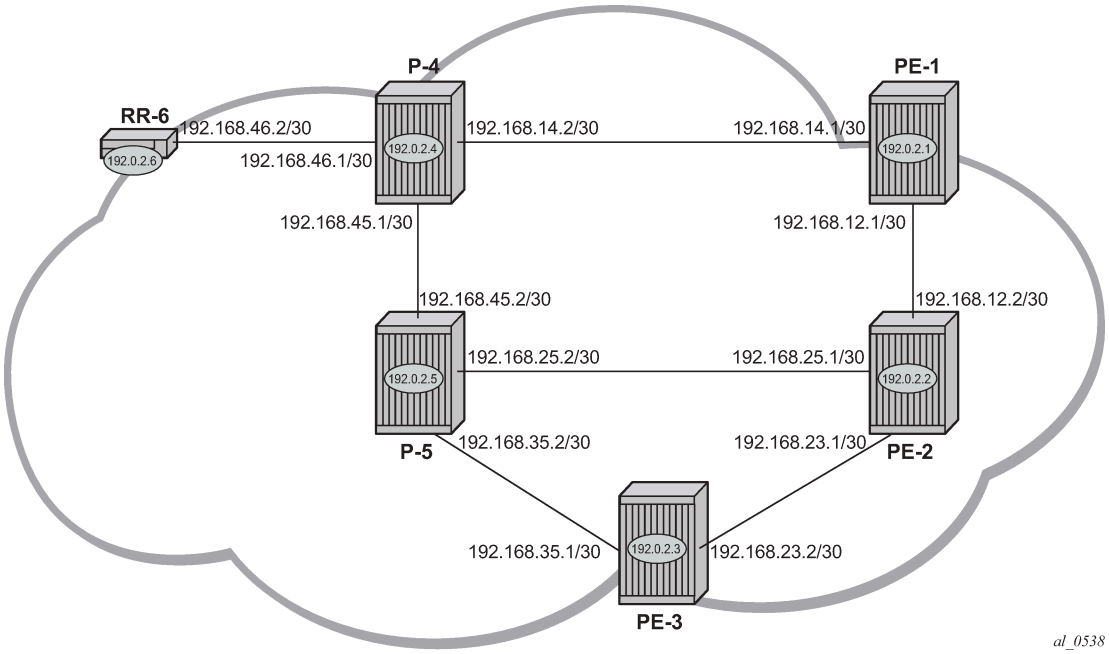

The following configuration tasks are completed as a prerequisite:

- IS-IS or OSPF is enabled on all network interfaces between each of the PE/P routers and route reflector RR-6.
- MPLS is configured on all interfaces between PE and P routers; MPLS is not required between P-4 and RR-6.
- LDP is configured on interfaces between PE and P routers; LDP is not required between P-4 and RR-6.
- The RSVP protocol must be enabled.

BGP-AD

In this architecture, a VPLS service is a collection of local VPLS instances present on a number of PEs in a provider network. In this context, VPLS-aware devices are PE routers. Each VPLS instance has a unique identifier known as the VPLS identifier (VPLS-ID). All PEs that have this VPLS instance present will have a common VPLS-ID configured.

Each VPLS instance within a PE contains a Virtual Switching Instance (VSI). The VPLS attachment circuits and pseudowires are associated with the VSI. Each VSI within a VPLS has a unique identifier called the VSI identifier (VSI-ID) and is a concatenation of the VPLS-ID plus an IP address, usually the system IP address.

The PEs communicate with each other at the control plane level by means of BGP updates containing BGP Layer 2 Network Layer Reachability Information (NLRI). Each update contains enough information for a PE to determine the presence of other local VPLS instances on peering PEs. In turn, this allows peer PE routers to set up pseudowire connectivity using LDP signaling for data flow between peers containing a local VPLS within the same VPLS instances.

Each update contains parameters usually associated with Multi-Protocol BGP updates:

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- NLRI encoded as route target (RT)—usually the VPLS-ID—and PE system address.
- Next-Hop — The system IP address of the sending PE router.
- Extended communities — Contains the RT extended community and the VPLS-ID as community values.

Each VPLS instance is configured with import and export RT extended communities to create the required pseudowire topology by controlling the distribution of each NLRI.

This chapter describes the provisioning of a VPLS instance across three PE routers. A full mesh of pseudowires interconnects the VSI of each PE within the VPLS instance. A single attachment circuit is also configured on each VSI.

Configuration

The first step is to configure an MP-iBGP session using the L2-VPN address family between each of the PEs and the RR.

The configuration on the PEs is as follows:

```
# on PE-1, PE-2, and PE-3:
configure {
    router "Base"
        autonomous-system 65536
        bgp {
            group "internal" {
                peer-as 65536
                family {
                    l2-vpn true
                }
            }
            neighbor "192.0.2.6" {
                group "internal"
            }
```

The IP addresses can be derived from Figure 175: Example topology.

The configuration for RR-6 is as follows:

```
# on RR-6:
configure {
    router "Base"
        autonomous-system 65536
        bgp {
            group "rr-internal" {
                peer-as 65536
                family {
                    l2-vpn true
                }
                cluster {
                    cluster-id 1.1.1.1
                }
            }
            neighbor "192.0.2.1" {
                group "rr-internal"
            }
            neighbor "192.0.2.2" {
                group "rr-internal"
            }
            neighbor "192.0.2.3" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                group "rr-internal"
            }
```

On PE-1, the BGP session with RR-6 is established with address family L2-VPN capability negotiated, as follows:

```
[]
A:admin@PE-1# show router bgp neighbor 192.0.2.6

===============================================================================
BGP Neighbor
===============================================================================
-------------------------------------------------------------------------------
Peer                 : 192.0.2.6
Description          : (Not Specified)
Group                : internal
-------------------------------------------------------------------------------
Peer AS              : 65536            Peer Port            : 50296
Peer Address         : 192.0.2.6
Local AS             : 65536            Local Port           : 179
Local Address        : 192.0.2.1
Peer Type            : Internal         Dynamic Peer         : No
State                : Established      Last State           : Established
Last Event           : recvOpen
Last Error           : Cease (Connection Collision Resolution)
Local Family         : L2-VPN
Remote Family        : L2-VPN
---snip---
```

On RR-6, the following BGP sessions are established with each PE for the L2-VPN address family:

```
[]
A:admin@RR-6# show router bgp summary all

===============================================================================
BGP Summary
===============================================================================
Legend : D - Dynamic Neighbor
===============================================================================
Neighbor
Description
ServiceId          AS PktRcvd InQ  Up/Down   State|Rcv/Act/Sent (Addr Family)
                      PktSent OutQ
-------------------------------------------------------------------------------
192.0.2.1
Def. Instance  65536       18    0 00h07m44s 0/0/0 (L2VPN)
                           18    0
192.0.2.2
Def. Instance  65536       18    0 00h07m44s 0/0/0 (L2VPN)
                           18    0
192.0.2.3
Def. Instance  65536       18    0 00h07m44s 0/0/0 (L2VPN)
                           18    0

-------------------------------------------------------------------------------
```

A full mesh of RSVP Label Switched Paths (LSPs) is configured between the PE routers. For reference, the MPLS interface configuration and LSPs for PE-1 to PE-2 and PE-3 is as follows:

```
# on PE-1:
configure {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    router "Base"
        mpls {
            admin-state enable
            interface "int-PE-1-P-4" {
            }
            interface "int-PE-1-PE-2" {
            }
            path "loose" {
                admin-state enable
            }
            lsp "LSP-PE-1-PE-2" {
                admin-state enable
                type p2p-rsvp
                to 192.0.2.2
                primary "loose" {
                }
            }
            lsp "LSP-PE-1-PE-3" {
                admin-state enable
                type p2p-rsvp
                to 192.0.2.3
                primary "loose" {
                }
            }
```

VPLS PE configuration

Pseudowire templates

Pseudowire templates are used by BGP to dynamically instantiate Service Distribution Point (SDP) bindings. For a given service, pseudowire templates signal the egress service de-multiplexer labels used by remote PEs to reach the local PE.

The template determines the signaling parameters of the pseudowire, control word presence, plus other usage characteristics such as Split Horizon Groups (SHGs), MAC-pinning, filters, and so on.

The MPLS transport tunnel between PE routers can be signaled using either LDP or RSVP.

LDP-based pseudowires can be automatically instantiated; RSVP-based SDPs have to be pre-provisioned.

Pseudowire templates for auto-SDP creation using LDP

In order to use an LDP transport tunnel for data flow between PEs, it is necessary for link layer LDP to be configured between all PEs/Ps so that a transport label for each PE system interface address is available. Using this mechanism, SDPs can be auto-instantiated with SDP-IDs starting at 32767. Any subsequent SDPs created use SDP-IDs decrementing from this value.

A pseudowire template is required which may contain an SHG. Each SDP created with this template is contained within the configured SHG so that traffic cannot be forwarded between them.

```
# on PE-1, PE-2, PE-3:
configure {
    service {
        pw-template "PW1" {
            pw-template-id 1
            split-horizon-group {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                name "vpls-shg"
            }
        }
```

A pseudowire template can also be created that does not contain a split horizon group. The split horizon group can then be specified when the pw-template is included within the service.

```
# on PE-1, PE-2, PE-3:
configure {
    service {
        pw-template "PW2" {
            pw-template-id 2
        }
```

Pseudowire templates for provisioned SDPs using RSVP

To use an RSVP tunnel as transport between PEs, it is necessary to bind the RSVP LSPs to the SDPs between each PE.

On PE-1, SDP 12 from PE-1 to PE-2 is configured as follows:

```
# on PE-1:
configure {
    service {
        sdp 12 {
            admin-state enable
            description "RSVP-based SDP from PE-1 to PE-2"
            delivery-type mpls
            far-end {
                ip-address 192.0.2.2
            }
            lsp "LSP-PE-1-PE-2" { }
        }
```

To create an SDP within a service that uses an RSVP transport tunnel, a pseudowire template is required that has the **provisioned-sdp use** parameter.

```
# on PE-1, PE-2, PE-3:
configure {
    service {
        pw-template "PW3" {
            pw-template-id 3
provisioned-sdp use
        }
```

Alternatively, the **provisioned-sdp prefer** parameter can be used, see chapter LDP VPLS Using BGP Auto-Discovery — Prefer Provisioned SDP.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

VPLS BGP-AD using auto-provisioned SDPs

Figure 176: VPLS instance with auto-provisioned SDPs

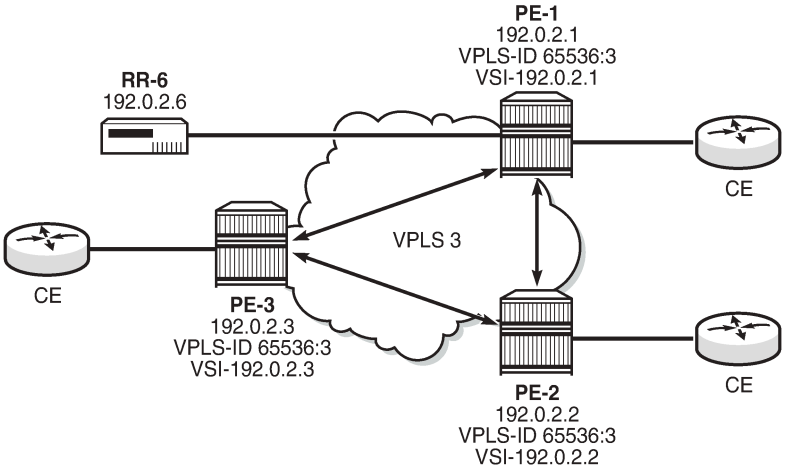

Figure 176: VPLS instance with auto-provisioned SDPs shows a schematic of a VPLS instance where the SDPs are auto-provisioned. SDPs are instantiated by a PE router using LDP signaling upon receipt of BGP Auto-Discovery (BGP-AD) updates from peer PE routers.

PE-1 configuration:

The following output shows the configuration required for a VPLS service using a pseudowire template configured for auto-provisioning of SDPs.

```
# on PE-1:
configure {
    service {
         vpls "VPLS-3" {
            admin-state enable
            service-id 3
            customer "1"
            bgp 1 {
                route-distinguisher "65536:3"
                route-target {
                    export "target:65536:3"
                    import "target:65536:3"
                }
                pw-template-binding "PW2" {
                    split-horizon-group "vpls-shg"
                    import-rt ["target:65536:3"]
                }
            }
            bgp-ad {
                admin-state enable
                vpls-id "65536:3"
                vsi-id-prefix 192.0.2.1
            }
            sap 1/1/4:3.0 {
            }
```

Within the **bgp** context, the pseudowire template is referenced which can be linked to an SHG and an import RT, if required.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Within the **bgp-ad** context, the signaling parameters are configured. These are two parameters used by each PE to determine the presence of a VPLS instance on a PE router. In turn, these are translated into endpoint identifiers for LDP signaling of pseudowires. As previously discussed, these parameters are:

- VPLS-ID — a unique identifier of the VPLS instance. Each PE that is a member of a VPLS must share the same VPLS-ID. This is inserted as an extended community value in the format AS:n. In this case, the VPLS-ID for VPLS 3 is 65536:3. This is a mandatory parameter and if it is not configured, it is not possible to enable BGP-AD (admin-state enable).
- Virtual Switching Instance (VSI) prefix — This identifies a specific instance of the VPLS. This must be unique within the VPLS instance, and is encoded using the 4 byte dotted-decimal notation. Generally, the system address is used as the VSI prefix. If this parameter is not configured, then the system address is used automatically.

The VPLS-ID and VSI prefix for VPLS 3 on each PE is shown in Figure 176: VPLS instance with auto-provisioned SDPs.

The VPLS-ID and VSI prefix are concatenated to form a unique VSI-ID. In this case, PE-1 has a VSI-ID of 65536:3:192.0.2.1. This uniquely identifies the VPLS instance on each individual PE and is advertised as an L2-VPN BGP update.

A BGP-AD update is transmitted to all other PEs via the RR, as follows:

```
[]
A:admin@PE-1# show router bgp routes l2-vpn rd 65536:3 hunt
===============================================================================
 BGP Router ID:192.0.2.1        AS:65536       Local AS:65536
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP L2VPN Routes
===============================================================================
---snip---

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
Route Type     : AutoDiscovery
Route Dist.    : 65536:3
Prefix         : 192.0.2.1
Nexthop        : 192.0.2.1
To             : 192.0.2.6
Res. Nexthop   : n/a
Local Pref.    : 100                    Interface Name : NotAvailable
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : 0
AIGP Metric    : None                   IGP Cost       : n/a
Connector      : None
Community      : target:65536:3 l2-vpn/vrf-imp:65536:3
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.6
Origin         : IGP
AS-Path        : No As-Path
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
Routes : 4
===============================================================================
```

The preceding BGP update is transmitted by PE-1 and has route type auto-discovery.

In this L2-VPN update, the VPLS-ID is encoded as the L2-VPN extended community 65536:3.

The VSI is seen as the prefix 192.0.2.1. The combination of the VPLS-ID and the VSI forms the VSI-ID and uniquely identifies the VPLS instance within this PE router.

The next-hop is also encoded as the local system IP address 192.0.2.1, which allows remote PEs to identify a suitable transport tunnel to PE-1 and for the targeted-LDP peer for instantiating the SDP.

As can be seen within the update, the VPLS-ID 65536:3 is also used to determine the RT extended community and the route distinguisher (RD).

PE-2 configuration

On PE-2, VPLS 3 is created using pseudowire template 1, with VPLS-ID 65536:3 and VSI-ID prefix 192.0.2.2 (system IP address), as follows”

```
# on PE-2:
configure {
    service {
        vpls "VPLS-3" {
            admin-state enable
            service-id 3
            customer "1"
            bgp 1 {
                route-distinguisher "65536:3"
                route-target {
                    export "target:65536:3"
                    import "target:65536:3"
                }
                pw-template-binding "PW2" {
                    split-horizon-group "vpls-shg"
                    import-rt ["target:65536:3"]
                }
            }
            bgp-ad {
                admin-state enable
                vpls-id "65536:3"
                vsi-id-prefix 192.0.2.2
            }
            sap 1/1/4:3.0 {
            }
```

PE-3 configuration

On PE-3, VPLS 3 is created using pseudowire template 2, with VPLS-ID 65536:3—identical to the VPLS-ID of PE-1 and PE-2—and VSI-ID 192.0.2.3 (system IP address), as follows:

```
# on PE-3:
configure {
    service {
         vpls "VPLS-3" {
            admin-state enable
            service-id 3
            customer "1"
            bgp 1 {
                route-distinguisher "65536:3"
                route-target {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    export "target:65536:3"
                    import "target:65536:3"
                }
                pw-template-binding "PW2" {
                    split-horizon-group "vpls-shg"
                    import-rt ["target:65536:3"]
                }
            }
            bgp-ad {
                admin-state enable
                vpls-id "65536:3"
                vsi-id-prefix 192.0.2.3
            }
            sap 1/1/4:3.0 {
            }
```

PE-1 service operation verification

The following output on PE-1 shows that the VPLS and its objects (SAP and auto-discovered spoke SDPs) are operationally up on PE-1:

```
[]
A:admin@PE-1# show service id 3 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 3                   Vpn Id            : 0
Service Type      : VPLS
---snip---

Admin State       : Up                  Oper State        : Up
MTU               : 1514
SAP Count         : 1                   SDP Bind Count    : 2
---snip---

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                             Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/4:3.0                          qinq         1522    1522    Up   Up
sdp:32766:4294967294 SB(192.0.2.3)     BgpAd        0       1556    Up   Up
sdp:32767:4294967295 SB(192.0.2.2)     BgpAd        0       1556    Up   Up
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

The **SB** flag indicates that the SDP is of type spoke-SDP (S flag) BGP (B flag).

BGP is used to discover the VPLS endpoints and exchange network reachability information. LDP is used to signal the pseudowires between the PEs.

LDP signaling occurs when each PE has discovered the endpoints of the VPLS instance. This compares with the use of the provisioned virtual circuit IDs used in an NMS provisioned VPLS instances as per RFC 4762.

The ability of PE-1 to reach the other PE routers with VSIs within the VPLS instance is verified from the following L2-route table:

```
[]
A:admin@PE-1# show service l2-route-table bgp-ad
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Services: L2 Route Information - Summary
===============================================================================
Svc Id     L2-Routes (RD-Prefix)                Next Hop        Origin
               Sdp Bind Id                          PW Temp Id
-------------------------------------------------------------------------------
3          *65536:3-192.0.2.2                   192.0.2.2       BGP-L2
               32767:4294967295                     2
3          *65536:3-192.0.2.3                   192.0.2.3       BGP-L2
               32766:4294967294                     2
-------------------------------------------------------------------------------
No. of L2 Route Entries: 2
===============================================================================
```

This output shows the presence of the signaled pseudowire SDPs. SDPs from PE-1 to PE-2 and PE-3 are signaled using LDP Forwarding Equivalence Class (FEC) Element 129.

Each PE router uses targeted LDP to signal the local and remote endpoints. If there is an endpoint match, then SDPs are instantiated. This compares with the use of LDP for NMS provisioned SDPs, which uses virtual circuit IDs to signal pseudowires using LDP FEC Element 128.

In order to signal the SDPs, the following parameters are required:

1. Attachment Group Identifier (AGI): this is used to carry the VPLS-ID of the local PE router VPLS instance. The VPLS-ID must be identical for all PEs in the same VPLS instance.
2. Source Attachment Individual Identifier (SAII) and Target Attachment Individual Identifier (TAII): these use AII type 1 (RFC 4446) and are used to carry the NLRI (VSI-ID minus the RD) of the remote PE router VPLS instance.

The AGI for each PE must be identical. SAII and TAII must be different.

The following shows the service LDP bindings for VPLS 3 on PE-1:

```
[]
A:admin@PE-1# show router ldp bindings services service-id 3

===============================================================================
LDP Bindings (IPv4 LSR ID 192.0.2.1)
             (IPv6 LSR ID ::)
===============================================================================
Label Status:
        U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        S - Status Signaled Up,  D - Status Signaled Down, e - Label ELC
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
Service Type:
        E - Epipe Service, V - VPLS Service, M - Mirror Service
        A - Apipe Service, F - Fpipe Service, I - IES Service, R - VPRN service
        P - Ipipe Service, C - Cpipe Service
FEC Flags:
        LF - Lower FEC, UF - Upper FEC, M - Community Mismatch,
        BA - ASBR Backup FEC
===============================================================================
LDP Service FEC 128 Bindings
===============================================================================
Type                                        VCId       SDPId             LMTU
Peer                                        SvcId      IngLbl            RMTU
                                                       EgrLbl
-------------------------------------------------------------------------------
No Matching Entries Found
===============================================================================

===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
LDP Service FEC 129 Bindings
===============================================================================
SAII                                            AGII        IngLbl         LMTU
TAII                                            Type        EgrLbl         RMTU
Peer                                            SvcId       SDPId
-------------------------------------------------------------------------------
192.0.2.1                                       1,8:020A00* 524280U        1500
192.0.2.2                                       V-Eth       524277S        1500
192.0.2.2:0                                     3           32767

192.0.2.1                                       1,8:020A00* 524279U        1500
192.0.2.3                                       V-Eth       524280S        1500
192.0.2.3:0                                     3           32766

-------------------------------------------------------------------------------
No. of FEC 129s: 2
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

This shows the two T-LDP bindings for PE-1 toward PE-2 and PE-3 for VPLS 3. The label bindings from this LDP output is identical to the SDP bindings output that follows. The following command can be used to list the SDP IDs and the SDP label bindings:

```
[]
A:admin@PE-1# show service id 3 sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
32766:4294967294 BgpAd    192.0.2.3       Up      Up        524279    524280
32767:4294967295 BgpAd    192.0.2.2       Up      Up        524280    524277
-------------------------------------------------------------------------------
Number of SDPs : 2
-------------------------------------------------------------------------------
===============================================================================
```

The SDP ID for the auto-provisioned SDP toward PE-2 is 32767, the SDP ID toward PE-3 is 32766.

The actual AGI, SAII, and TAII values are seen in the following detailed SDP output.

- AGI — 65536:3
- SAII — Local system IP address 192.0.2.1
- TAII — Remote system IP address 192.0.2.2 or 192.0.2.3

```
[]
A:admin@PE-1# show service id 3 sdp 32767:4294967295 detail

===============================================================================
Service Destination Point (Sdp Id : 32767:4294967295) Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 32767:4294967295  -(192.0.2.2)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id            : 32767:4294967295         Type              : BgpAd
PW-Template Id    : 2
AGI               : 65536:3                  SDP Bind Source   : bgp-l2vpn
Local AII         : 192.0.2.1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Remote AII         : 192.0.2.2
Split Horiz Grp    : vpls-shg
Etree Root Leaf Tag: Disabled                Etree Leaf AC     : Disabled
VC Type            : Ether                   VC Tag            : n/a
Admin Path MTU     : 0                       Oper Path MTU     : 1556
Delivery           : MPLS
Far End            : 192.0.2.2               Tunnel Far End    : n/a
Oper Tunnel Far End: 192.0.2.2
LSP Types          : LDP/BGP
---snip---
```

PE-2 service operation verification

For completeness, the following shows that the VPLS service is operationally up on PE-2.

```
[]
A:admin@PE-2# show service id 3 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 3                     Vpn Id            : 0
Service Type      : VPLS
---snip---

Admin State       : Up                    Oper State        : Up
MTU               : 1514
SAP Count         : 1                     SDP Bind Count    : 2
---snip---

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                               Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/4:3.0                            qinq         1522    1522    Up   Up
sdp:32766:4294967293 SB(192.0.2.3)       BgpAd        0       1556    Up   Up
sdp:32767:4294967294 SB(192.0.2.1)       BgpAd        0       1556    Up   Up
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

```
[]
A:admin@PE-2# show service l2-route-table bgp-ad

===============================================================================
Services: L2 Route Information - Summary
===============================================================================
Svc Id     L2-Routes (RD-Prefix)                Next Hop        Origin
               Sdp Bind Id                          PW Temp Id
-------------------------------------------------------------------------------
3          *65536:3-192.0.2.1                   192.0.2.1       BGP-L2
               32767:4294967294                     2
3          *65536:3-192.0.2.3                   192.0.2.3       BGP-L2
               32766:4294967293                     2
-------------------------------------------------------------------------------
No. of L2 Route Entries: 2
===============================================================================
```

```
[]
A:admin@PE-2# show router ldp bindings services service-id 3

===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
LDP Bindings (IPv4 LSR ID 192.0.2.2)
             (IPv6 LSR ID ::)
===============================================================================
Label Status:
        U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        S - Status Signaled Up,  D - Status Signaled Down, e - Label ELC
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
Service Type:
        E - Epipe Service, V - VPLS Service, M - Mirror Service
        A - Apipe Service, F - Fpipe Service, I - IES Service, R - VPRN service
        P - Ipipe Service, C - Cpipe Service
FEC Flags:
        LF - Lower FEC, UF - Upper FEC, M - Community Mismatch,
        BA - ASBR Backup FEC
===============================================================================
LDP Service FEC 128 Bindings
===============================================================================
Type                                          VCId       SDPId             LMTU
Peer                                          SvcId      IngLbl            RMTU
                                                         EgrLbl
-------------------------------------------------------------------------------
No Matching Entries Found
===============================================================================

===============================================================================
LDP Service FEC 129 Bindings
===============================================================================
SAII                                            AGII        IngLbl         LMTU
TAII                                            Type        EgrLbl         RMTU
Peer                                            SvcId       SDPId
-------------------------------------------------------------------------------
192.0.2.2                                       1,8:020A00* 524277U        1500
192.0.2.1                                       V-Eth       524280S        1500
192.0.2.1:0                                     3           32767

192.0.2.2                                       1,8:020A00* 524278U        1500
192.0.2.3                                       V-Eth       524279S        1500
192.0.2.3:0                                     3           32766

-------------------------------------------------------------------------------
No. of FEC 129s: 2
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

```
[]
A:admin@PE-2# show service id 3 sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
32766:4294967293 BgpAd    192.0.2.3       Up      Up        524278    524279
32767:4294967294 BgpAd    192.0.2.1       Up      Up        524277    524280
-------------------------------------------------------------------------------
Number of SDPs : 2
-------------------------------------------------------------------------------
===============================================================================
```

PE-3 service operation verification

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

For completeness, the same commands are launched on PE-3, as follows:

```
[]
A:admin@PE-3# show service id 3 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 3                   Vpn Id            : 0
Service Type      : VPLS
---snip---

Admin State       : Up                  Oper State        : Up
MTU               : 1514
SAP Count         : 1                   SDP Bind Count    : 2
---snip---

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                             Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/4:3.0                          qinq         1522    1522    Up   Up
sdp:32766:4294967294 SB(192.0.2.2)     BgpAd        0       1556    Up   Up
sdp:32767:4294967295 SB(192.0.2.1)     BgpAd        0       1556    Up   Up
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

```
[]
A:admin@PE-3# show service l2-route-table bgp-ad

===============================================================================
Services: L2 Route Information - Summary
===============================================================================
Svc Id     L2-Routes (RD-Prefix)                Next Hop        Origin
           Sdp Bind Id                              PW Temp Id
-------------------------------------------------------------------------------
3          *65536:3-192.0.2.1                   192.0.2.1       BGP-L2
           32767:4294967295                         2
3          *65536:3-192.0.2.2                   192.0.2.2       BGP-L2
           32766:4294967294                         2
-------------------------------------------------------------------------------
No. of L2 Route Entries: 2
===============================================================================
```

```
[]
A:admin@PE-3# show router ldp bindings services service-id 3

===============================================================================
LDP Bindings (IPv4 LSR ID 192.0.2.3)
            (IPv6 LSR ID ::)
===============================================================================
Label Status:
        U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        S - Status Signaled Up,  D - Status Signaled Down, e - Label ELC
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
Service Type:
        E - Epipe Service, V - VPLS Service, M - Mirror Service
        A - Apipe Service, F - Fpipe Service, I - IES Service, R - VPRN service
        P - Ipipe Service, C - Cpipe Service
FEC Flags:
        LF - Lower FEC, UF - Upper FEC, M - Community Mismatch,
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        BA - ASBR Backup FEC
===============================================================================
LDP Service FEC 128 Bindings
===============================================================================
Type                                        VCId       SDPId             LMTU
Peer                                        SvcId      IngLbl            RMTU
                                                       EgrLbl
-------------------------------------------------------------------------------
No Matching Entries Found
===============================================================================

===============================================================================
LDP Service FEC 129 Bindings
===============================================================================
SAII                                          AGII        IngLbl         LMTU
TAII                                          Type        EgrLbl         RMTU
Peer                                          SvcId       SDPId
-------------------------------------------------------------------------------
192.0.2.3                                     1,8:020A00* 524280U        1500
192.0.2.1                                     V-Eth       524279S        1500
192.0.2.1:0                                   3           32767

192.0.2.3                                     1,8:020A00* 524279U        1500
192.0.2.2                                     V-Eth       524278S        1500
192.0.2.2:0                                   3           32766

-------------------------------------------------------------------------------
No. of FEC 129s: 2
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

```
[]
A:admin@PE-3# show service id 3 sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
32766:4294967294 BgpAd    192.0.2.2       Up      Up        524279    524278
32767:4294967295 BgpAd    192.0.2.1       Up      Up        524280    524279
-------------------------------------------------------------------------------
Number of SDPs : 2
-------------------------------------------------------------------------------
===============================================================================
```

BGP AD using pre-provisioned SDPs

It is possible to configure BGP-AD instances that use RSVP transport tunnels. In this case, the LSPs and SDPs must be manually created.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 177: VPLS instance using pre-provisioned SDPs

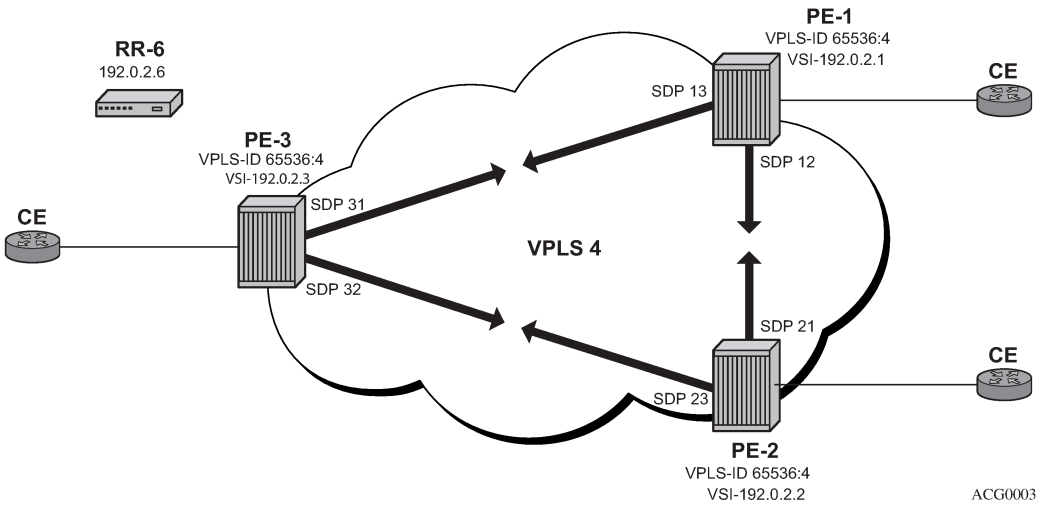

Figure 177: VPLS instance using pre-provisioned SDPs shows a VPLS instance configured across three PE routers as before.

The following SDPs are configured on the three PEs:

```
# on PE-1:
configure {
    service {
        sdp 12 {
            admin-state enable
            description "RSVP-based SDP from PE-1 to PE-2"
            delivery-type mpls
            far-end {
                ip-address 192.0.2.2
            }
            lsp "LSP-PE-1-PE-2" { }
        }
        sdp 13 {
            admin-state enable
            description "RSVP-based SDP from PE-1 to PE-3"
            delivery-type mpls
            far-end {
                ip-address 192.0.2.3
            }
            lsp "LSP-PE-1-PE-3" { }
        }
```

```
# on PE-2:
configure {
    service {
        sdp 21 {
            admin-state enable
            description "RSVP-based SDP from PE-2 to PE-1"
            delivery-type mpls
            far-end {
                ip-address 192.0.2.1
            }
            lsp "LSP-PE-2-PE-1" { }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
        sdp 23 {
            admin-state enable
            description "RSVP-based SDP from PE-2 to PE-3"
            delivery-type mpls
            far-end {
                ip-address 192.0.2.3
            }
            lsp "LSP-PE-2-PE-3" { }
        }
```

```
# on PE-3:
configure {
    service {
        sdp 31 {
            admin-state enable
            description "RSVP-based SDP from PE-3 to PE-1"
            delivery-type mpls
            far-end {
                ip-address 192.0.2.1
            }
            lsp "LSP-PE-3-PE-1" { }
        }
        sdp 32 {
            admin-state enable
            description "RSVP-based SDP from PE-3 to PE-2"
            delivery-type mpls
            far-end {
                ip-address 192.0.2.2
            }
            lsp "LSP-PE-3-PE-2" { }
        }
```

The PW template to be used within each VPLS instance must be provisioned on all PEs and must use the keyword provisioned-sdp use, as follows:

```
# on PE-1, PE-2, PE-3:
configure {
    service {
        pw-template "PW3" {
            pw-template-id 3
            provisioned-sdp use
        }
```

The following output shows the configuration required for a VPLS service using a pseudowire template configured for pre-provisioned RSVP SDPs.

```
# on PE-1:
configure {
    service {
        vpls "VPLS-4" {
            admin-state enable
            service-id 4
            customer "1"
            bgp 1 {
                route-distinguisher "65536:4"
                route-target {
                    export "target:65536:4"
                    import "target:65536:4"
                }
                pw-template-binding "PW3" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    split-horizon-group "vpls-shg"
                    import-rt ["target:65536:4"]
                }
            }
            bgp-ad {
                admin-state enable
                vpls-id "65536:4"
                vsi-id-prefix 192.0.2.1
            }
            sap 1/1/4:4.0 {
            }
```

Similarly, on PE-2 the configuration is as follows:

```
# on PE-2:
configure {
    service {
        vpls "VPLS-4" {
            admin-state enable
            service-id 4
            customer "1"
            bgp 1 {
                route-distinguisher "65536:4"
                route-target {
                    export "target:65536:4"
                    import "target:65536:4"
                }
                pw-template-binding "PW3" {
                    split-horizon-group "vpls-shg"
                    import-rt ["target:65536:4"]
                }
            }
            bgp-ad {
                admin-state enable
                vpls-id "65536:4"
                vsi-id-prefix 192.0.2.2
            }
            sap 1/1/4:4.0 {
            }
```

On PE-3, VPLS 4 is configured as follows:

```
# on PE-3:
configure {
    service {
        vpls "VPLS-4" {
            admin-state enable
            service-id 4
            customer "1"
            bgp 1 {
                route-distinguisher "65536:4"
                route-target {
                    export "target:65536:4"
                    import "target:65536:4"
                }
                pw-template-binding "PW3" {
                    split-horizon-group "vpls-shg"
                    import-rt ["target:65536:4"]
                }
            }
            bgp-ad {
                admin-state enable
                vpls-id "65536:4"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                vsi-id-prefix 192.0.2.3
            }
            sap 1/1/4:4.0 {
            }
```

The following output shows that the service and its objects are operationally up on PE-1.

```
[]
A:admin@PE-1# show service id 4 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 4                   Vpn Id            : 0
Service Type      : VPLS
---snip---

Admin State       : Up                  Oper State        : Up
MTU               : 1514
SAP Count         : 1                   SDP Bind Count    : 2
---snip---

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                             Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/4:4.0                          qinq         1522    1522    Up   Up
sdp:12:4294967293 S(192.0.2.2)         BgpAd        0       1556    Up   Up
sdp:13:4294967292 S(192.0.2.3)         BgpAd        0       1556    Up   Up
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

The SDP identifiers are the pre-provisioned SDPs: SDP 12 and 13.

The following command shows that the service and its objects are operationally up on PE-2.

```
[]
A:admin@PE-2# show service id 4 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 4                   Vpn Id            : 0
Service Type      : VPLS
---snip---

Admin State       : Up                  Oper State        : Up
MTU               : 1514
SAP Count         : 1                   SDP Bind Count    : 2
---snip---

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                             Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/4:4.0                          qinq         1522    1522    Up   Up
sdp:21:4294967292 S(192.0.2.1)         BgpAd        0       1556    Up   Up
sdp:23:4294967291 S(192.0.2.3)         BgpAd        0       1556    Up   Up
===============================================================================
* indicates that the corresponding row element may have been truncated.
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The following command shows that the service and its objects are operationally up on PE-3.

```
[]
A:admin@PE-3# show service id 4 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 4                   Vpn Id            : 0
Service Type      : VPLS
---snip---

Admin State       : Up                  Oper State        : Up
MTU               : 1514
SAP Count         : 1                   SDP Bind Count    : 2
---snip---

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                             Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/4:4.0                          qinq         1522    1522    Up   Up
sdp:31:4294967291 S(192.0.2.1)         BgpAd        0       1556    Up   Up
sdp:32:4294967292 S(192.0.2.2)         BgpAd        0       1556    Up   Up
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

Conclusion

BGP-AD coupled with LDP pseudowire signaling allows the delivery of L2-VPN services to customers where BGP is commonly used. This example shows the configuration of BGP-AD together with the associated show outputs which can be used for verification and troubleshooting.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

LDP VPLS Using BGP Auto-Discovery — Prefer Provisioned SDP

This chapter provides information about LDP VPLS using BGP auto-discovery — prefer provisioned SDP.

Topics in this chapter include:

Applicability

This chapter was initially written for SR OS Release 14.0.R6, but the MD-CLI in the current edition is based on SR OS Release 21.2.R1. BGP Auto-Discovery (BGP-AD) based on RFC 6074 is supported in SR OS Release 6.0, and later. The **provisioned-sdp prefer** option is supported in SR OS Release 14.0.R1, and later.

Overview

As described in chapter LDP VPLS Using BGP Auto-Discovery, BGP-AD based on RFC 6074 can auto-create SDP bindings, but an operator can force the system to use a provisioned SDP by specifying the **provisioned-sdp use** option. This chapter compares the **provisioned-sdp use** option with the **provisioned-sdp prefer** option. The chapter describes a migration scenario for a VPLS service with a pseudowire (PW) template binding, restricted to using provisioned SDPs toward a PW template binding preferring to use provisioned SDPs, but auto-creating SDPs in case there is no suitable manually created SDP available.

PW templates

PW templates can be configured with the following command:

```
[ex:/configure service]
A:admin@PE-1# pw-template "PW 1" ?

 pw-template

 Immutable fields      - pw-template-id, provisioned-sdp, auto-gre-sdp
---snip---
```

When provisioned SDPs are to be used, the **provisioned-sdp** context must be configured:

```
*[ex:/configure service pw-template "PW 1"]
A:admin@PE-1# provisioned-sdp ?

 provisioned-sdp <keyword>
 <keyword>  - (use|prefer)

 'provisioned-sdp' is: immutable
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Provisioned SDP type

Warning: Modifying this element recreates 'configure service pw-template "PW 1"'
automatically for the new value to take effect.
```

- When the **provisioned-sdp use** option is configured, the tunnel manager is forced to look for a provisioned and active SDP to the far-end PE. The far-end PE is auto-discovered from the BGP next hop. If multiple SDPs are active to this far-end PE, the tunnel manager chooses the SDP template with the best metric. If there is a tie, the SDP ID is used as a tie-breaker and the highest SDP ID wins. However, if no provisioned SDP exists, the SDP binding will not be instantiated.
- When the **provisioned-sdp prefer** option is configured, the behavior is the same when a provisioned SDP exists. When the tunnel manager finds an existing matching SDP, it will use it even if it is operationally down. Only when no provisioned SDP exists, will the SDP binding be auto-created.
- When a PW template is configured without the **provisioned-sdp use** or **provisioned-sdp prefer** option, the SDP bindings will be auto-created.

Figure 178: LDP VPLS using BGP-AD with provisioned-sdp use option shows the following use case: the metro Ethernet networks were initially built with provisioned SDPs. Intra-metro services are provisioned using provisioned SDPs; for example, customer X has a VPLS service defined in the metro Ethernet networks, using BGP-AD with a PW template to use the provisioned SDPs in the metro Ethernet networks.

Figure 178: LDP VPLS using BGP-AD with provisioned-sdp use option

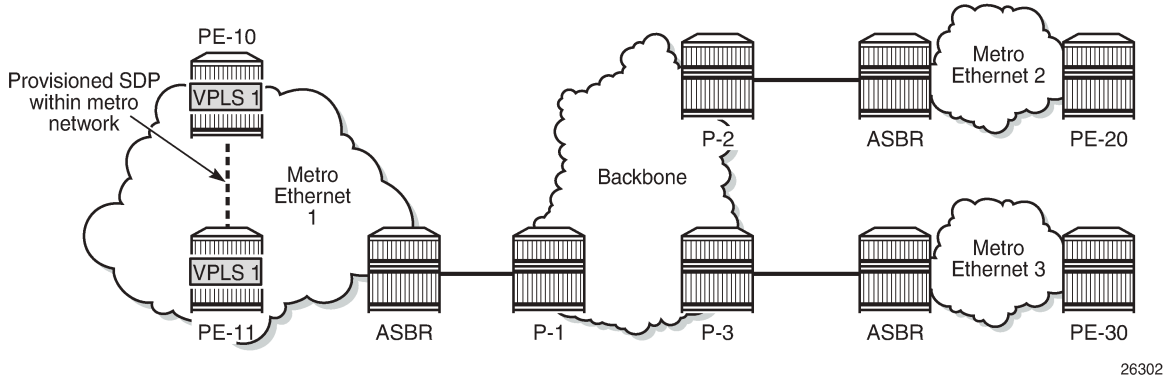

The service provider initially started with PE-10 and PE-11 in metro Ethernet 1, but now wants to add PE-20 and PE-30 as new sites to the VPLS service. Therefore, the BGP-AD routes should propagate beyond the boundaries of the metro Ethernet network. The backbone network may be in a different AS, but in this example, all networks are in the same AS. VPLS 1 of customer X can have sites added to the service on PEs in different metro Ethernet networks. A new PW template is configured with the **provisioned-sdp prefer** option and applied to the VPLS service.

- When a new site within the metro Ethernet network is added, an SDP is already provisioned to this site and this SDP is used for the SDP binding in the VPLS.
- When a new site in a different metro Ethernet network is added, no SDP is available to the site in the remote metro Ethernet network and the SDP binding is auto-created.

Figure 179: LDP VPLS using BGP-AD with provisioned-sdp prefer option shows the SDP bindings in VPLS 1 between PE-10 and the other PEs. For simplicity, the SDP bindings between the other PEs are not shown.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 179: LDP VPLS using BGP-AD with provisioned-sdp prefer option

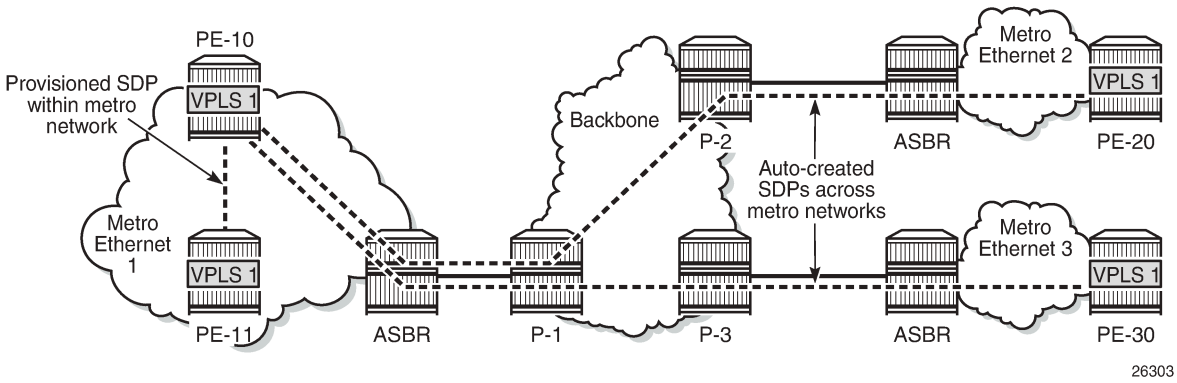

When PW templates are in use, it is not possible to modify the **provisioned-sdp prefer** option to **provisioned-sdp use** and vice versa. To support migration from one PW template to another with minimal service impact, two PW templates can be applied in parallel, as shown in the Configuration section.

Configuration

Figure 180: Example topology shows the example topology. For simplicity, all nodes are in the same AS.

Figure 180: Example topology

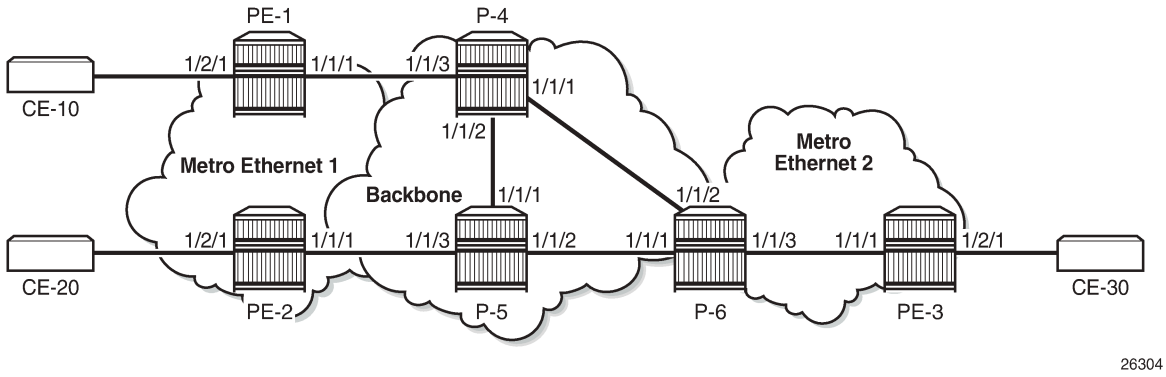

The initial configuration includes the following:

- Cards, MDAs, ports
- Router interfaces
- IS-IS as IGP (or OSPF) on all interfaces
- MPLS and RSVP on all interfaces, except "int-P-4-P-6" and "int-P-5-P-6".
- LDP on all interfaces

BGP is configured on all PE routers for address family l2-vpn, as follows:

```
# on PE-1, PE-2, PE-3:
configure {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    router "Base" {
        autonomous-system 64496
        bgp {
            group "internal" {
                peer-as 64496
                family {
                    l2-vpn true
                }
            }
            neighbor "192.0.2.6" {
                group "internal"
            }
        }
```

The BGP configuration on the route reflector (RR) P-6 is as follows:

```
# on P-6:
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            group "rr-internal" {
                peer-as 64496
                family {
                    l2-vpn true
                }
                cluster {
                    cluster-id 1.1.1.1
                }
            }
            neighbor "192.0.2.1" {
                group "rr-internal"
            }
            neighbor "192.0.2.2" {
                group "rr-internal"
            }
            neighbor "192.0.2.3" {
                group "rr-internal"
            }
        }
```

On PE-1 and PE-2 in metro Ethernet network 1, an RSVP LSP is created that is used in a manually created SDP. The LSP configuration on PE-1 is as follows:

```
# on PE-1:
configure {
    router "Base" {
        mpls {
            admin-state enable
            path "loose" {
                admin-state enable
            }
            lsp "LSP-PE-1-PE-2" {
                admin-state enable
                type p2p-rsvp
                to 192.0.2.2
                primary "loose" {
                }
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

On PE-1, SDP 12 is configured as follows:

```
# on PE-1:
configure {
    service {
        sdp 12 {
            admin-state enable
            description "SDP12 to 192.0.2.2"
            delivery-type mpls
            far-end {
                ip-address 192.0.2.2
            }
            lsp "LSP-PE-1-PE-2" { }
        }
```

The configuration on PE-2 is similar.

LDP VPLS using AD without provisioned-sdp prefer option

Initially, the following two PW templates are configured on all PEs: PW template 1 has the **provisioned-sdp use** option and PW template 2 is configured without any option; therefore, SDP bindings will be auto-created.

```
# on PE-1, PE-2, PE-3:
configure {
    service {
        pw-template "PW 1" {
            pw-template-id 1
            provisioned-sdp use
        }
        pw-template "PW 2" {
            pw-template-id 2
        }
```

The following lists the PW templates configured on PE-1:

```
[/]
A:admin@PE-1# show service pw-template

===============================================================================
PW Template information
===============================================================================
PW Template Id      SDP                           Last Update
-------------------------------------------------------------------------------
1                   Use-provisioned               04/01/2021 16:35:27
2                   Auto-mpls                     04/01/2021 15:33:56
===============================================================================
```

On all PEs, two VPLS services are configured: VPLS 1 with BGP-AD PW template 1 and VPLS 2 with PW template 2, as follows:

```
# on PE-1, PE-2, PE-3:
configure {
    service {
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            bgp 1 {
                route-distinguisher "64496:1"
                route-target {
                    export "target:64496:1"
                    import "target:64496:1"
                }
                pw-template-binding "PW 1" {
                }
            }
            bgp-ad {
                admin-state enable
                vpls-id "64496:1"
            }
            sap 1/2/1:1 {
            }
        }
        vpls "VPLS 2" {
            admin-state enable
            service-id 2
            customer "1"
            bgp 1 {
                route-distinguisher "64496:2"
                route-target {
                    export "target:64496:2"
                    import "target:64496:2"
                }
                pw-template-binding "PW 2" {
                    import-rt ["target:64496:2"]
                }
            }
            bgp-ad {
                admin-state enable
                vpls-id "64496:2"
            }
            sap 1/2/1:2 {
            }
        }
```

On PE-1, the following SDP bindings have been created:

```
[/]
A:admin@PE-1# show service sdp-using

===============================================================================
SDP Using
===============================================================================
SvcId      SdpId              Type   Far End             Opr   I.Label E.Label
                                                        State
-------------------------------------------------------------------------------
1          12:4294967295      BgpAd  192.0.2.2           Up    524280  524280
2          32766:4294967293   BgpAd  192.0.2.3           Up    524278  524280
2          32767:4294967294   BgpAd  192.0.2.2           Up    524279  524279
-------------------------------------------------------------------------------
Number of SDPs : 3
-------------------------------------------------------------------------------
===============================================================================
```

The first SDP binding is created by BGP-AD in VPLS 1 and uses the configured SDP 12 with far-end PE-2; the other two SDP bindings have been auto-created by BGP-AD in VPLS 2 and have far-end PE-2 and PE-3.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The list of SDP bindings on PE-2 looks similar:

```
[/]
A:admin@PE-2# show service sdp-using

===============================================================================
SDP Using
===============================================================================
SvcId      SdpId              Type   Far End             Opr   I.Label E.Label
                                                        State
-------------------------------------------------------------------------------
1          21:4294967295      BgpAd  192.0.2.1           Up    524280  524280
2          32766:4294967293   BgpAd  192.0.2.3           Up    524278  524281
2          32767:4294967294   BgpAd  192.0.2.1           Up    524279  524279
-------------------------------------------------------------------------------
Number of SDPs : 3
-------------------------------------------------------------------------------
===============================================================================
```

On PE-3, there are only two SDP bindings, both in VPLS 2:

```
[/]
A:admin@PE-3# show service sdp-using

===============================================================================
SDP Using
===============================================================================
SvcId      SdpId              Type   Far End             Opr   I.Label E.Label
                                                        State
-------------------------------------------------------------------------------
2          32766:4294967294   BgpAd  192.0.2.2           Up    524281  524278
2          32767:4294967295   BgpAd  192.0.2.1           Up    524280  524278
-------------------------------------------------------------------------------
Number of SDPs : 2
-------------------------------------------------------------------------------
===============================================================================
```

Log "99" on PE-3 shows that the system failed to create a dynamic BGP-L2VPN SDP binding because no provisioned SDP was found, as follows:

```
69 2021/04/01 16:36:44.405 CEST MAJOR: SVCMGR #2322 Base
"The system failed to create a dynamic bgp-l2vpn SDP Bind  in service 1 with SDP pw-template
 policy 1 for the following reason: suitable manual SDP not found."
```

Figure 181: SDP bindings in VPLS 1 with provisioned-sdp use option shows the SDPs used in VPLS 1. PE-1 and PE-2 both used the provisioned SDP. PE-3 has no SDP bindings in VPLS 1.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 181: SDP bindings in VPLS 1 with provisioned-sdp use option

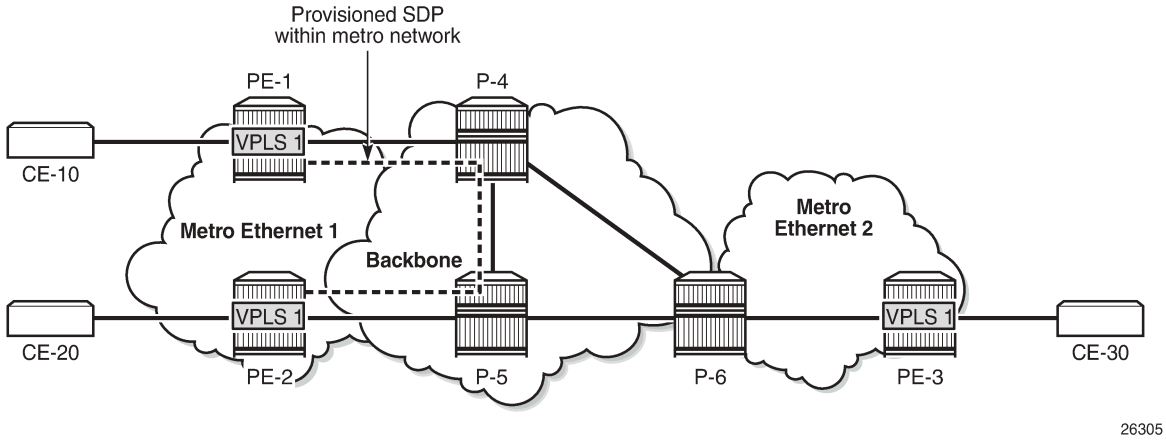

Figure 182: Auto-created SDP bindings in VPLS 2 shows the auto-created SDP bindings in VPLS 2. Each PE has two auto-created SDP bindings to each other PE.

Figure 182: Auto-created SDP bindings in VPLS 2

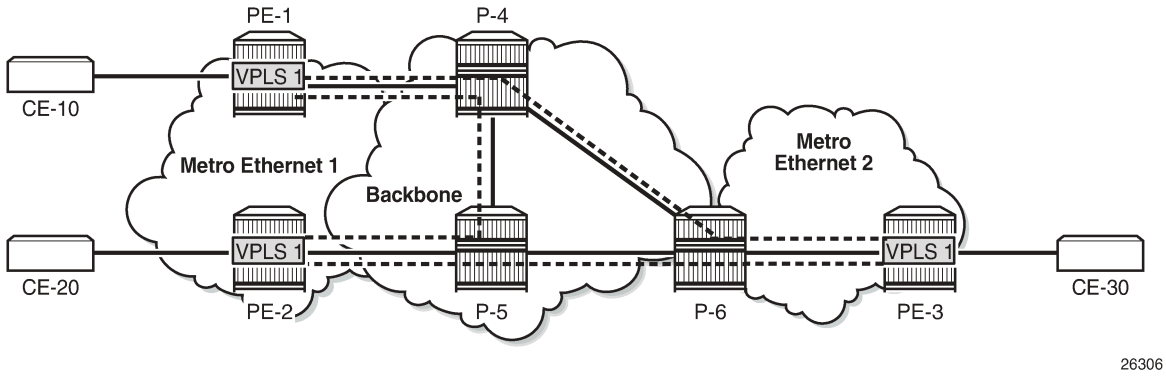

Migrate VPLS 1 to provisioned-sdp prefer option

VPLS 1 uses PW template 1 with the **provisioned-sdp use** option. This option cannot be modified when the PW template is in use, as follows:

```
[ex:/configure service pw-template "PW 1"]
A:admin@PE-1# provisioned-sdp prefer

*[ex:/configure service pw-template "PW 1"]
A:admin@PE-1# commit
MINOR: SVCMGR #5609: configure service vpls "VPLS 1" bgp 1 pw-template-binding "PW 1" - PW
 Template is in use
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The following steps are needed to migrate to another PW template with the **provisioned-sdp prefer** option without service outage:

1. Configure new PW template with **provisioned-sdp prefer** option.
2. Add new PW template binding to VPLS and verify which PW template is used.
3. Modify old PW template binding to make it not usable.
4. Launch tools command to re-evaluate old PW template in the VPLS.
5. When the old PW template is not used anymore, remove PW template binding from the VPLS configuration.

A new PW template with the provisioned-sdp prefer option is configured on all PEs, as follows:

```
# on PE-1, PE-2, PE-3:
configure {
    service {
        pw-template "PW 10" {
            pw-template-id 10
            provisioned-sdp prefer
        }
```

An additional PW template binding is configured in VPLS 1 on all PEs, as follows:

```
# on PE-1, PE-2, PE-3:
configure {
    service {
        vpls "VPLS 1" {
            bgp 1 {
                pw-template-binding "PW 10" {
                }
            }
        }
```

The configuration of VPLS 1 includes two PW template bindings, as follows:

```
[ex:/configure service vpls "VPLS 1"]
A:admin@PE-1# info
    admin-state enable
    service-id 1
    customer "1"
    bgp 1 {
        route-distinguisher "64496:1"
        route-target {
            export "target:64496:1"
            import "target:64496:1"
        }
        pw-template-binding "PW 1" {
        }
        pw-template-binding "PW 10" {
        }
    }
    bgp-ad {
        admin-state enable
        vpls-id "64496:1"
    }
    sap 1/2/1:1 {
    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The following shows that no additional SDP bindings have been created. The only SDP binding in VPLS 1 on PE-1 uses the provisioned SDP 12.

```
[/]
A:admin@PE-1# show service id 1 sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
12:4294967295    BgpAd    192.0.2.2       Up      Up        524280    524280
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
```

The following shows that PW template 1 was used for the creation of the SDP binding:

```
[/]
A:admin@PE-1# show service id 1 sdp detail | match 'SDP Id|PW-Template Id'
SDP Id             : 12:4294967295            Type              : BgpAd
PW-Template Id     : 1
```

The PW template 10 has a higher ID than PW template 1 and is not used. Re-evaluating the PW template binding for PW template 1 in VPLS 1 will make no difference if both PW templates are usable. However, PW template 1 can be made unusable by adding a dummy **import-rt** not matching any route in the VPLS, as follows:

```
# on PE-1, PE-2, PE-3:
configure {
    service {
        vpls "VPLS 1" {
            bgp 1 {
                pw-template-binding "PW 1" {
                    import-rt "target:111:111"
                }
```

As a result, PW template 10 with the **provisioned-sdp prefer** option is used for the automatic creation of SDP bindings where no provisioned SDP is available, as follows:

```
[/]
A:admin@PE-1# show service id 1 sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
12:4294967295    BgpAd    192.0.2.2       Up      Up        524280    524280
32766:4294967292 BgpAd    192.0.2.3       Up      Up        524277    524279
-------------------------------------------------------------------------------
Number of SDPs : 2
-------------------------------------------------------------------------------
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

For the first SDP binding, PW template 1 is used, and for the second SDP binding, PW template 10 is used, as follows:

```
[/]
A:admin@PE-1# show service id 1 sdp detail | match 'SDP Id|PW-Template Id'
SDP Id             : 12:4294967295           Type              : BgpAd
PW-Template Id     : 1
SDP Id             : 32766:4294967292        Type              : BgpAd
PW-Template Id     : 10
```

The following command forces the system to re-evaluate PW template 1 in VPLS 1:

```
[/]
A:admin@PE-1# tools perform service id 1 eval-pw-template 1 allow-service-impact
eval-pw-template succeeded for Svc 1 12:4294967295 Policy 1
```

As a result, only PW template 10 is used for the creation of SDP bindings in VPLS 1, as follows:

```
[/]
A:admin@PE-1# show service id 1 sdp detail | match 'SDP Id|PW-Template Id'
SDP Id             : 12:4294967291           Type              : BgpAd
PW-Template Id     : 10
SDP Id             : 32766:4294967292        Type              : BgpAd
PW-Template Id     : 10
```

PW template 1 is not used anymore and can be removed from the VPLS configuration, as follows:

```
# on PE-1, PE-2, PE-3:
configure {
    service {
        vpls "VPLS 1" {
            bgp 1 {
                delete pw-template-binding "PW 1" {
                }
```

The configuration of VPLS 1 on PE-1 contains only a PW template binding for PW template 10, as follows:

```
[ex:/configure service vpls "VPLS 1" bgp 1]
A:admin@PE-1# info
    route-distinguisher "64496:1"
    route-target {
        export "target:64496:1"
        import "target:64496:1"
    }
    pw-template-binding "PW 10" {
    }
```

Figure 183: SDP bindings in VPLS 1 with provisioned-sdp prefer option shows the SDP bindings in VPLS 1 with the **provisioned-sdp prefer** option. Within metro Ethernet network 1, the provisioned SDP is used, and between metro Ethernet networks, auto-created SDP bindings are used.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 183: SDP bindings in VPLS 1 with provisioned-sdp prefer option

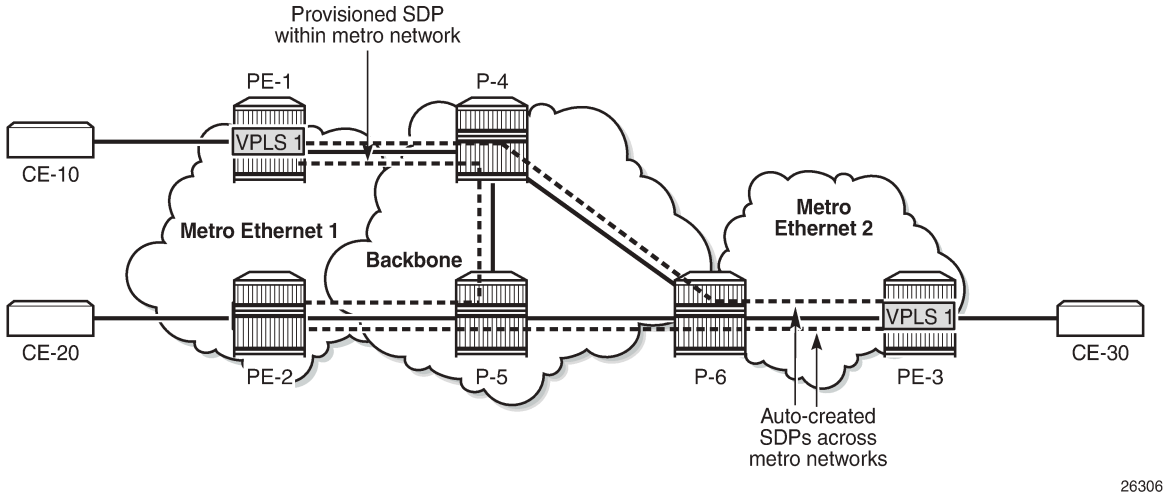

Conclusion

LDP VPLS using BGP-AD allows the creation of SDP bindings that are either auto-created or that use provisioned SDPs. When the **provisioned-sdp prefer** option is used, the tunnel manager will look for a provisioned and active SDP to the far end and use it, if available, even if it is down. When no provisioned SDP is available, the system will auto-create an SDP binding.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Mobility for EVPN Hosts Within an R-VPLS

This chapter provides information about Mobility for EVPN Hosts Within an R-VPLS.

Topics in this chapter include:

Applicability

The information and configuration in this chapter are based on SR OS Release 21.10.R2. Efficient EVPN host mobility without tromboning or hairpinning in an R-VPLS is supported for IPv4 in SR OS Release 19.10.R3 and later and is supported for IPv6 in SR OS Release 20.5.R1 and later.

Overview

SR OS can populate a VPRN route table with host routes learned from the IPv4 Address Resolution Protocol (ARP) messages or IPv6 Neighbor Discovery (ND) protocol messages. The host routes can be advertised in the VPRN context as IP-VPN or EVPN route type 5 (RT5), to be used by an IP-VPN or EVPN core network for inter-subnet forwarding. SR OS supports *draft-ietf-bess-evpn-inter-subnet-forwarding* for a dynamic and efficient routing between remote hosts, avoiding hairpinning.

In SR OS releases earlier than Release 19.10.R3, inefficient hairpinning situations may occur when the VPRN is configured to advertise IPv4 host routes as IP-VPN or EVPN RT5 routes. Figure 184: Hairpinning in a broadcast domain after switchover for SR OS releases earlier than Release 19.10.R3 shows hairpinning in an EVPN broadcast domain with PE-1, PE-2, and PE-3.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 184: Hairpinning in a broadcast domain after switchover for SR OS releases earlier than Release 19.10.R3

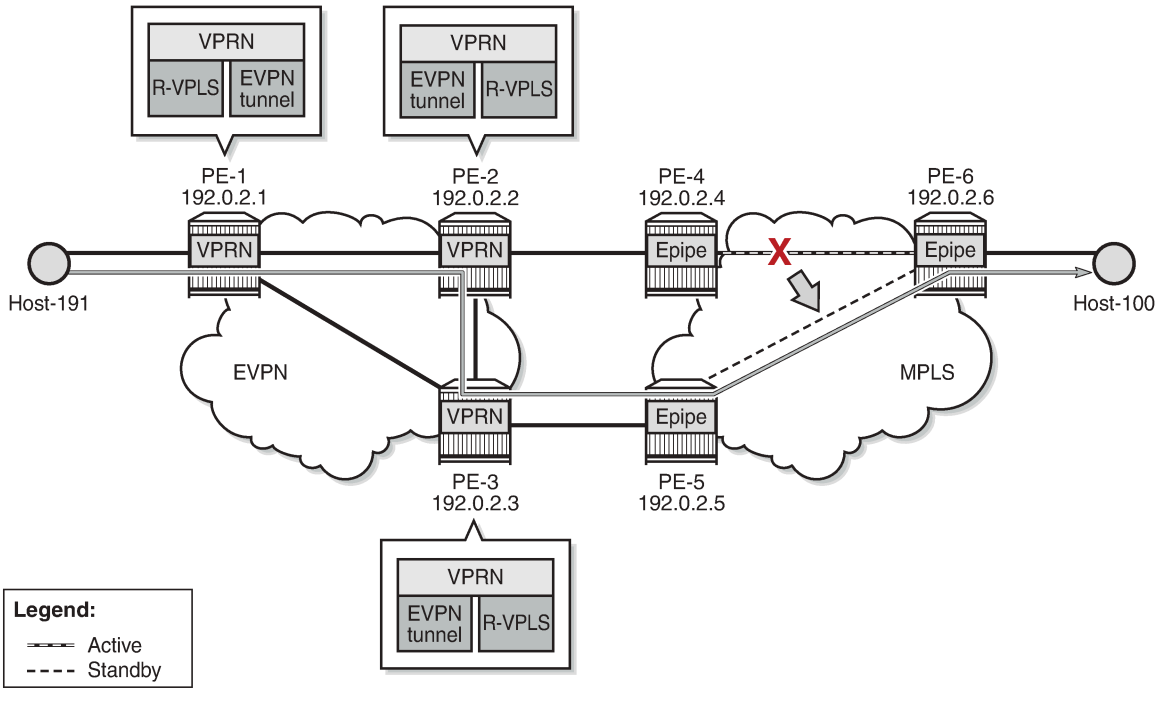

When host-100 comes up, it sends a Gratuitous Address Resolution Protocol (GARP) message that is then learned on PE-2 and PE-3. PE-2 and PE-3 are configured to advertise host routes, so they generate an RT5 host route for prefix 10.0.0.100/32 of host-100. PE-3 selects its best RT5 for 10.0.0.100/32 and traffic from host-191 to host-100 uses the path via PE-1, PE-2, PE-4, and PE-6. However, when the active path between PE-4 and PE-6 fails, the standby path between PE-5 and PE-6 takes over and hairpinning occurs when PE-1 continues selecting PE-2 as the next hop, while a more efficient path is possible via next-hop PE-3. Traffic from host-191 to host-100 uses the path via PE-1, PE-2, PE-3, PE-5, and PE-6.

In SR OS Release 19.10.R3 and later, the more efficient path from host-191 via PE-1, PE-3, PE-5, and PE-6 to host-100 is used, as shown in Figure 185: Forwarding in a broadcast domain after switchover for SR OS Release 19.10.R3 and later.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 185: Forwarding in a broadcast domain after switchover for SR OS Release 19.10.R3 and later

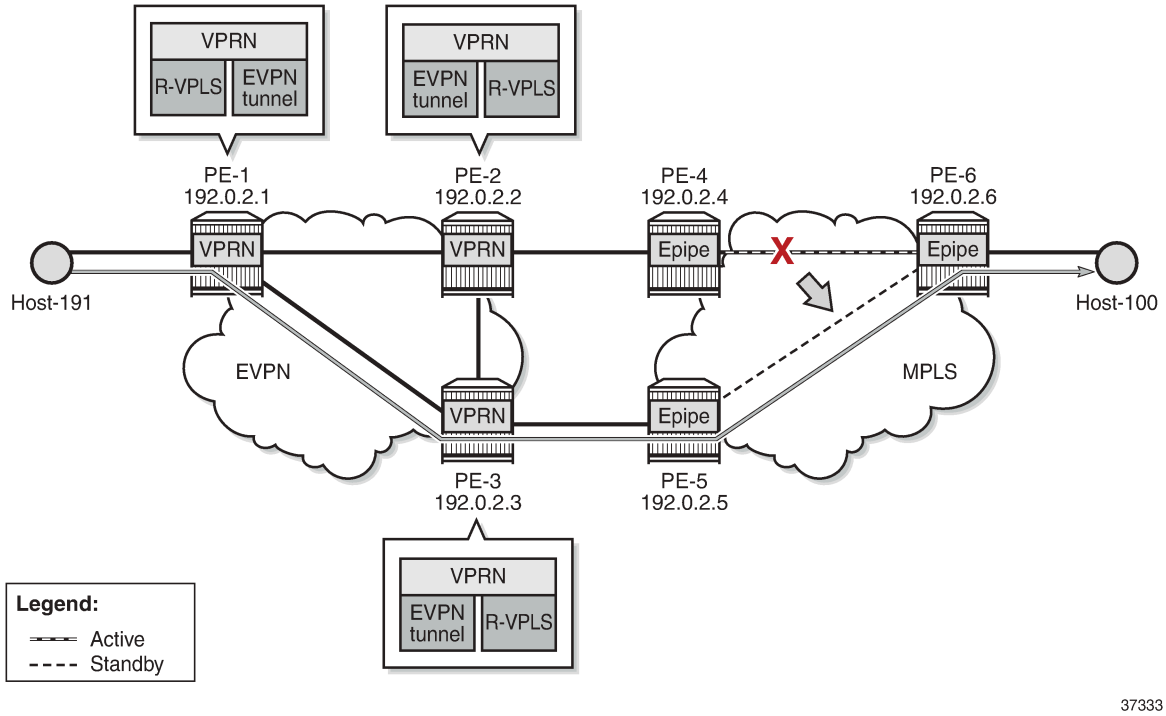

In SR OS Release 19.10.R3 and later, EVPN host mobility is supported for IPv4 as described in section "Symmetric and Asymmetric IRB" of *draft-ietf-bess-evpn-inter-subnet-forwarding*. When a host moves from a source PE to a target PE in the same broadcast domain, the behavior for IPv4 hosts is one of the following.

1. The host initiates an ARP request or GARP.
2. The host sends a data packet without first initiating an ARP request or GARP.
3. The host does not send any traffic and the source PE generates an ARP request when the MAC address of the host expires and the EVPN-MAC is withdrawn.

All three scenarios are described in more detail later, where the move of host-100 from source PE-2 to target PE-3 is simulated.

For the first of these scenarios, the VPRN configuration on PE-2 is as follows:

```
# on PE-2:
configure {
    service {
        vprn "ip-vrf-16" {
            admin-state enable
            service-id 16
            customer "1"
            interface "evi-15" {
                mac 00:00:00:00:00:02
                vpls "sbd-15" {
                    evpn-tunnel {
                    }
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            interface "evi-17" {
                mac 00:00:00:00:2f:17
                ipv4 {
                    primary {
                        address 10.0.0.2
                        prefix-length 24
                    }
                    neighbor-discovery {
                        timeout 300
                        learn-unsolicited true
                        proactive-refresh true
                        local-proxy-arp true
                        host-route {
                            populate dynamic {
                            }
                        }
                    }
                    vrrp 1 {
                        backup [10.0.0.254]
                        passive true
                        ping-reply true
                        traceroute-reply true
                    }
                }
                vpls "evi-17" {
                    evpn {
                        arp {
                            learn-dynamic false
                            flood-garp-and-unknown-req false
                            advertise dynamic {
                            }
                        }
                    }
                }
            }
        }
```

For IPv4, the behavior is controlled by the following commands.

- **ipv4>neighbor-discovery>host-route>populate [dynamic | evpn | static]** configures PE-2 to advertise host routes. The type of ARP entry that can create a host route can be dynamic, EVPN, static, or a combination of these.
- **ipv4>neighbor-discovery>learn-unsolicited** triggers the learning of an ARP entry upon receiving an ARP or GARP message that was not requested by the router.
- **ipv4>neighbor-discovery>proactive-refresh** triggers the refresh of the ARP entry 30 seconds before aging out.
- **ipv4>neighbor-discovery>local-proxy-arp** ensures that PE-2 replies to any received ARP request on behalf of the other hosts in the R-VPLS broadcast domain.
- **vpls>evpn>arp>learn-dynamic** controls whether data path ARP messages received on EVPN connections can populate the ARP tables.
- **vpls>evpn>arp>flood-garp-and-unknown-req** controls the flooding of Control Processing Module (CPM)-generated ARP requests to EVPN destinations.
- **vpls>evpn>arp>advertise [dynamic | static]** enables PE-2 to advertise MAC and IP in EVPN-MAC routes for ARP entries of the dynamic or static type.

For IPv6, the corresponding commands are as follows:

- **ipv6>neighbor-discovery>host-route>populate [dynamic | evpn | static]**

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- **ipv6>neighbor-discovery>learn-unsolicited [global | link-local | both]**
- **ipv6>neighbor-discovery>proactive-refresh [global | link-local | both]** triggers the refresh of the ND entry upon aging out.
- **ipv6>neighbor-discovery>local-proxy-nd**
- **vpls>evpn>nd>learn-dynamic**
- **vpls>evpn>nd>advertise [dynamic | static]**

For IPv6, CPM-generated Neighbor Solicitation (NS) messages are always flooded to EVPN destinations. This is not configurable in the **vpls>evpn>nd** context of the VPRN service, in contrast to the **flood-garp-and-unknown-req** command in the **vpls>evpn>arp** context for IPv4.

The behavior for IPv6 hosts when moving from a source PE to a target PE is one of the following.

1. The host initiates an unsolicited Neighbor Advertisement (NA).
2. The host sends traffic, without first initiating NA or NS messages.
3. The host does not send any traffic, and the source PE generates an NS message when the MAC address of the host expires and the EVPN-MAC is withdrawn.

All three scenarios are described in more detail later, where the move of host-66 from source PE-2 to target PE-3 is simulated.

Configuration

Figure 186: Example topology with system IP addresses shows the example topology with PE-1, PE-2, and PE-3 in an EVPN-MPLS network and PE-4, PE-5, and PE-6 in an MPLS network.

Figure 186: Example topology with system IP addresses

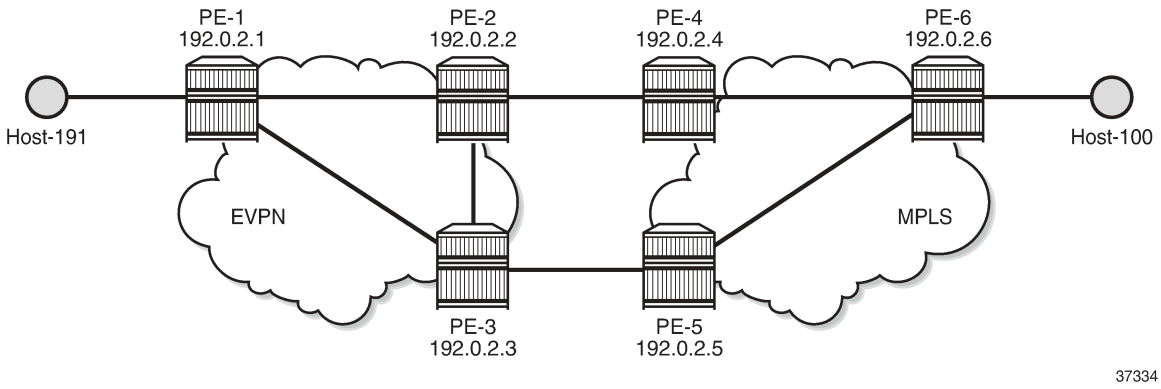

The initial configuration includes:

- Cards, MDAs, ports
- Router interfaces
- IS-IS between PE-1, PE-2, PE-3 and between PE-4, PE-5, and PE-6
- LDP between PE-1, PE-2, PE-3 and between PE-4, PE-5, and PE-6
- BGP configured for the EVPN address family on PE-1, PE-2, and PE-3

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

On PE-1, BGP is configured as follows:

```
# on PE-1:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            family {
                ipv4 false
                evpn true
            }
            rapid-update {
                evpn true
            }
            group "dc" {
                type internal
            }
            neighbor "192.0.2.2" {
                group "dc"
            }
            neighbor "192.0.2.3" {
                group "dc"
            }
        }
```

The BGP configuration is similar on PE-2 and PE-3.

IPv4 host mobility

The following use cases for IPv4 host mobility are described:

1. Host initiates ARP request or GARP after moving
2. Host initiates non-ARP traffic after moving
3. Host does not send any traffic after moving

IPv4 host mobility case 1: host initiates ARP request or GARP after moving

The service configuration on PE-1 is as follows:

```
# on PE-1:
configure {
    service {
        vpls "sbd-15" {
            admin-state enable
            description "R-VPLS 15"
            service-id 15
            customer "1"
            routed-vpls {
            }
            bgp 1 {
                route-distinguisher "192.0.2.1:15"
            }
            bgp-evpn {
                evi 15
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                routes {
                    ip-prefix {
                        advertise true
                    }
                }
                mpls 1 {
                    admin-state enable
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
        }
        vprn "ip-vrf-16" {
            admin-state enable
            service-id 16
            customer "1"
            interface "evi-15" {
                mac 00:00:00:00:00:01
                vpls "sbd-15" {
                    evpn-tunnel {
                    }
                }
            }
            interface "evi-20" {
                mac 00:00:00:00:1e:20
                ipv4 {
                    primary {
                        address 10.0.20.1
                        prefix-length 24
                    }
                }
                vpls "evi-20" {
                }
            }
        }
        vpls "evi-20" {
            admin-state enable
            description "R-VPLS 20"
            service-id 20
            customer "1"
            routed-vpls {
            }
            sap pxc-10.a:20 {
            }
        }
```

VPRN "ip-vrf-16" has two interfaces: interface "evi-15" toward R-VPLS "sbd-15" and interface "evi-20" toward R-VPLS "evi-20". Host-191 is connected to interface "evi-20" of R-VPLS "evi-20".

PE-2 and PE-3 are configured with an anycast gateway, that is, a VRRP passive instance with the same backup IP address 10.0.0.254 on interface "evi-17" in VPRN "ip-vrf-16". The MAC address under VRRP is by default derived from the Virtual Router ID (VRID), so both PE-2 and PE-3 get MAC address 00:00:5E:00:01:01. The service configuration on PE-2 and PE-3 is similar.

```
# on PE-2:
configure {
    service {
        vpls "evi-17" {
            admin-state enable
            description "R-VPLS 17"
            service-id 17
            customer "1"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            routed-vpls {
            }
            bgp 1 {
                route-distinguisher "192.0.2.2:17"        # on PE-3: 192.0.2.3:17
            }
            bgp-evpn {
                evi 17
                mpls 1 {
                    admin-state enable
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            sap 1/1/1:17 {                                # on PE-3: sap 1/1/2:17
            }
        }
        vpls "sbd-15" {
            admin-state enable
            description "R-VPLS 15"
            service-id 15
            customer "1"
            routed-vpls {
            }
            bgp 1 {
                route-distinguisher "192.0.2.2:15"         # on PE-3: 192.0.2.3:15
            }
            bgp-evpn {
                evi 15
                routes {
                    ip-prefix {
                        advertise true
                    }
                }
                mpls 1 {
                    admin-state enable
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
        }
        vprn "ip-vrf-16" {
            admin-state enable
            service-id 16
            customer "1"
            interface "evi-15" {
                mac 00:00:00:00:00:02              # on PE-3: 00:00:00:00:00:03
                vpls "sbd-15" {
                    evpn-tunnel {
                    }
                }
            }
            interface "evi-17" {
                mac 00:00:00:00:2f:17              # on PE-3: 00:00:00:00:3f:17
                ipv4 {
                    primary {
                        address 10.0.0.2           # on PE-3: 10.0.0.3
                        prefix-length 24
                    }
                    neighbor-discovery {
                        timeout 300
                        learn-unsolicited true
                        proactive-refresh true
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        local-proxy-arp true
                        host-route {
                            populate dynamic {
                            }
                        }
                    }
                    vrrp 1 {
                        backup [10.0.0.254]      # anycast IP address on PE-2, PE-3
                        passive true
                        ping-reply true
                        traceroute-reply true
                    }
                }
                vpls "evi-17" {
                    evpn {
                        arp {
                            learn-dynamic false
                            flood-garp-and-unknown-req false
                            advertise dynamic {
                            }
                        }
                    }
                }
            }
        }
```

The **ipv4>neighbor-discovery>host-route>populate dynamic** ensures that route-table ARP-ND host routes are created for dynamic entries, not for static or EVPN entries. The **learn-dynamic false** command prevents PE-2 and PE-3 from learning ARP entries from ARP messages received on an EVPN destination. The **flood-garp-and-unknown-req false** command suppresses CPM-generated ARP to reduce unnecessary ARP flooding.

In this sample topology, an Epipe is used where a failover from the primary to the secondary path simulates a move of host-100 from PE-2 to PE-3. SAP 1/1/1:17 in R-VPLS "evi-17" on PE-2 is connected to a SAP of Epipe 17 on PE-4; SAP 1/1/2:17 in R-VPLS "evi-17" on PE-3 to a SAP of Epipe 17 on PE-5. The service configuration on PE-4 is as follows. The configuration on PE-5 is similar.

```
# on PE-4:
configure {
    service {
        oper-group "op-grp-1" {
        }
        epipe "Epipe 17" {
            admin-state enable
            service-id 17
            customer "1"
            spoke-sdp 46:17 {
                oper-group "op-grp-1"
            }
            sap 1/1/2:17 {
                description "SAP connected to SAP 1/1/1:17 on PE-2"
                monitor-oper-group "op-grp-1"
            }
        }
        sdp 46 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.6
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
```

On PE-6, the service configuration is as follows:

```
# on PE-6:
configure {
    service {
        epipe "Epipe 17" {
            admin-state enable
            service-id 17
            customer "1"
            endpoint "EP17" {
            }
            spoke-sdp 64:17 {
                endpoint {
                    name "EP17"
                    precedence primary
                }
            }
            spoke-sdp 65:17 {
                endpoint {
                    name "EP17"
                }
            }
            sap 1/2/1:17 {
            }
        }
        sdp 64 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.4
            }
        }
        sdp 65 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.5
            }
        }
```

Host-100 is connected to SAP 1/2/1:17 in Epipe 17.

On PE-2 and PE-3, debugging is enabled (in classic CLI):

```
# on PE-2, PE-3:
debug
    router "Base"
        bgp
            update
        exit
    exit
    router service-name "ip-vrf-16"
        ip
            arp
            route-table
        exit
    exit
exit
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Initial phase

Figure 187: Initial situation with forwarding path via PE-2 shows that traffic from host-191 to host-100 is forwarded via PE-1, PE-2, PE-4, and PE-6.

Figure 187: Initial situation with forwarding path via PE-2

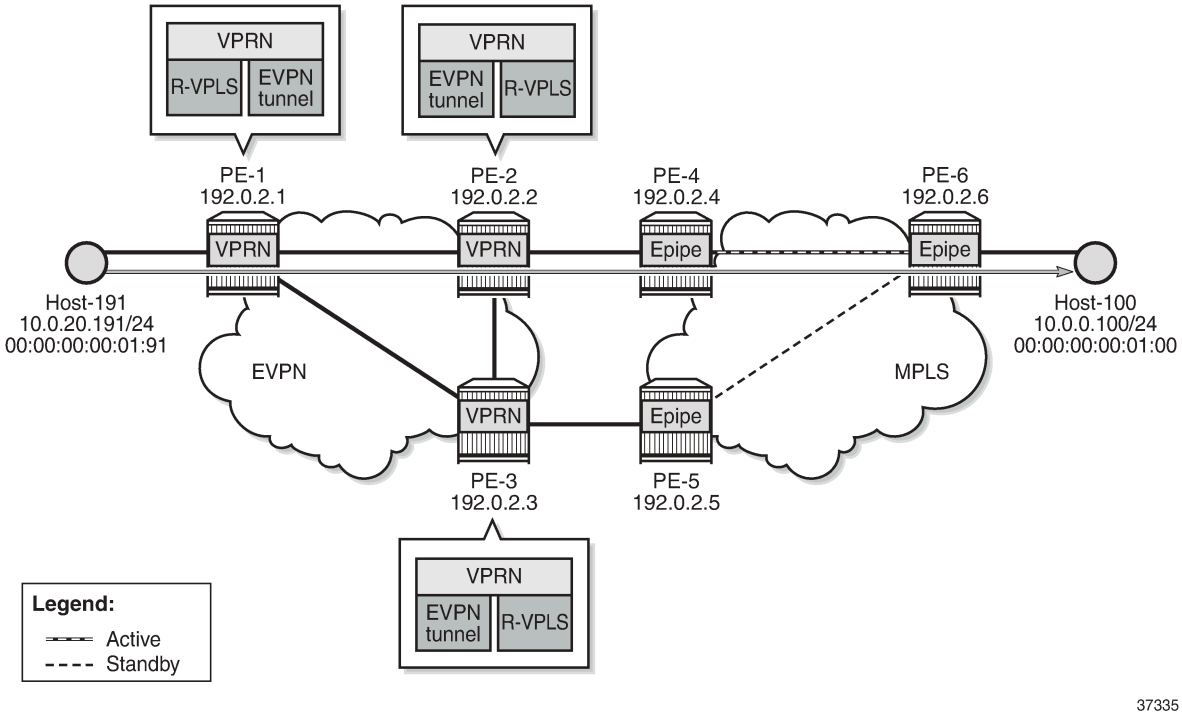

Host-191 sends a traceroute to host-100 via PE-2 (10.0.0.2):

```
[/]
A:admin@PE-1# traceroute 10.0.0.100 router-instance "H-191" source-address 10.0.20.191
traceroute to 10.0.0.100 from 10.0.20.191, 30 hops max, 40 byte packets
  1  10.0.20.1 (10.0.20.1)    3.96 ms  2.27 ms  2.07 ms
  2  10.0.0.2 (10.0.0.2)    2.96 ms  2.95 ms  2.82 ms
  3  10.0.0.100 (10.0.0.100)    10.9 ms  5.54 ms  5.07 ms
```

The ARP table for VPRN "ip-vrf-16" on PE-2 shows that IP address 10.0.0.100 corresponds to MAC address 00:00:00:00:01:00 and is learned dynamically:

```
[/]
A:admin@PE-2# show router 16 arp 10.0.0.100

===============================================================================
ARP Table (Service: 16)
===============================================================================
IP Address      MAC Address       Expiry    Type   Interface
-------------------------------------------------------------------------------
10.0.0.100      00:00:00:00:01:00 00h04m49s Dyn[I] evi-17
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The ARP table for VPRN "ip-vrf-16" on PE-3 shows that IP address 10.0.0.100 is advertised through EVPN:

```
[/]
A:admin@PE-3# show router 16 arp 10.0.0.100

===============================================================================
ARP Table (Service: 16)
===============================================================================
IP Address      MAC Address       Expiry    Type   Interface
-------------------------------------------------------------------------------
10.0.0.100      00:00:00:00:01:00 00h00m00s Evp[I] evi-17
===============================================================================
```

On PE-2, MAC address 00:00:00:00:01:00 is learned on SAP 1/1/1:17 in R-VPLS "evi-17":

```
[/]
A:admin@PE-2# show service id 17 fdb detail

===============================================================================
Forwarding Database, Service 17
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
17         00:00:00:00:01:00 sap:1/1/1:17            L/30     01/28/22 12:45:29
17         00:00:00:00:2e:17 cpm                     Intf     01/28/22 12:43:20
17         00:00:00:00:3f:17 mpls-1:                 EvpnS:P  01/28/22 12:43:29
                             192.0.2.3:524283
           ldp:65538
17         00:00:5e:00:01:01 cpm                     Intf     01/28/22 12:43:20
-------------------------------------------------------------------------------
No. of MAC Entries: 4
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

On PE-3, the FDB for R-VPLS "evi-17" shows that MAC address 00:00:00:00:01:00 is advertised through EVPN:

```
[/]
A:admin@PE-3# show service id 17 fdb mac 00:00:00:00:01:00

===============================================================================
Forwarding Database, Service 17
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
17         00:00:00:00:01:00 mpls-1:                 Evpn     01/28/22 12:45:29
                             192.0.2.2:524283
           ldp:65537
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

The route table for VPRN "ip-vrf-16" on PE-2 shows an ARP-ND host route with preference 1 for prefix 10.0.0.100/32:

```
[/]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
A:admin@PE-2# show router 16 route-table 10.0.0.100

===============================================================================
Route Table (Service: 16)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
10.0.0.100/32                                 Remote  ARP-ND    00h02m44s  1
       10.0.0.100                                                 0
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

The route table for VPRN "ip-vrf-16" on PE-3 shows an EVPN Interface-ful (EVPN-IFF) host route for prefix 10.0.0.100/32 with preference 169:

```
[/]
A:admin@PE-3# show router 16 route-table 10.0.0.100

===============================================================================
Route Table (Service: 16)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
10.0.0.100/32                                 Remote  EVPN-IFF  00h02m43s  169
       evi-15 (ET-00:00:00:00:00:02)                              0
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

PE-3 receives the IP and MAC addresses of host-100 as EVPN type. PE-3 must not learn these IP and MAC addresses as dynamic because PE-3 must be prevented from advertising an RT5 route. If PE-3 advertised prefix 10.0.0.100, then PE-1 could select PE-3 as next hop to reach host-100, causing an undesired hairpinning forwarding behavior.

Host-100 sends an ARP request or GARP after moving

Figure 188: Host-100 sends an ARP request or GARP after switchover shows a switchover from the active to the standby path where host-100 sends an ARP request or GARP and its IP and MAC addresses are learned on PE-3 instead of PE-2. The failure is simulated by disabling the SDP from PE-4 to PE-6.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 188: Host-100 sends an ARP request or GARP after switchover

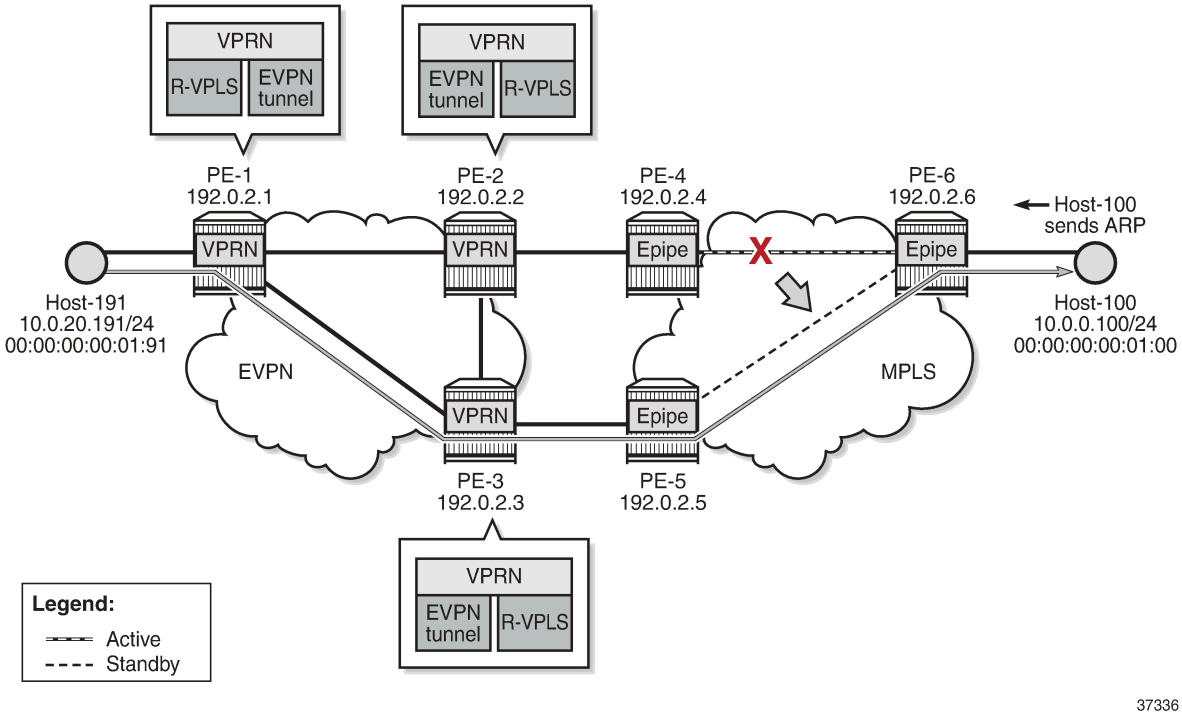

Due to the SDP failure on PE-4, the initial path can no longer be used. Host-100 sends an ARP request or GARP with its IP and MAC addresses. In the following example, PE-3 receives the following ARP request and replies to it:

```
1 2022/01/28 12:49:45.684 UTC MINOR: DEBUG #2001 vprn16 PIP
"PIP: ARP
instance 2 (16), interface index 6 (evi-17),
ARP ingressing on evi-17
   Who has 10.0.0.254 ? Tell 10.0.0.100
"
2 2022/01/28 12:49:45.684 UTC MINOR: DEBUG #2001 vprn16 PIP
"PIP: ARP
instance 2 (16), interface index 6 (evi-17),
ARP egressing on evi-17
   10.0.0.254 is at 00:00:5e:00:01:01
"
```

The Route Table Manager (RTM) for prefix 10.0.0.100 in VPRN "ip-vrf-16" is modified with preference 1 and owner ARP-ND. This behavior is due to the **ipv4>neighbor-discovery>host-route>populate dynamic** command.

```
3 2022/01/28 12:49:45.684 UTC MINOR: DEBUG #2001 vprn16 PIP
"PIP: ROUTE
instance 2 (16), RTM MODIFY event
   New Route Info
      prefix: 10.0.0.100/32 (0x119549018)  preference: 1   metric: 0
                            backup metric: 0  owner: ARP-ND ownerId: 0
      1 ecmp hops  0 backup hops:
         hop 0:  10.0.0.100 @ if 6, weight 0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
"
```

PE-3 sends an RT5 for prefix 10.0.0.100/32 to PE-1 and PE-2:

```
4 2022/01/28 12:49:45.685 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 90
    Flag: 0x90 Type: 14 Len: 45 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.3
        Type: EVPN-IP-PREFIX Len: 34 RD: 192.0.2.3:15, tag: 0,
                  ip_prefix: 10.0.0.100/32 gw_ip 0.0.0.0
                  Label: 8388544 (Raw Label: 0x7fffc0)
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:15
        mac-nh:00:00:00:00:00:03
        bgp-tunnel-encap:MPLS
"
```

PE-3 sends EVPN-MAC routes for MAC 00:00:00:00:01:00 with an increased sequence number for MAC mobility: one EVPN-MAC route with MAC address 00:00:00:00:01:00 and IP address 10.0.0.100 and another EVPN-MAC route with MAC address 00:00:00:00:01:00 only and a null IP address.

```
5 2022/01/28 12:49:45.685 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 128
    Flag: 0x90 Type: 14 Len: 83 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.3
        Type: EVPN-MAC Len: 37 RD: 192.0.2.3:17 ESI: ESI-0, tag: 0, mac len: 48
                  mac: 00:00:00:00:01:00, IP len: 4, IP: 10.0.0.100, label1: 8388528
        Type: EVPN-MAC Len: 33 RD: 192.0.2.3:17 ESI: ESI-0, tag: 0, mac len: 48
                  mac: 00:00:00:00:01:00, IP len: 0, IP: NULL, label1: 8388528
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:17
        bgp-tunnel-encap:MPLS
        mac-mobility:Seq:1
"
```

The FDB for R-VPLS "evi-17" shows that MAC address 00:00:00:00:01:00 is dynamically learned on SAP 1/1/2:17 on PE-3:

```
[/]
A:admin@PE-3# show service id 17 fdb mac 00:00:00:00:01:00

===============================================================================
Forwarding Database, Service 17
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
17         00:00:00:00:01:00 sap:1/1/2:17           L/14     01/28/22 12:49:46
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

On PE-2, the FDB for R-VPLS "evi-17" is updated and PE-2 withdraws its EVPN-MAC route based on the higher sequence number of the received EVPN-MAC route for MAC address 00:00:00:00:01:00 with next hop 192.0.2.3:

```
[/]
A:admin@PE-2# show service id 17 fdb mac 00:00:00:00:01:00

===============================================================================
Forwarding Database, Service 17
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
17         00:00:00:00:01:00 mpls-1:                 Evpn     01/28/22 12:49:46
                             192.0.2.3:524283
           ldp:65538
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

On PE-3, the ARP table for VPRN "ip-vrf-16" shows that IP address 10.0.0.100 is learned dynamically on interface "evi-17":

```
[/]
A:admin@PE-3# show router 16 arp 10.0.0.100

===============================================================================
ARP Table (Service: 16)
===============================================================================
IP Address      MAC Address       Expiry    Type   Interface
-------------------------------------------------------------------------------
10.0.0.100      00:00:00:00:01:00 00h04m40s Dyn[I] evi-17
===============================================================================
```

On PE-2, the ARP table for VPRN "ip-vrf-16" shows that the entry for IP address 10.0.0.100 is updated from dynamic to type EVPN:

```
[/]
A:admin@PE-2# show router 16 arp 10.0.0.100

===============================================================================
ARP Table (Service: 16)
===============================================================================
IP Address      MAC Address       Expiry    Type   Interface
-------------------------------------------------------------------------------
10.0.0.100      00:00:00:00:01:00 00h00m00s Evp[I] evi-17
===============================================================================
```

An ARP entry's change from dynamic to EVPN triggers a CPM-generated ARP request from PE-2, but the configured **flood-garp-and-unknown-req false** command prevents PE-2 from flooding the ARP request to EVPN destinations such as PE-3.

On PE-3, the route table for VPRN "ip-vrf-16" shows an ARP-ND host route for prefix 10.0.0.100:

```
[/]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
A:admin@PE-3# show router 16 route-table 10.0.0.100

===============================================================================
Route Table (Service: 16)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
10.0.0.100/32                                 Remote  ARP-ND    00h01m32s  1
       10.0.0.100                                                 0
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

The route table for VPRN "ip-vrf-16" for prefix 10.0.0.100 shows that PE-2 removed its ARP-ND host route and the received EVPN route from PE-3 is used instead:

```
[/]
A:admin@PE-2# show router 16 route-table 10.0.0.100

===============================================================================
Route Table (Service: 16)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
10.0.0.100/32                                 Remote  EVPN-IFF  00h01m31s  169
       evi-15 (ET-00:00:00:00:00:03)                              0
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

IPv4 host mobility case 2: host sends traffic without first initiating an ARP request or GARP after moving

In use cases 2 and 3, the configuration of VPRN "ip-vrf-16" is modified on PE-2 and PE-3. The only difference from case 1 is that the **flood-garp-and-unknown-req** is configured, which is the default setting. The VPRN configuration on PE-2 is as follows:

```
# on PE-2:
configure {
    service {
        vprn "ip-vrf-16" {
            admin-state enable
            service-id 16
            customer "1"
            interface "evi-15" {
                mac 00:00:00:00:00:02
                vpls "sbd-15" {
                    evpn-tunnel {
                    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
        }
        interface "evi-17" {
            mac 00:00:00:00:2f:17
            ipv4 {
                primary {
                    address 10.0.0.2
                    prefix-length 24
                }
                neighbor-discovery {
                    timeout 300
                    learn-unsolicited true
                    proactive-refresh true
                    local-proxy-arp true
                    host-route {
                        populate dynamic {
                        }
                    }
                }
                vrrp 1 {
                    backup [10.0.0.254]
                    passive true
                    ping-reply true
                    traceroute-reply true
                }
            }
            vpls "evi-17" {
                evpn {
                    arp {
                        learn-dynamic false
                        flood-garp-and-unknown-req true  # default behavior
                        advertise dynamic {
                        }
                    }
                }
            }
        }
```

Initial forwarding path

The initial forwarding path via PE-2 is restored by enabling the SDP from PE-4 to PE-6. The route table for VPRN "ip-vrf-16" on PE-2 shows the following ARP-ND host route for prefix 10.0.0.100/32:

```
[/]
A:admin@PE-2# show router 16 route-table 10.0.0.100

===============================================================================
Route Table (Service: 16)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
10.0.0.100/32                                 Remote  ARP-ND    00h03m46s  1
       10.0.0.100                                                 0
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
```

Host-100 generates non-ARP traffic after moving

On PE-4, the SDP from PE-4 to PE-6 is disabled, causing a switchover to the standby path. Figure 189: Host sends non-ARP frame after switchover shows the path after switchover. Host-100 generates non-ARP traffic after moving.

Figure 189: Host sends non-ARP frame after switchover

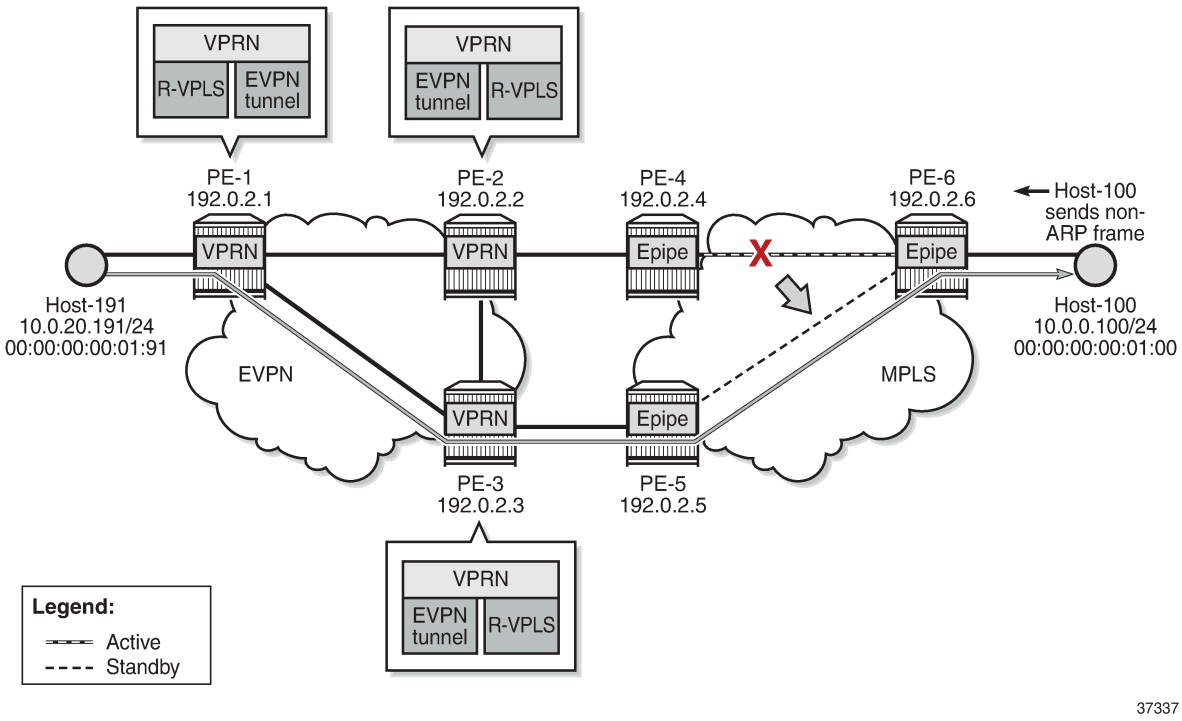

Host-100 sends a non-ARP frame with MAC source address 00:00:00:00:01:00 to host-191. The following steps occur:

1. PE-3 receives this frame with MAC 00:00:00:00:01:00 and updates its FDB.
2. PE-3 advertises an EVPN-MAC route for MAC 00:00:00:00:01:00 (with a null IP address) with a higher sequence number.
3. PE-2 receives this EVPN MAC route, updates its FDB and withdraws its EVPN-MAC routes for MAC 00:00:00:00:01:00.
4. The FDB update for MAC 00:00:00:00:01:00 triggers PE-2 to send an ARP request for MAC 00:00:00:00:01:00.
5. PE-2 is configured with **flood-garp-and-unknown-req**, so the ARP request is flooded to the EVPN destinations PE-1 and PE-3. PE-3 floods this ARP request to its SAPs and SDP-bindings; in this case, to SAP 1/1/2:17.
6. When the ARP request reaches host-100, it sends an ARP reply to the anycast IP address 10.0.0.254. This ARP reply is received by PE-3.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

7. When PE-3 receives the ARP reply, it updates the ARP entry for 10.0.0.100 to type dynamic instead of type EVPN.
8. PE-3 is configured with **populate dynamic**, so it advertises an RT5 for prefix 10.0.0.100/32. Also, MAC 00:00:00:00:01:00 is now learned in ARP as local, so PE-3 sends an EVPN-MAC route with MAC 00:00:00:00:01:00 and IP prefix 10.0.0.100.
9. PE-2 receives the EVPN routes and updates the ARP entry for prefix 10.0.0.100 from type dynamic to type EVPN. PE-2 also removes its ARP-ND host route from the route table and withdraws its RT5 for prefix 10.0.0.100/32.

On PE-3, the route for prefix 10.0.0.100/32 is an ARP-ND host route:

```
[/]
A:admin@PE-3# show router 16 route-table 10.0.0.100

===============================================================================
Route Table (Service: 16)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
10.0.0.100/32                                 Remote  ARP-ND    00h01m05s  1
       10.0.0.100                                                  0
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

IPv4 host mobility case 3: host does not send any traffic after moving

The service configuration on PE-2 and PE-3 remains the same as in use case 2.

The forwarding path is restored by enabling the SDP from PE-4 to PE-6, so the initial situation is the same as in the preceding cases. PE-2 learns MAC address 00:00:00:00:01:00 on its local SAP 1/1/1:17, as follows:

```
[/]
A:admin@PE-2# show service id 17 fdb detail

===============================================================================
Forwarding Database, Service 17
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
17         00:00:00:00:01:00 sap:1/1/1:17            L/0      01/28/22 13:10:34
17         00:00:00:00:2f:17 cpm                     Intf     01/28/22 12:53:11
17         00:00:00:00:3f:17 mpls-1:                 EvpnS:P  01/28/22 12:43:29
                             192.0.2.3:524283
           ldp:65538
17         00:00:5e:00:01:01 cpm                     Intf     01/28/22 12:43:20
-------------------------------------------------------------------------------
No. of MAC Entries: 4
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
```

The SDP from PE-4 to PE-6 is disabled and host-100 does not send any traffic, as shown in Figure 190: Host does not send any traffic after switchover.

Figure 190: Host does not send any traffic after switchover

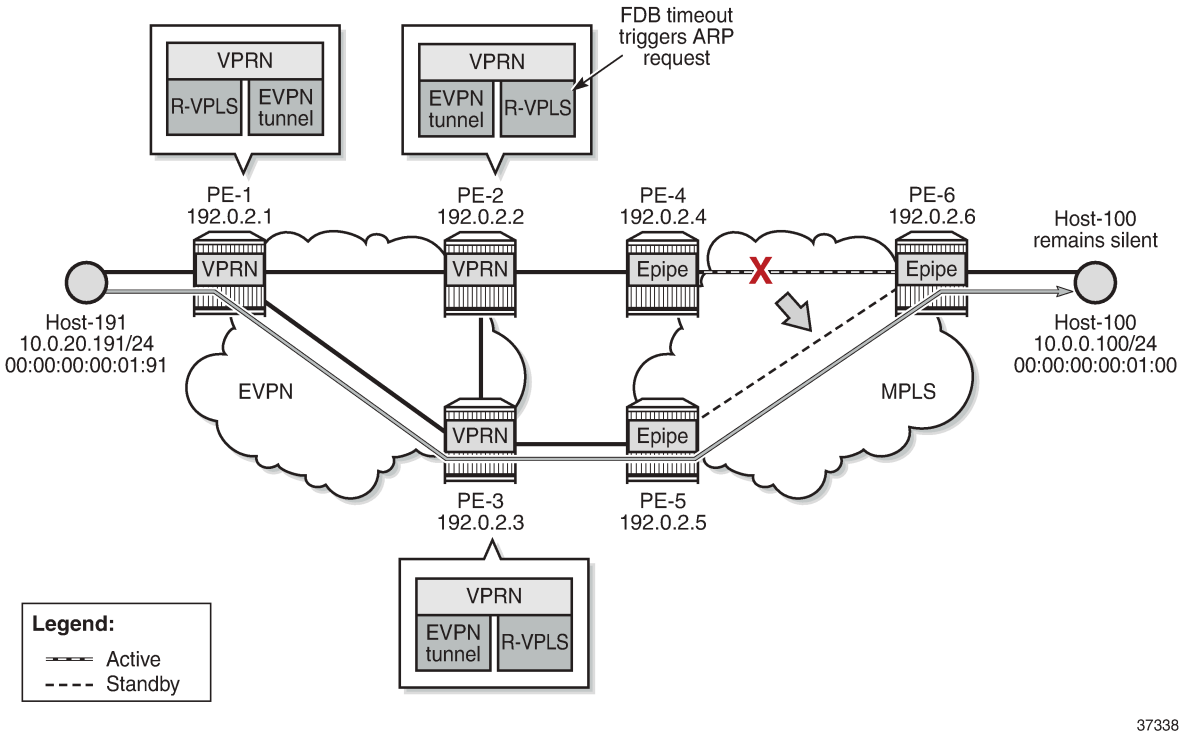

The following steps occur:

1. When MAC 00:00:00:00:01:00 ages out in the FDB of R-VPLS 17 on PE-2, PE-2 withdraws the EVPN-MAC routes for MAC 00:00:00:00:01:00. The update for MAC 00:00:00:00:01:00 triggers PE-2 to send an ARP request for 10.0.0.100.

```
# on PE-2:
99 2022/01/28 13:12:04.028 UTC MINOR: DEBUG #2001 vprn16 PIP
"PIP: ARP
instance 2 (16), interface index 6 (evi-17),
ARP egressing on evi-17
   Who has 10.0.0.100 ? Tell 10.0.0.254
"
```

```
101 2022/01/28 13:12:04.028 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 42
    Flag: 0x90 Type: 15 Len: 38 Multiprotocol Unreachable NLRI:
        Address Family EVPN
        Type: EVPN-MAC Len: 33 RD: 192.0.2.2:17 ESI: ESI-0, tag: 0, mac len: 48
                       mac: 00:00:00:00:01:00, IP len: 0, IP: NULL, label1: 0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
"
```

2. PE-2 is configured with **flood-garp-and-unknown-req true**. PE-2 floods the CPM-generated ARP request to PE-3. PE-3 forwards the ARP request to host-100.
3. Host-100 sends an ARP reply that is received by PE-3. PE-3 updates its FDB and ARP tables.
4. The FDB update on PE-3 makes PE-3 advertise an EVPN-MAC route for MAC 00:00:00:00:01:00 (with a null IP address). The ARP update makes PE-3 advertise an EVPN-MAC route with MAC 00:00:00:00:01:00 and IP prefix 10.0.0.100. PE-2 receives two EVPN-MAC routes from PE-3:

```
103 2022/01/28 13:12:04.033 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 128
    Flag: 0x90 Type: 14 Len: 83 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.3
        Type: EVPN-MAC Len: 37 RD: 192.0.2.3:17 ESI: ESI-0, tag: 0, mac len: 48
                  mac: 00:00:00:00:01:00, IP len: 4, IP: 10.0.0.100, label1: 8388528
        Type: EVPN-MAC Len: 33 RD: 192.0.2.3:17 ESI: ESI-0, tag: 0, mac len: 48
                  mac: 00:00:00:00:01:00, IP len: 0, IP: NULL, label1: 8388528
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:17
        bgp-tunnel-encap:MPLS
        mac-mobility:Seq:5
"
```

5. PE-3 is configured with **populate dynamic**, so it advertises an RT5 for prefix 10.0.0.100/32. In the route table for VPRN "ip-vrf-16", the route for IP prefix 10.0.0.100/32 is ARP-ND host route. PE-2 receives the following RT5 route from PE-3:

```
102 2022/01/28 13:12:04.033 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 90
    Flag: 0x90 Type: 14 Len: 45 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.3
        Type: EVPN-IP-PREFIX Len: 34 RD: 192.0.2.3:15, tag: 0,
         ip_prefix: 10.0.0.100/32 gw_ip 0.0.0.0 Label: 8388544 (Raw Label: 0x7fffc0)
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:15
        mac-nh:00:00:00:00:00:03
        bgp-tunnel-encap:MPLS
"
```

6. PE-2 receives the EVPN routes and updates its FDB and ARP tables. When the ARP entry changes its type from dynamic to EVPN, PE-2 withdraws its RT5 route.

On PE-2, the FDB for R-VPLS 17 shows an EVPN route for MAC 00:00:00:00:01:00:

```
[/]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
A:admin@PE-2# show service id 17 fdb detail

===============================================================================
Forwarding Database, Service 17
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
17         00:00:00:00:01:00 mpls-1:                 Evpn     01/28/22 13:12:04
                             192.0.2.3:524283
           ldp:65538
17         00:00:00:00:2f:17 cpm                     Intf     01/28/22 12:53:11
17         00:00:00:00:3f:17 mpls-1:                 EvpnS:P  01/28/22 12:43:29
                             192.0.2.3:524283
           ldp:65538
17         00:00:5e:00:01:01 cpm                     Intf     01/28/22 12:43:20
-------------------------------------------------------------------------------
No. of MAC Entries: 4
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

IPv6 host mobility

The following use cases for IPv6 host mobility are described:

1. Host initiates an unsolicited NA message after moving
2. Host sends non-ND traffic after moving
3. Host does not send any traffic after moving

The configuration is identical in these use cases.

Figure 191: Example topology for initial forwarding path via PE-2 with IPv6 addresses shows the topology with IPv6 addresses for host-61 and host-66.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 191: Example topology for initial forwarding path via PE-2 with IPv6 addresses

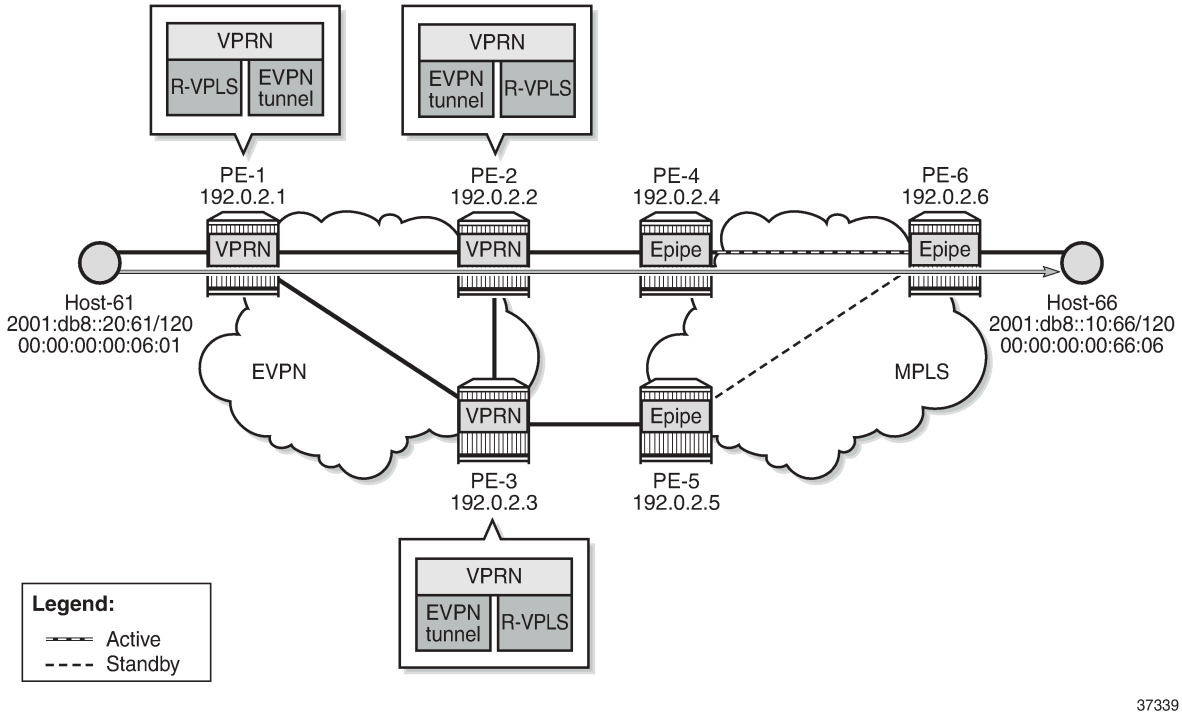

The services are the following:

- R-VPLS "sbd-5" on PE-1, PE-2, and PE-3
- VPRN "ip-vrf-6" on PE-1, PE-2, and PE-3
- R-VPLS "evi-10" on PE-1; R-VPLS "evi-7" on PE-2 and PE-3
- Epipe "Epipe 7" on PE-4, PE-5, and PE-6
- Host-61 is connected to R-VPLS "evi-10" on PE-1
- Host-66 is connected to Epipe "Epipe 7" on PE-6

The service configuration on PE-1 is as follows:

```
# on PE-1:
configure {
    service {
        vpls "evi-10" {
            admin-state enable
            description "R-VPLS 10"
            service-id 10
            customer "1"
            routed-vpls {
            }
            sap pxc-10.a:10 {
            }
        }
        vpls "sbd-5" {
            admin-state enable
            description "R-VPLS 5"
            service-id 5
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            customer "1"
            routed-vpls {
            }
            bgp 1 {
                route-distinguisher "192.0.2.1:5"
            }
            bgp-evpn {
                evi 5
                routes {
                    ip-prefix {
                        advertise true
                    }
                }
                mpls 1 {
                    admin-state enable
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
        }
        vprn "ip-vrf-6" {
            admin-state enable
            service-id 6
            customer "1"
            interface "evi-10" {
                mac 00:00:00:06:1e:20
                vpls "evi-10" {
                }
                ipv6 {
                    address 2001:db8::20:1 {
                        prefix-length 120
                    }
                }
            }
            interface "evi-5" {
                mac 00:00:00:00:06:01
                vpls "sbd-5" {
                    evpn-tunnel {
                    }
                }
                ipv6 {
                }
            }
        }
```

The service configuration on PE-2 is as follows. The service configuration on PE-3 is similar.

```
# on PE-2:
configure {
    service {
        vpls "sbd-5" {
            admin-state enable
            description "R-VPLS 5"
            service-id 5
            customer "1"
            routed-vpls {
            }
            bgp 1 {
                route-distinguisher "192.0.2.2:5"         # on PE-3: 192.0.2.3:5
            }
            bgp-evpn {
                evi 5
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                routes {
                    ip-prefix {
                        advertise true
                    }
                }
                mpls 1 {
                    admin-state enable
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
        }
        vprn "ip-vrf-6" {
            admin-state enable
            service-id 6
            customer "1"
            interface "evi-5" {
                mac 00:00:00:00:06:02                   # on PE-3: 00:00:00:00:06:03
                vpls "sbd-5" {
                    evpn-tunnel {
                    }
                }
                ipv6 {
                }
            }
            interface "evi-7" {
                mac 00:00:00:00:2f:07                   # on PE-3: 00:00:00:00:3f:07
                vpls "evi-7" {
                    evpn {
                        nd {
                            learn-dynamic false
                            advertise dynamic {
                            }
                        }
                    }
                }
                ipv6 {
                    link-local-address {
                        address fe80::10:2           # on PE-3: fe80::10:3
                        duplicate-address-detection false
                    }
                    address 2001:db8::10:2 {         # on PE-3: 2001:db8::10:3
                        prefix-length 120
                    }
                    neighbor-discovery {
                        learn-unsolicited both
                        proactive-refresh both
                        local-proxy-nd true
                        host-route {
                            populate dynamic {
                            }
                        }
                    }
                    vrrp 1 {
                        backup [fe80::10:fe]
                        passive true
                        ping-reply true
                        traceroute-reply true
                    }
                }
            }
        }
        vpls "evi-7" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            admin-state enable
            description "R-VPLS 7"
            service-id 7
            customer "1"
            routed-vpls {
            }
            bgp 1 {
                route-distinguisher "192.0.2.2:7"        # on PE-3: 192.0.2.3:7
            }
            bgp-evpn {
                evi 7
                mpls 1 {
                    admin-state enable
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            sap 1/1/1:7 {                                # on PE-3: sap 1/1/2:7
            }
        }
```

Debugging is enabled on PE-2 and PE-3 (in classic CLI):

```
# on PE-2, PE-3:
debug
    router "Base"
        bgp
            update
        exit
    exit
    router service-name "ip-vrf-6"
        ip
            route-table
            icmp6 "evi-7"
            neighbor "evi-7"
        exit
    exit
exit
```

Initially, the traceroute from host-66 to host-61 is via PE-2 (2001:db8::10:2):

```
[/]
A:admin@PE-6# traceroute 2001:db8::20:61 router-instance "H-66"
traceroute to 2001:db8::20:61, 30 hops max, 60 byte packets
  1  2001:db8::10:2 (2001:db8::10:2)    1034 ms  2.80 ms  4.17 ms
  2  ::  * * *
  3  2001:db8::20:61 (2001:db8::20:61)    11.1 ms  5.59 ms  4.95 ms
```

The following route table on PE-2 shows an ARP-ND host route for prefix 2001:db8::10:66/128:

```
[/]
A:admin@PE-2# show router 6 route-table 2001:db8::10:66

===============================================================================
IPv6 Route Table (Service: 6)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
2001:db8::10:66/128                           Remote  ARP-ND    00h00m33s  1
       2001:db8::10:66                                             0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

IPv6 host mobility case 1: host initiates an unsolicited NA message after moving

On PE-2 and PE-3, the **learn-unsolicited** command is configured on interface "evi-7" in VPRN "ip-vrf-6". When an unsolicited NA message is received, a stale neighbor is created. If **host-route>populate dynamic** is enabled, a confirmation message is sent for all the neighbor entries created as stale, and if confirmed, the corresponding ARP-ND routes are added to the route table.

Disabling SDP 46 on PE-4 causes a failover from the primary path via PE-4 to the secondary path via PE-5, simulating host-66 moving from PE-2 to PE-3. To trigger an unsolicited NA message from host-66, its MAC address 00:00:00:00:66:06 is replaced by MAC address 00:00:00:00:06:06. Figure 192: Host-66 sends unsolicited NA message after switchover shows that host-66 sends an unsolicited NA message.

Figure 192: Host-66 sends unsolicited NA message after switchover

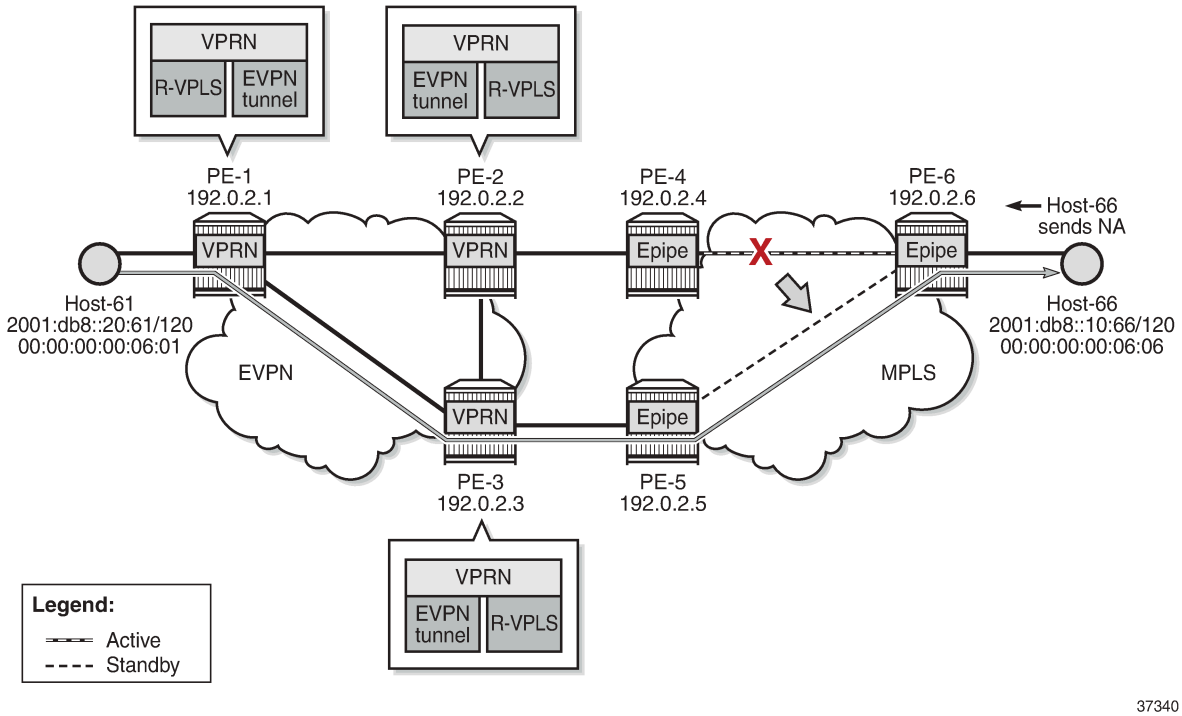

Host-66 advertises its new MAC address in unsolicited NA messages. PE-3 receives the following NA messages from host-66. PE-2 also receives the NA messages, but it rejects NA messages received on interface "evi-7" when **learn-dynamic false** is configured:

```
# on PE-3:
3 2022/01/28 13:19:55.747 UTC MINOR: DEBUG #2001 vprn6 TIP
"TIP: ICMP6_PKT
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
ICMP6 ingressing on evi-7 (vprn6):
   fe80::10:6 -> ff02::1
   Type: Neighbor Advertisement (136)
   Code: No Code (0)
      Tgt Addr: 2001:db8::10:66
      Flags   : Router Override
      Option  : Tgt Link Layer Addr 00:00:00:00:06:06
"
```

```
1 2022/01/28 13:19:55.747 UTC MINOR: DEBUG #2001 vprn6 TIP
"TIP: ICMP6_PKT
ICMP6 ingressing on evi-7 (vprn6):
   fe80::10:6 -> ff02::1
   Type: Neighbor Advertisement (136)
   Code: No Code (0)
      Tgt Addr: fe80::10:6
      Flags   : Router Override
      Option  : Tgt Link Layer Addr 00:00:00:00:06:06
"
```

PE-3 learns the MAC address dynamically:

```
[/]
A:admin@PE-3# show service id 7 fdb mac 00:00:00:00:06:06

===============================================================================
Forwarding Database, Service 7
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
7          00:00:00:00:06:06 sap:1/1/2:7             L/90     01/28/22 13:19:56
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

PE-3 sends CPM-generated NS messages that are also flooded to the EVPN destinations. The **learn-dynamic false** command prevents PE-2 from learning MAC addresses dynamically on an EVPN connection.

PE-3 sends an EVPN-MAC update to PE-2 and MAC address 00:00:00:00:06:06 appears in the FDB on PE-2 as an EVPN entry:

```
[/]
A:admin@PE-2# show service id 7 fdb mac 00:00:00:00:06:06

===============================================================================
Forwarding Database, Service 7
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
7          00:00:00:00:06:06 mpls-1:                 Evpn     01/28/22 13:19:56
                             192.0.2.3:524281
           ldp:65538
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The route table for VPRN "ip-vrf-6" on PE-3 shows an ARP-ND entry for destination prefix 2001:db8::10:66/128, as follows:

```
[/]
A:admin@PE-3# show router 6 route-table 2001:db8::10:66

===============================================================================
IPv6 Route Table (Service: 6)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
2001:db8::10:66/128                           Remote  ARP-ND    00h02m37s  1
       2001:db8::10:66                                            0
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

On PE-2, the route table for VPRN "ip-vrf-6" shows an EVPN entry for prefix 2001:db8::10:66/128:

```
[/]
A:admin@PE-2# show router 6 route-table 2001:db8::10:66

===============================================================================
IPv6 Route Table (Service: 6)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
2001:db8::10:66/128                           Remote  EVPN-IFF  00h02m39s  169
       fe80::7:b0d1:3fa3:2f60-"evi-5"                             0
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

IPv6 host mobility case 2: host sends non-ND traffic after moving,

The service configuration is the same as in the use case 1. The only difference from use case 1 is the type of message that is sent by host-66 after moving.

Initially, the traceroute from host-66 to host-61 is via PE-2 (2001:db8::10:2):

```
[/]
A:admin@PE-6# traceroute 2001:db8::20:61 router-instance "H-66"
traceroute to 2001:db8::20:61, 30 hops max, 60 byte packets
  1  2001:db8::10:2 (2001:db8::10:2)    7.88 ms  3.85 ms  3.67 ms
  2  ::  * * *
  3  2001:db8::20:61 (2001:db8::20:61)    11.0 ms  5.19 ms  5.30 ms
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

A switchover from the primary path to the secondary path takes place, so host-66 moves from PE-2 to PE-3. Figure 193: Host generates non-ND traffic after switchover shows that host-66 sends non-ND traffic after moving.

Figure 193: Host generates non-ND traffic after switchover

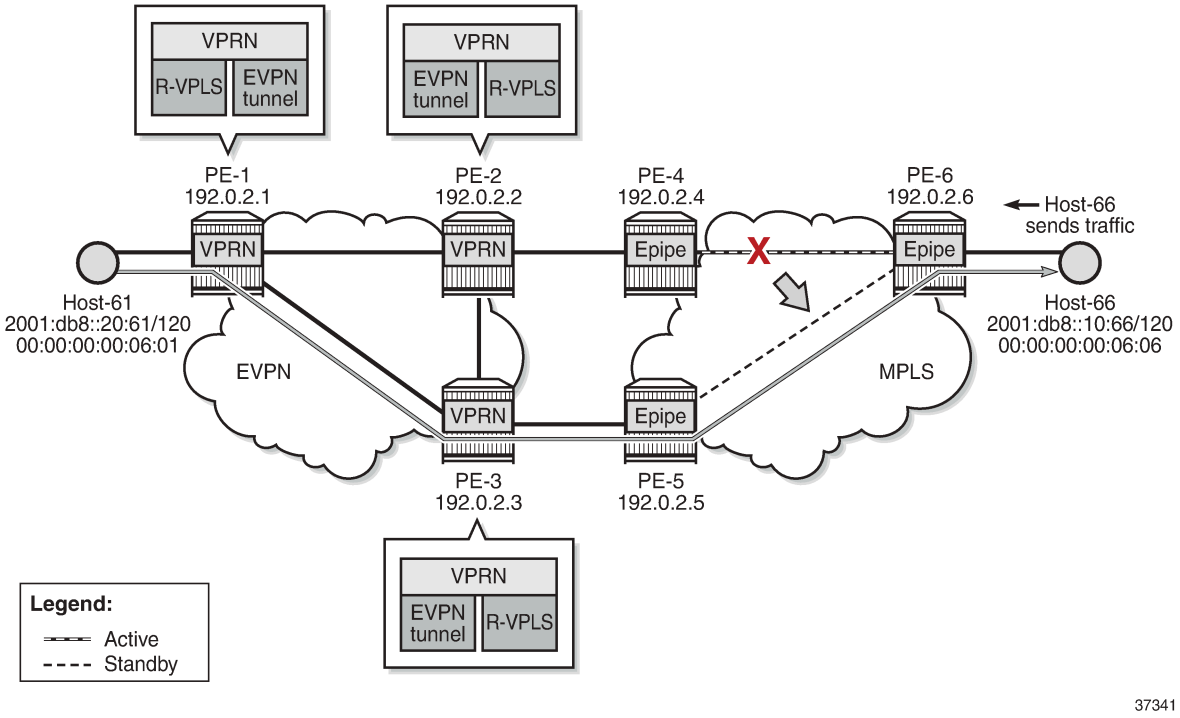

The traceroute from host-66 to host-61 is via PE-3 (2001:db8::10:3) instead of PE-2, as follows:

```
[/]
A:admin@PE-6# traceroute 2001:db8::20:61 router-instance "H-66"
traceroute to 2001:db8::20:61 from 2001:db8::10:66, 30 hops max, 60 byte packets
  1  2001:db8::10:3 (2001:db8::10:3)    7.68 ms  3.84 ms  3.65 ms
  2  ::  * * *
  3  2001:db8::20:61 (2001:db8::20:61)    5.54 ms  5.55 ms  5.64 ms
```

On PE-3, MAC address 00:00:00:00:06:06 from host-66 is learned on the local SAP 1/1/2:17:

```
[/]
A:admin@PE-3# show service id 7 fdb mac 00:00:00:00:06:06

===============================================================================
Forwarding Database, Service 7
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
7          00:00:00:00:06:06 sap:1/1/2:7             L/0      01/28/22 13:29:59
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

PE-3 advertises MAC address 00:00:00:00:06:06 from host-66 in three EVPN-MAC routes: one with the global IP address 2001:db8::10:66, one with the link local IP address fe80::200:ff:fe00:606, and one with a null IP address. PE-2 receives the following EVPN-MAC routes from PE-3:

```
# on PE-2:
59 2022/01/28 13:29:59.437 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 191
    Flag: 0x90 Type: 14 Len: 146 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.3
        Type: EVPN-MAC Len: 49 RD: 192.0.2.3:7 ESI: ESI-0, tag: 0, mac len: 48
            mac: 00:00:00:00:06:06, IP len: 16, IP: fe80::10:6, label1: 8388496
        Type: EVPN-MAC Len: 49 RD: 192.0.2.3:7 ESI: ESI-0, tag: 0, mac len: 48
            mac: 00:00:00:00:06:06, IP len: 16, IP: 2001:db8::10:66, label1: 8388496
        Type: EVPN-MAC Len: 33 RD: 192.0.2.3:7 ESI: ESI-0, tag: 0, mac len: 48
            mac: 00:00:00:00:06:06, IP len: 0, IP: NULL, label1: 8388496
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:7
        bgp-tunnel-encap:MPLS
        mac-mobility:Seq:2
"
```

On PE-2, the following EVPN entry for MAC 00:00:00:00:06:06 is added to the FDB:

```
[/]
A:admin@PE-2# show service id 7 fdb mac 00:00:00:00:06:06

===============================================================================
Forwarding Database, Service 7
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
7          00:00:00:00:06:06 mpls-1:                 Evpn     01/28/22 13:29:59
                             192.0.2.3:524281
           ldp:65538
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

The route table on PE-3 shows an ARP-ND host route for prefix 2001:db8::10:66/128:

```
[/]
A:admin@PE-3# show router 6 route-table 2001:db8::10:66

===============================================================================
IPv6 Route Table (Service: 6)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
2001:db8::10:66/128                           Remote  ARP-ND    00h01m38s  1
       2001:db8::10:66                                            0
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

PE-2 receives the following RT5 route from PE-3 for prefix 2001:db8::10:66/128:

```
95 2022/01/28 13:30:00.438 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 106
    Flag: 0x90 Type: 14 Len: 69 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.3
        Type: EVPN-IP-PREFIX Len: 58 RD: 192.0.2.3:5, tag: 0,
                  ip_prefix: 2001:db8::10:66/128 gw_ip fe80::7:b0d1:3fa3:2f60
                  Label: 8388512 (Raw Label: 0x7fffa0)
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:5
        bgp-tunnel-encap:MPLS
"
```

In the route table on PE-2, the route for prefix 2001:db8::10:66/128 is an EVPN route:

```
[/]
A:admin@PE-2# show router 6 route-table 2001:db8::10:66

===============================================================================
IPv6 Route Table (Service: 6)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
2001:db8::10:66/128                           Remote  EVPN-IFF  00h01m37s  169
       fe80::7:b0d1:3fa3:2f60-"evi-5"                                 0
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

IPv6 host mobility case 3: host does not send any traffic after moving

The service configuration is the same as use cases 1 and 2. SDP 46 is enabled on PE-4, so host-66 moves back to PE-2. The following traceroute shows that the forwarding path from host-66 to host-61 is via PE-2:

```
[/]
A:admin@PE-6# traceroute router 8 2001:db8::20:61 source 2001:db8::10:66
traceroute to 2001:db8::20:61 from 2001:db8::10:66, 30 hops max, 60 byte packets
  1  2001:db8::10:2 (2001:db8::10:2)    3.76 ms  4.23 ms  4.00 ms
  2  ::  * * *
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
  3  2001:db8::20:61 (2001:db8::20:61)    5.89 ms  5.40 ms  5.61 ms
```

The FDB on PE-2 shows that MAC address 00:00:00:00:06:06 is learned on the local SAP 1/1/1:7, as follows:

```
[/]
A:admin@PE-2# show service id 7 fdb detail

===============================================================================
Forwarding Database, Service 7
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
7          00:00:00:00:06:06 sap:1/1/1:7             L/30     01/28/22 13:36:03
7          00:00:00:00:2f:07 cpm                     Intf     01/28/22 13:17:40
7          00:00:00:00:3f:07 mpls-1:                 EvpnS:P  01/28/22 13:17:47
                             192.0.2.3:524281
           ldp:65538
-------------------------------------------------------------------------------
No. of MAC Entries: 3
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

A failure is simulated, causing a failover from the primary path via PE-4 to the secondary path via PE-5. Host-66 does not send any traffic after switchover. Figure 194: Host does not send any traffic after switchover shows that PE-2 sends an NS message when the FDB entry for host-66 ages out.

Figure 194: Host does not send any traffic after switchover

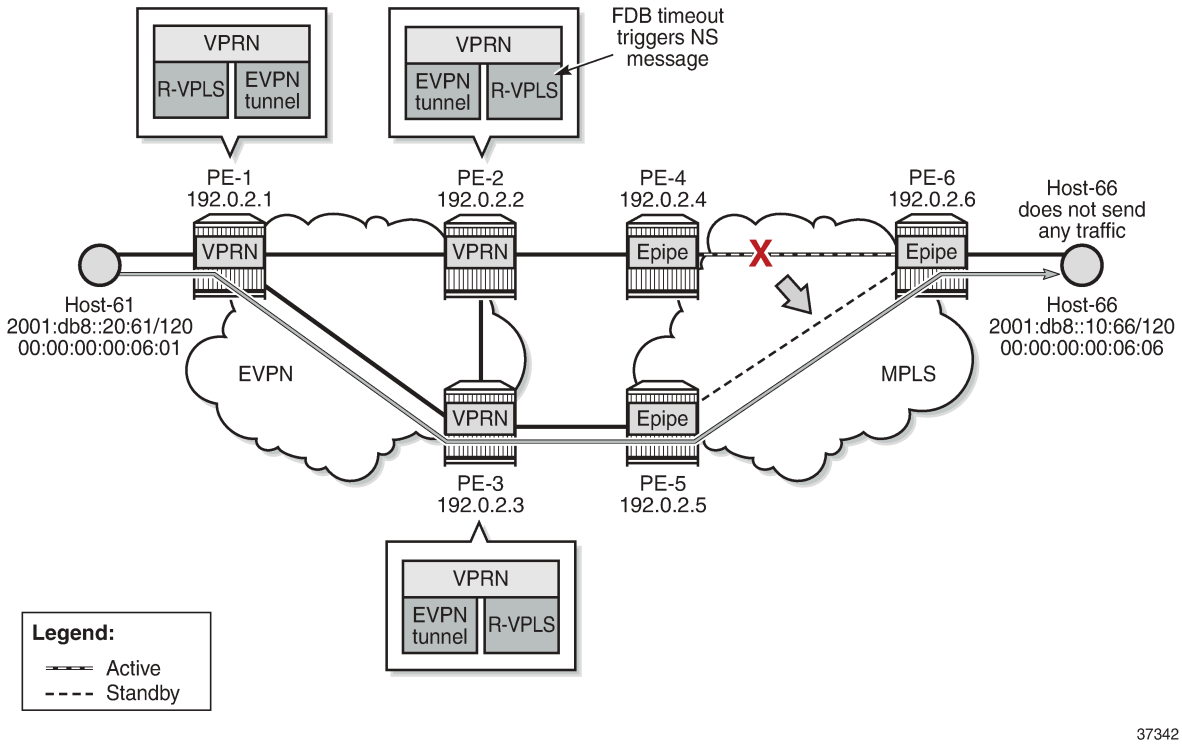

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

On PE-2, MAC address 00:00:00:00:06:06 expires in the FDB for R-VPLS "evi-7", which triggers PE-2 to send an NS message for 2001:db8::10:66. This CPM-generated NS message is flooded to the EVPN destinations PE-1 and PE-3.

```
# on PE-2:
202 2022/01/28 13:37:34.634 UTC MINOR: DEBUG #2001 vprn6 TIP
"TIP: NBR
Sending NS for nbr addr 2001:db8::10:66 nbr type dynamic"
```

The NS message reaches host-66, which replies with an NA message. PE-3 receives the NA message and updates its FDB and ND tables. PE-2 also receives the NA message, but it rejects NA messages received on interface "evi-7" when no learn-dynamic is configured:

```
225 2022/01/28 13:37:34.638 UTC MINOR: DEBUG #2001 vprn6 TIP
"TIP: NBR
Ignore NA for target address 2001:db8::10:66 on evpn endpoint evi-7 because learn-dynamic is
 disabled."
```

PE-2 receives the following EVPN-MAC routes from PE-3:

```
237 2022/01/28 13:43:03.001 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 191
    Flag: 0x90 Type: 14 Len: 146 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.3
        Type: EVPN-MAC Len: 49 RD: 192.0.2.3:7 ESI: ESI-0, tag: 0, mac len: 48
            mac: 00:00:00:00:06:06, IP len: 16, IP: 2001:db8::10:66, label1: 8388512
        Type: EVPN-MAC Len: 49 RD: 192.0.2.3:7 ESI: ESI-0, tag: 0, mac len: 48
            mac: 00:00:00:00:06:06, IP len: 16, IP: fe80::10:6, label1: 8388512
        Type: EVPN-MAC Len: 33 RD: 192.0.2.3:7 ESI: ESI-0, tag: 0, mac len: 48
            mac: 00:00:00:00:06:06, IP len: 0, IP: NULL, label1: 8388512
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:7
        bgp-tunnel-encap:MPLS
        mac-mobility:Seq:3
"
```

PE-2 receives the following RT5 route from PE-3:

```
228 2022/01/28 13:37:35.638 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 106
    Flag: 0x90 Type: 14 Len: 69 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.3
        Type: EVPN-IP-PREFIX Len: 58 RD: 192.0.2.3:5, tag: 0,
                  ip_prefix: 2001:db8::10:66/128 gw_ip fe80::7:b0d1:3fa3:2f60
                  Label: 8388512 (Raw Label: 0x7fffa0)
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        target:64500:5
        bgp-tunnel-encap:MPLS
"
```

Upon receiving the routes, PE-2 updates its FDB and ARP tables. After the switchover, MAC address 00:00:00:00:06:06 is no longer learned on a local SAP on PE-2, but is learned via an EVPN-MAC route from PE-3, as follows:

```
[/]
A:admin@PE-2# show service id 7 fdb detail

===============================================================================
Forwarding Database, Service 7
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
7          00:00:00:00:06:06 mpls-1:                 Evpn     01/28/22 13:37:35
                             192.0.2.3:524281
           ldp:65538
7          00:00:00:00:2f:07 cpm                     Intf     01/28/22 13:17:40
7          00:00:00:00:3f:07 mpls-1:                 EvpnS:P  01/28/22 13:17:47
                             192.0.2.3:524281
           ldp:65538
-------------------------------------------------------------------------------
No. of MAC Entries: 3
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

Conclusion

EVPN host mobility is supported in SR OS as described in draft-ietf-bess-evpn-inter-subnet-forwarding. This chapter describes several cases when a host moves from a source PE to a target PE within the same broadcast domain.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Operational Groups for EVPN-VXLAN VPWS Services

This chapter describes the Operational Groups for EVPN-VXLAN VPWS Services.

Topics in this chapter include:

Applicability

This chapter was initially written based on SR OS Release 16.0.R5, but the MD-CLI in the current edition corresponds to SR OS Release 21.7.R1. EVPN-VXLAN VPWS and service-level operational groups for VPWS services are supported in SR OS Release 16.0.R1, or later.

Overview

Operational groups on Epipe services are used for fault propagation to other services, such as I-VPLS or R-VPLS services. Epipes with VXLAN destinations are used in some edge PE applications along with port cross-connect (PXC) so that VXLAN networks can be terminated in other VPLS or VPRN services. In such cases, the operational status of the Epipe services terminating VXLAN must override the operational status of the SAPs of the VPLS or VPRN where the Epipe is stitched to.

Operational group on egress VTEP in Epipes with static VXLAN bindings

The Static VXLAN Termination in Epipe Services chapter describes how Epipes with static VXLAN termination are stitched to I-VPLS services. In Epipes with static VXLAN bindings, operational groups can be configured in the egress VTEP context. Figure 195: Epipe with static VXLAN termination shows the example topology with a static VXLAN tunnel between PE-1 and an anycast address on PE-2 and PE-3. The All-Active Multi-Homing Ethernet Segments (AA MH ESs) "vES23_1.101" and "vES23_1.102" are used by the I-VPLSs 101 and 102, which are both stitched to Epipe 1 in PE-2 and PE-3. The SAPs in these I-VPLSs monitor the operational group configured in the egress VTEP context of the Epipe service, so the SAPs will go operationally down when the operational group of the VTEP goes operationally down.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 195: Epipe with static VXLAN termination

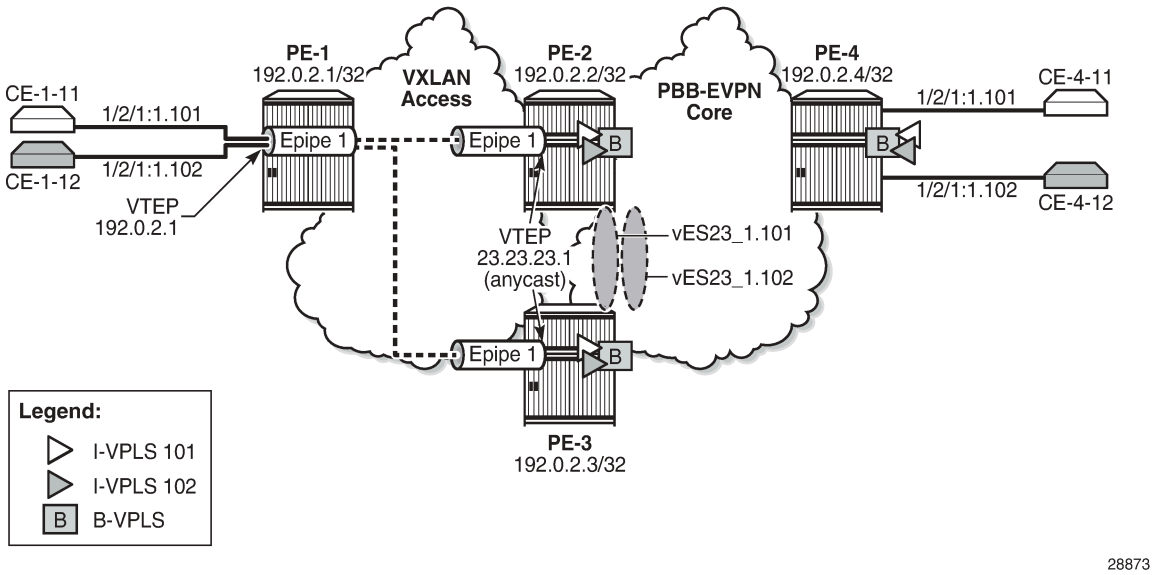

On PE-2 and PE-3, Epipe 1 is configured with static VXLAN bindings, as follows. The egress VTEP is 192.0.2.1, which is the system IP address of PE-1. Operational group "op-grp-1" in configured for this egress VTEP. LAG 2 combines PXC ports and is used to stitch Epipe 1 to the I-VPLS services 101 and 102. For a detailed description of the configuration, see the Static VXLAN Termination in Epipe Services chapter.

```
# on PE-2, PE-3:
configure {
    service {
        oper-group "op-grp-1" {
        }
        epipe "Epipe 1" {
            admin-state enable
            description "Epipe 1 with static VXLAN bindings"
            service-id 1
            customer "1"
            sap lag-2:1.* {
            }
            vxlan {
                source-vtep 23.23.23.1
                instance 1 {
                    vni 1
                    egress-vtep {
                        ip-address 192.0.2.1
                        oper-group "op-grp-1"
                    }
                }
            }
        }
```

For failure propagation to the stitched I-VPLSs, the SAPs in the I-VPLSs can monitor the operational group "op-grp-1", for I-VPLS 101 on PE-2 and PE-3, as follows:

```
# on PE-2, PE-3:
configure {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    service {
        vpls "I-VPLS 101" {
            admin-state enable
            service-id 101
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS-100" {
                    isid 101
                }
            }
            sap lag-1:1.101 {
                monitor-oper-group "op-grp-1"
            }
        }
```

When the egress VTEP prefix 192.0.2.1 disappears from the global route-table on PE-2, the VXLAN binding goes down, as follows:

```
[/]
A:admin@PE-2# show service id 1 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address                          Egress VNI          Oper      Vxlan
                                                          State     Type
-------------------------------------------------------------------------------
192.0.2.1                             1                   Down      static
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 1
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN VXLAN ES Dest
===============================================================================
I Eth Seg Id                    TEP Address       VNI       Last Changed
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

When the egress VTEP 192.0.2.1 goes down, the operational group "op-grp-1" goes down too, as follows:

```
[/]
A:admin@PE-2# show service oper-group "op-grp-1"

===============================================================================
Service Oper Group Information
===============================================================================
Oper Group       : op-grp-1
Creation Origin  : manual                          Oper Status: down
Hold DownTime    : 0 secs                          Hold UpTime: 4 secs
Members          : 1                               Monitoring : 2
===============================================================================
```

When the operational group "op-grp-1" goes down, the monitoring SAP in I-VPLS 101 goes operationally down with flag OperGroupDown, as follows:

```
[/]
A:admin@PE-2# show service id 101 sap lag-1:1.101 detail | match 'Flags | Oper State'
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Admin State        : Up                        Oper State        : Down
Flags              : OperGroupDown
Stp Admin State    : Up                        Stp Oper State    : Down
```

When this SAP goes down, the entire I-VPLS 101 service goes down on PE-2, as follows:

```
[/]
A:admin@PE-2# show service id 101 base | match 'State'
Admin State       : Up                  Oper State        : Down
```

Epipes with static VXLAN bindings impose the following restrictions, which cannot be overcome unless a control plane protocol such as BGP-EVPN is used for the VXLAN bindings.

- When anycast VTEPs on the PEs are used, a change in the vES preference on the DF PE triggers a DF switchover for the I-VPLS service. However, the access PE (PE-1 in Figure 1) is unaware and keeps sending the VXLAN traffic to the same PE, unless a change in DF comes with an automatic change in the underlay IGP metrics, which cannot be easily accomplished.
- Without a control plane, Eth-CFM must be used between PEs and access PEs to detect end-to-end service-level failures.
- Traffic from the access PE is forwarded to the anycast VTEP, based on underlay IGP metrics. There is no control on a per-service basis.
- The architecture does not support AA MH for the Epipe service, so the access PEs always send the traffic to one single PE.

The preceding challenges can be addressed by using different VTEPs on the PEs and adding a BGP-EVPN control plane on the Epipe, as described in the next section.

Operational groups in EVPN-VXLAN Epipes

Figure 196: Epipe 2 with EVPN-VXLAN and all-active multi-homing shows EVPN-VXLAN Epipe 2 stitched to I-VPLSs 201 and 202. AA MH ESs "vES23_2.201" and "vES23_2.202" are used by the I-VPLSs 201 and 202 respectively.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 196: Epipe 2 with EVPN-VXLAN and all-active multi-homing

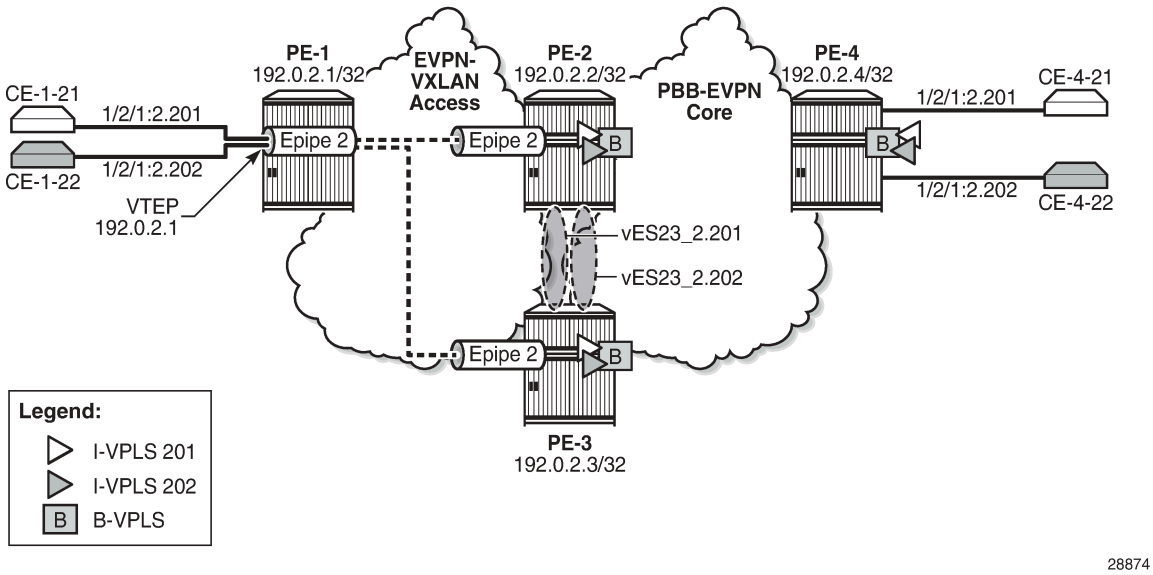

The EVPN-VXLAN VPWS chapter describes the configuration of Epipes with EVPN-VXLAN bindings instead of static VXLAN bindings. The egress VTEP is not configured manually, but dynamically learned through BGP-EVPN. Therefore, the operational group cannot be configured in the egress VTEP context. However, it is possible to configure an operational group in an Epipe at the service level, as follows:

```
# on PE-2:
configure {
    service {
        oper-group "op-grp-2" {
        }
        epipe "Epipe 2" {
            admin-state enable
            description "Epipe 2 with EVPN-VXLAN"
            service-id 2
            customer "1"
            oper-group "op-grp-2"
            bgp 1 {
            }
            sap lag-2:2.* {
            }
            vxlan {
                instance 1 {
                    vni 2
                }
            }
            bgp-evpn {
                evi 2
                local-attachment-circuit "AC-23" {
                    eth-tag 123
                }
                remote-attachment-circuit "AC-1" {
                    eth-tag 101
                }
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
            }
        }
```

The following shows the error messages raised when attempting to configure the egress VTEP manually in an Epipe service with BGP-EVPN enabled:

```
[ex:/configure service epipe "Epipe 2" vxlan instance 1 egress-vtep]
A:admin@PE-2# ip-address 192.0.2.1

*[ex:/configure service epipe "Epipe 2" vxlan instance 1 egress-vtep]
A:admin@PE-2# commit
MINOR: SVCMGR #12: configure service epipe "Epipe 2" vxlan instance 1 egress-vtep ip-address -
 Inconsistent Value error - configuration incompatible with bgp-evpn - configure service epipe
 "Epipe 2" bgp-evpn
MINOR: MGMT_CORE #4001: configure service epipe "Epipe 2" vxlan instance 1 egress-vtep ip-
address - egress vtep not supported with bgp-evpn - configure service epipe "Epipe 2" bgp-evpn
```

An operational group can be associated with the entire Epipe or with specific objects, such as SAPs or spoke-SDPs, but not simultaneously. The following error is raised when attempting to associate the operational group "op-grp-2"—that is already associated with the Epipe with the SAP on PE-2:

```
[ex:/configure service epipe "Epipe 2" sap lag-2:2.*]
A:admin@PE-2# oper-group "op-grp-2"

*[ex:/configure service epipe "Epipe 2" sap lag-2:2.*]
A:admin@PE-2# commit
MINOR: SVCMGR #12: configure service epipe "Epipe 2" sap lag-2:2.* oper-group - Inconsistent
 Value error - cannot monitor or belong to a service oper-group within the same service -
 configure service epipe "Epipe 2" oper-group
```

The service-level operational group status is derived from the service operational status: when Epipe 2 is operationally down, the operational group "op-grp-2" will be down.

For fault propagation to the stitched I-VPLSs 201 and 202, the SAPs in the I-VPLSs monitor the operational group "op-grp-2", for I-VPLS 201 on PE-2 and PE-3, as follows:

```
# on PE-2, PE-3:
configure {
    service {
        vpls "I-VPLS 201" {
            admin-state enable
            service-id 201
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS-100" {
                    isid 201
                }
            }
            sap lag-1:2.201 {
                 monitor-oper-group "op-grp-2"
            }
        }
```

Configuration

In this section, the following use cases are described:

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- operational group on egress VTEP in Epipes with static VXLAN bindings stitched to I-VPLSs using AA MH ESs
- service-level operational group in EVPN-VXLAN Epipes stitched to I-VPLSs using AA MH ESs
- service-level operational group in EVPN-VXLAN Epipes stitched to I-VPLSs using Single-Active (SA) MH ESs

Figure 197: Example topology shows the example topology with four PEs in an autonomous system.

Figure 197: Example topology

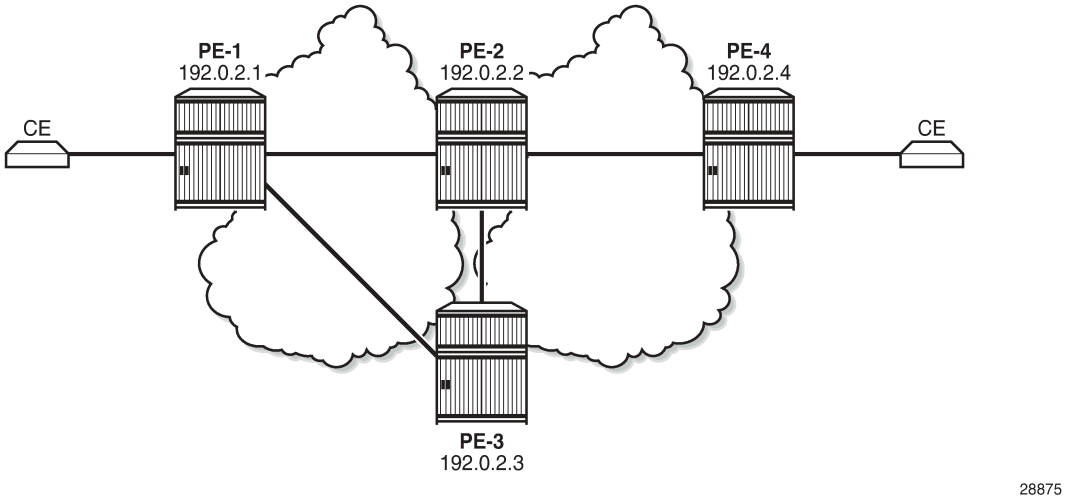

The initial configuration includes:

- Cards, MDAs, ports
- Router interfaces
- IS-IS as IGP (level capability 2 in the core between PE-2, PE-3, and PE-4; level capability 1 in the access toward PE-1)
- LDP between the core routers PE-2, PE-3, and PE-4

Oper group on egress VTEP in Epipes with static VXLAN bindings stitched to I-VPLSs using AA MH ESs

When static VXLAN bindings are used, no BGP-EVPN is required in the access network to and from PE-1; BGP is only configured in the core network. When PE-2 acts as the route reflector, its BGP configuration is as follows:

```
# on RR PE-2:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-update {
                evpn true
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            group "CORE" {
                type internal
                split-horizon true
                family {
                    evpn true
                }
                cluster {
                    cluster-id 192.0.2.2
                }
            }
            neighbor "192.0.2.3" {
                group "CORE"
            }
            neighbor "192.0.2.4" {
                group "CORE"
            }
```

Figure 195: Epipe with static VXLAN termination shows that Epipe 1 is configured in the access network: on PE-1, the VTEP is the system address 192.0.2.1 (default), and on PE-2 and PE-3, the VTEP is a unicast address 23.23.23.1.

On PE-1, Epipe 1 is configured with egress VTEP 23.23.23.1, as follows:

```
# on PE-1:
configure {
    service {
        epipe "Epipe 1" {
            admin-state enable
            description "Epipe 1 with static VXLAN bindings"
            service-id 1
            customer "1"
            sap 1/2/1:1.* {
            }
            vxlan {
                instance 1 {
                    vni 1
                    egress-vtep {
                        ip-address 23.23.23.1
                    }
                }
            }
        }
```

On PE-2 and PE-3, the following unicast address is configured:

```
# on PE-2, PE-3:
configure {
    router "Base" {
        interface "lo23" {
            loopback
            ipv4 {
                primary {
                    address 23.23.23.0
                    prefix-length 31
                }
            }
```

On PE-2 and PE-3, three ports are configured as PXC. PXC 1 is used as Forwarding Path Extension (FPE) and the VXLAN tunnel termination 23.23.23.1 is configured with this FPE, as follows:

```
# on PE-2, PE-3:
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
configure {
    service {
        system {
            vxlan {
                tunnel-termination 23.23.23.1 {
                    fpe-id 1
                }
            }
```

PXCs 2 and 3 are used in the internal LAGs that are used to stitch Epipe 1 to I-VPLSs 101 and 102, as follows. LAG "lag-1" will be used in the I-VPLS services; LAG "lag-2" in the Epipe services.

```
# on PE-2, PE-3:
configure {
    lag "lag-1" {
        admin-state enable
        encap-type qinq
        mode hybrid
        max-ports 64
        port pxc-2.a {
        }
        port pxc-3.a {
        }
    }
    lag "lag-2" {
        admin-state enable
        encap-type qinq
        mode hybrid
        max-ports 64
        port pxc-2.b {
        }
        port pxc-3.b {
        }
    }
```

On PE-2 and PE-3, Epipe 1 is configured with source VTEP 23.23.23.1 and egress VTEP 192.0.2.1, as follows. The operational group "op-grp-1" is associated with the egress VTEP. The SAP stitches Epipe 1 to the I-VPLSs 101 and 102.

```
# on PE-2, PE-3:
configure {
    service {
        oper-group "op-grp-1" {
        }
        epipe "Epipe 1" {
            admin-state enable
            description "Epipe 1 with static VXLAN bindings"
            service-id 1
            customer "1"
            sap lag-2:1.* {
            }
            vxlan {
                source-vtep 23.23.23.1
                instance 1 {
                    vni 1
                    egress-vtep {
                        ip-address 192.0.2.1
                        oper-group "op-grp-1"
                    }
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
```

On PE-2, B-VPLS 100 is configured as follows. The configuration is similar on PE-3 and PE-4.

```
# on PE-2:
configure {
    service {
        vpls "B-VPLS-100" {
            admin-state enable
            service-id 100
            customer "1"
            service-mtu 1532
            pbb-type b-vpls
            pbb {
                source-bmac {
                    address 00:00:00:00:00:02
                    use-es-bmac-lsb true
                }
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 100
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
        }
```

On PE-2, I-VPLS 101 is configured as follows. The SAP monitors the operational group "op-grp-1" that is configured in Epipe 1. The AA MH ES "vES23_1.101" is used. The configuration of I-VPLS 102 is similar, but it uses AA MH ES "vES23_1.102" with preference value 50 instead. On PE-3, the preference values are reversed: preference value 50 for "vES23_1.101" and 100 for "vES23_1.102".

```
# on PE-2:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "vES23_1.101" {
                        admin-state enable
                        type virtual
                        esi 0x01000000002300000111
                        multi-homing-mode all-active
                        df-election {
                            es-activation-timer 3
                            service-carving-mode manual
                            manual {
                                preference {
                                    value 100
                                }
                            }
                        }
                        association {
                            lag "lag-1" {
                                virtual-ranges {
                                    qinq {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                                    s-tag-c-tag 1 c-tag-start 101 {
                                        c-tag-end 101
                                    }
                                }
                            }
                        }
                    }
                    pbb {
                        source-bmac-lsb 0x2311
                    }
                }
            }
        }
        vpls "I-VPLS 101" {
            admin-state enable
            service-id 101
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS-100" {
                    isid 101
                }
            }
            sap lag-1:1.101 {
                monitor-oper-group "op-grp-1"
            }
        }
```

On PE-4, I-VPLS 101 is configured as follows:

```
# on PE-4:
configure {
    service {
        vpls "I-VPLS 101" {
            admin-state enable
            service-id 101
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS-100" {
                    isid 101
                }
            }
            sap 1/2/1:1.101 {
            }
        }
```

To emulate a failure that affects the operational state of the egress VTEP (and also of the Epipe service), the SAP of Epipe 1 on PE-2 is disabled, as follows:

```
# on PE-2:
configure {
    service {
        epipe "Epipe 1" {
            sap lag-2:1.* {
                admin-state disable
```

When the SAP is operationally down, Epipe 1 goes down, as follows:

```
[/]
A:admin@PE-2# show service id 1 base | match 'State'
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Admin State       : Up                  Oper State        : Down
```

The egress VTEP 192.0.2.1 is operationally down, as follows:

```
[/]
A:admin@PE-2# show service id 1 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address                          Egress VNI          Oper      Vxlan
                                                          State     Type
-------------------------------------------------------------------------------
192.0.2.1                             1                   Down      static
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 1
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN VXLAN ES Dest
===============================================================================
I Eth Seg Id                    TEP Address       VNI       Last Changed
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

The operational group "op-grp-1" is associated with the egress VTEP, so it goes operationally down, as follows:

```
[/]
A:admin@PE-2# show service oper-group "op-grp-1"

===============================================================================
Service Oper Group Information
===============================================================================
Oper Group       : op-grp-1
Creation Origin  : manual                          Oper Status: down
Hold DownTime    : 0 secs                          Hold UpTime: 4 secs
Members          : 1                               Monitoring : 2
===============================================================================
```

This operational group is monitored by the SAPs in I-VPLSs 101 and 102, so these SAPs go down with flag OperGroupDown; for example, for I-VPLS 101 on PE-2:

```
[/]
A:admin@PE-2# show service id 101 sap lag-1:1.101 detail | match 'Flags'
Flags              : OperGroupDown
```

When the SAP goes down, the I-VPLS service goes down, as follows:

```
[/]
A:admin@PE-2# show service id 101 base | match 'State'
Admin State       : Up                  Oper State        : Down
```

Even though Epipe 1 on PE-2 is operationally down while Epipe 1 on PE-3 is up, PE-1 is unaware because the VXLAN destination in Epipe 1 remains up, as follows:

```
[/]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
A:admin@PE-1# show service id 1 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address                            Egress VNI          Oper      Vxlan
                                                            State     Type
-------------------------------------------------------------------------------
23.23.23.1                              1                   Up        static
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 1
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN VXLAN ES Dest
===============================================================================
I Eth Seg Id                    TEP Address       VNI       Last Changed
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

With ECMP=1, all traffic from PE-1 is directed to PE-2, regardless of the state of the Epipe on PE-2. The following route table on PE-1 shows that destination prefix 23.23.23.0/31 has next-hop 192.168.12.2, which is an interface address on PE-2.

```
[/]
A:admin@PE-1# show router route-table 23.23.23.0/31

===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
23.23.23.0/31                                   Remote  ISIS      00h03m57s  15
       192.168.12.2                                                    10
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

Traffic from the CEs attached to PE-1 is forwarded by PE-1 to PE-2, where it is dropped.

Service-level operational group in EVPN-VXLAN Epipes stitched to I-VPLS using AA MH ESs

BGP must be enabled on all nodes for the EVPN address family, also in the access to and from PE-1. The BGP configuration on PE-1 is as follows:

```
# on PE-1:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-import true
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            vpn-apply-export true
            rapid-update {
                evpn true
            }
            group "ACCESS" {
                type internal
                split-horizon true
                family {
                    evpn true
                }
            }
            neighbor 192.0.2.2 {
                group "ACCESS"
            }
            neighbor 192.0.2.3 {
                group "ACCESS"
            }
```

On PE-1, the following EVPN-VXLAN Epipe 2 is configured with local Ethernet tag 101 and remote Ethernet tag 123.

```
# on PE-1:
configure {
    service {
        epipe "Epipe 2" {
            admin-state enable
            description "Epipe 2 with EVPN-VXLAN"
            service-id 2
            customer "1"
            bgp 1 {
            }
            sap 1/2/1:2.* {
            }
            vxlan {
                instance 1 {
                    vni 2
                }
            }
            bgp-evpn {
                evi 2
                local-attachment-circuit "AC-1" {
                    eth-tag 101
                }
                remote-attachment-circuit "AC-23" {
                    eth-tag 123
                }
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
        }
```

On PE-2, the following EVPN-VXLAN Epipe 2 is configured with local Ethernet tag 123 and remote Ethernet tag 101. The operational group "op-grp-2" is associated with Epipe 2. The configuration on PE-3 is identical.

```
# on PE-2:
configure {
    service {
        oper-group "op-grp-2" {
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        epipe "Epipe 2" {
            admin-state enable
            description "Epipe 2 with EVPN-VXLAN"
            service-id 2
            customer "1"
            oper-group "op-grp-2"
            bgp 1 {
            }
            sap lag-2:2.* {
            }
            vxlan {
                instance 1 {
                    vni 2
                }
            }
            bgp-evpn {
                evi 2
                local-attachment-circuit "AC-23" {
                    eth-tag 123
                }
                remote-attachment-circuit "AC-1" {
                    eth-tag 101
                }
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
        }
```

The configuration of B-VPLS 100 remains unchanged and the configuration of the I-VPLSs 201 and 202 resembles the configuration of VPLSs 101.

When there is no failure, the egress VTEP for Epipe 2 on PE-1 is 192.0.2.2, which is the system IP address of PE-2, as follows:

```
[/]
A:admin@PE-1# show service id 2 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address                          Egress VNI          Oper      Vxlan
                                                          State     Type
-------------------------------------------------------------------------------
192.0.2.2                             2                   Up        evpn
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 1
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN VXLAN ES Dest
===============================================================================
I Eth Seg Id                    TEP Address       VNI       Last Changed
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

To emulate a failure that affects the operational state of the Epipe service, the SAP in Epipe 2 is disabled, as follows:

```
# on PE-2:
configure {
    service {
        epipe "Epipe 2" {
            sap lag-2:2.* {
                admin-state disable
```

When the SAP goes down, the Epipe goes down, as follows:

```
[/]
A:admin@PE-2# show service id 2 base | match 'State'
Admin State       : Up                 Oper State        : Down
```

On PE-1, the egress VTEP for Epipe 2 is 192.0.2.3, which is the system IP address of PE-3, as follows:

```
[/]
A:admin@PE-1# show service id 2 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address                          Egress VNI          Oper      Vxlan
                                                          State     Type
-------------------------------------------------------------------------------
192.0.2.3                             2                   Up        evpn
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 1
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN VXLAN ES Dest
===============================================================================
I Eth Seg Id                    TEP Address       VNI       Last Changed
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

The operational group "op-grp-2" follows the state of Epipe 2, so it goes down, as follows. As a consequence, the monitoring SAPs for this operational group also go down.

```
[/]
A:admin@PE-2# show service oper-group "op-grp-2" detail

===============================================================================
Service Oper Group Information
===============================================================================
Oper Group       : op-grp-2
Creation Origin  : manual                        Oper Status: down
Hold DownTime    : 0 secs                        Hold UpTime: 4 secs
Members          : 1                             Monitoring : 2
===============================================================================

===============================================================================
Member Services for OperGroup: op-grp-2
===============================================================================
Svc Id
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
2
-------------------------------------------------------------------------------
Service Entries found: 1
===============================================================================

===============================================================================
Monitoring SAPs for OperGroup: op-grp-2
===============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.  Egr.   Adm  Opr
                                           QoS   Fltr    QoS   Fltr
-------------------------------------------------------------------------------
lag-1:2.201                     201        1     none    1     none   Up   Down
lag-1:2.202                     202        1     none    1     none   Up   Down
-------------------------------------------------------------------------------
SAP Entries found: 2
===============================================================================
```

The SAPs in I-VPLSs 201 and 202 go down with the OperGroupDown flag, as follows:

```
[/]
A:admin@PE-2# show service id 201 sap lag-1:2.201 detail | match 'Flags'
Flags              : OperGroupDown

[/]
A:admin@PE-2# show service id 202 sap lag-1:2.202 detail | match 'Flags'
Flags              : OperGroupDown
```

When the SAPs go down, the I-VPLSs 201 and 202 also go down, as follows:

```
[/]
A:admin@PE-2# show service id 201 base | match 'State'
Admin State       : Up                 Oper State        : Down

[/]
A:admin@PE-2# show service id 202 base | match 'State'
Admin State       : Up                 Oper State        : Down
```

Even with this failure on PE-2, traffic can still flow between the CEs, as follows:

```
[/]
A:admin@PE-4# ping 172.16.21.11 router-instance "VPRN 21" interval 0.1 output-format summary
PING 172.16.21.11 56 data bytes
!!!!!
---- 172.16.21.11 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 3.40ms, avg = 4.02ms, max = 4.57ms, stddev = 0.388ms
```

The following FDB for I-VPLS 201 on PE-4 shows that MAC address 00:ca:fe:00:21:11 of CE-1-21 is reachable via AA MH ES with ES-BMAC 00:00:00:00:23:21:

```
[/]
A:admin@PE-4# show service id 201 fdb detail

===============================================================================
Forwarding Database, Service 201
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
201        00:ca:fe:00:21:11 eES-BMAC:               L/90     08/10/21 16:56:31
                             00:00:00:00:23:21
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
201        00:ca:fe:00:21:41 sap:1/2/1:2.201         L/90     08/10/21 16:56:24
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

Service-level operational group in EVPN-VXLAN Epipes stitched to I-VPLS using SA MH ESs

Figure 198: Epipe 3 with EVPN-VXLAN and SA MH ES shows the example topology with an SA MH ES used by the I-VPLSs.

Figure 198: Epipe 3 with EVPN-VXLAN and SA MH ES

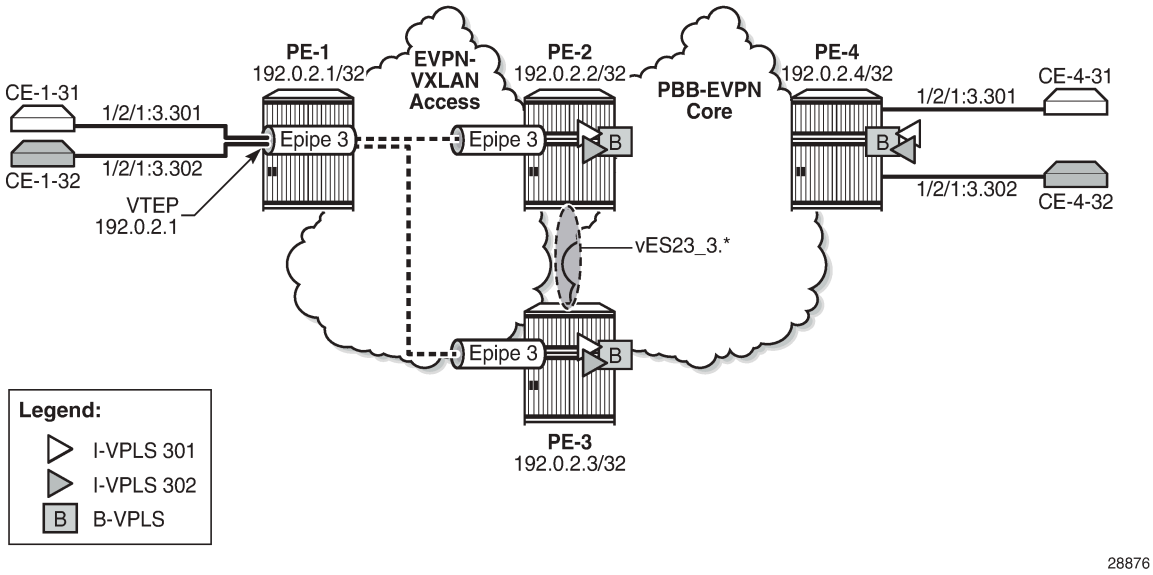

The configuration of Epipe 3 resembles the configuration of Epipe 2: the same Ethernet tags are used, only the VNI, EVI, and SAPs are different.

On PE-2 and PE-3, PXC 4 is configured to stitch Epipe 3 to I-VPLSs 301 and 302. The PXC port will be used in the SA MH ES.

On PE-2 and PE-3, Epipe 3 is configured as follows:

```
# on PE-2, PE-3:
configure {
    service {
        oper-group "op-grp-3" {
        }
        epipe "Epipe 3" {
            admin-state enable
            description "EVPN-VXLAN Epipe 3"
            service-id 3
            customer "1"
            oper-group "op-grp-3"
            sap pxc-4.b:3.* {
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            vxlan {
                instance 1 {
                    vni 3
                }
            }
            bgp-evpn {
                evi 3
                local-attachment-circuit "AC-23" {
                    eth-tag 123
                }
                remote-attachment-circuit "AC-1" {
                    eth-tag 101
                }
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
        }
```

On PE-2, I-VPLS 301 uses SA MH ES "vES23_3.*", and is configured as follows.

```
# on PE-2:
configure {
    service {
        oper-group "op-grp-3" {
        }
        system {
            bgp {
                evpn {
                    ethernet-segment "vES23_3.*" {
                        admin-state enable
                        type virtual
                        esi 0x01000000002300000301
                        multi-homing-mode single-active
                        df-election {
                            es-activation-timer 3
                            service-carving-mode manual
                            manual {
                                preference {
                                    value 100
                                }
                            }
                        }
                        association {
                            port pxc-4.a {
                                virtual-ranges {
                                    qinq {
                                        s-tag 3 {
                                            end 3
                                        }
                                    }
                                }
                            }
                        }
                        pbb {
                            source-bmac-lsb 0x2334
                        }
                    }
                }
            }
        }
        vpls "I-VPLS 301" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            admin-state enable
            service-id 301
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS-100" {
                    isid 301
                }
            }
            sap pxc-4.a:3.301 {
                monitor-oper-group "op-grp-3"
            }
        }
```

The configuration is similar on PE-3, but with source-bmac-lsb 23-35 and preference 50.

When Epipe 3 on PE-2 is operationally up, the egress VTEP for Epipe 3 on PE-1 is 192.0.2.2, as follows:

```
[/]
A:admin@PE-1# show service id 3 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address                            Egress VNI          Oper      Vxlan
                                                            State     Type
-------------------------------------------------------------------------------
192.0.2.2                               3                   Up        evpn
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 1
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN VXLAN ES Dest
===============================================================================
I Eth Seg Id                    TEP Address       VNI       Last Changed
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

To emulate a failure on PE-2 that affects the operational state of the Epipe service, the SAP in Epipe 3 is disabled, as follows:

```
# on PE-2:
configure {
    service {
        epipe "Epipe 3" {
            sap pxc-4.b:3.* {
                admin-state disable
```

When the SAP goes down, Epipe 3 goes down on PE-2, as follows:

```
[/]
A:admin@PE-2# show service id 3 base | match 'State'
Admin State       : Up                  Oper State        : Down
```

When Epipe 3 on PE-2 goes operationally down, the egress VTEP for Epipe 3 on PE-1 is 192.0.2.3, as follows:

```
[/]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
A:admin@PE-1# show service id 3 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address                          Egress VNI          Oper      Vxlan
                                                          State     Type
-------------------------------------------------------------------------------
192.0.2.3                             3                   Up        evpn
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 1
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN VXLAN ES Dest
===============================================================================
I Eth Seg Id                    TEP Address       VNI       Last Changed
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

The operational group "op-grp-3" follows the state of Epipe 3 on PE-2, so it goes down. Also, the monitoring SAPs for this operational group go down.

```
[/]
A:admin@PE-2# show service oper-group "op-grp-3" detail

===============================================================================
Service Oper Group Information
===============================================================================
Oper Group       : op-grp-3
Creation Origin  : manual                          Oper Status: down
Hold DownTime    : 0 secs                          Hold UpTime: 4 secs
Members          : 1                               Monitoring : 2
===============================================================================

===============================================================================
Member Services for OperGroup: op-grp-3
===============================================================================
Svc Id
-------------------------------------------------------------------------------
3
-------------------------------------------------------------------------------
Service Entries found: 1
===============================================================================

===============================================================================
Monitoring SAPs for OperGroup: op-grp-3
===============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.  Egr.   Adm  Opr
                                           QoS   Fltr    QoS   Fltr
-------------------------------------------------------------------------------
pxc-4.a:3.301                   301        1     none    1     none   Up   Down
pxc-4.a:3.302                   302        1     none    1     none   Up   Down
-------------------------------------------------------------------------------
SAP Entries found: 2
===============================================================================
```

The SAPs in I-VPLSs 301 and 302 on PE-2 go down with the OperGroupDown flag, as follows:

```
[/]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
A:admin@PE-2# show service id 301 sap pxc-4.a:3.301 detail | match 'Flags' post-lines 1
Flags              : StandByForMHProtocol
                     OperGroupDown

[/]
A:admin@PE-2# show service id 302 sap pxc-4.a:3.302 detail | match 'Flags' post-lines 1
Flags              : StandByForMHProtocol
                     OperGroupDown
```

When the SAPs go down, the I-VPLSs go down on PE-2, as follows:

```
[/]
A:admin@PE-2# show service id 301 base | match 'State'
Admin State       : Up                 Oper State        : Down

[/]
A:admin@PE-2# show service id 302 base | match 'State'
Admin State       : Up                 Oper State        : Down
```

When the initial DF PE-2 goes down for the I-VPLSs 301 and 302, PE-3 becomes the new DF. The connectivity between the CEs is preserved, as follows:

```
[/]
A:admin@PE-4# ping 172.16.31.11 router-instance "VPRN 31" interval 0.1 output-format summary
PING 172.16.31.11 56 data bytes
!!!!!
---- 172.16.31.11 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 3.56ms, avg = 3.89ms, max = 4.94ms, stddev = 0.526ms
```

The following FDB for I-VPLS 301 on PE-4 shows that the frames toward MAC address 00:ca:fe:00:31:11 of CE-1-31 are sent via PE-3 (192.0.2.3):

```
[/]
A:admin@PE-4# show service id 301 fdb detail

===============================================================================
Forwarding Database, Service 301
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
301        00:ca:fe:00:31:11 b-mpls:                 L/270    08/10/21 17:08:17
                             192.0.2.3:524281
           ldp:65539
301        00:ca:fe:00:31:41 sap:1/2/1:3.301         L/270    08/10/21 17:04:55
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

PE-3 is now the DF for I-VPLS 301, as follows:

```
[/]
A:admin@PE-3# show service id 301 ethernet-segment

===============================================================================
SAP Ethernet-Segment Information
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
SAP                   Eth-Seg                         Status
-------------------------------------------------------------------------------
pxc-4.a:3.301         vES23_3.*                       DF
===============================================================================
No sdp entries
No vxlan instance entries
```

Conclusion

Some service providers use VXLAN as a next-generation access technology used between the MSANs (or access PEs) and core PE routers. EVPN-VXLAN Epipes can be stitched using PXC to other services, such as I-VPLS. Operational groups can be defined in the Epipe for fault propagation to the SAPs of the services where the Epipe is stitched to.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Operational Groups in EVPN Services

This chapter provides information about Operational Groups in EVPN Services.

Topics in this chapter include:

Applicability

The information and configuration in this chapter are based on SR OS Release 21.10.R1. EVPN operational groups are supported in EVPN-VXLAN and EVPN-MPLS VPLS and R-VPLS services in SR OS Release 19.10.R2 and later; in EVPN-MPLS Epipes in SR OS Release 19.5.R1 and later.

Overview

An operational group includes objects and drives the status of service endpoints (such as pseudowires, SAPs, IP interfaces) located in the same or in different service instances. The operational group status is derived from the status of the individual components. Other service objects can monitor the operational group status. The status of the operational group influences the status of the monitoring objects.

If the operational group goes down, the monitoring objects are also brought operationally down. When one of the objects included in the operational group comes up, the entire operational group comes up, as well as the monitoring objects.

Operational groups for EVPN destinations

Figure 199: EVPN mesh going down triggers DF switchover from PE-5 to PE-4 shows a sample topology with VPLS 1 configured on all nodes. PE-4 and PE-5 share a single-active Ethernet Segment (ES) "ESI-45_1" where PE-5 is the Designated Forwarder (DF).

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 199: EVPN mesh going down triggers DF switchover from PE-5 to PE-4

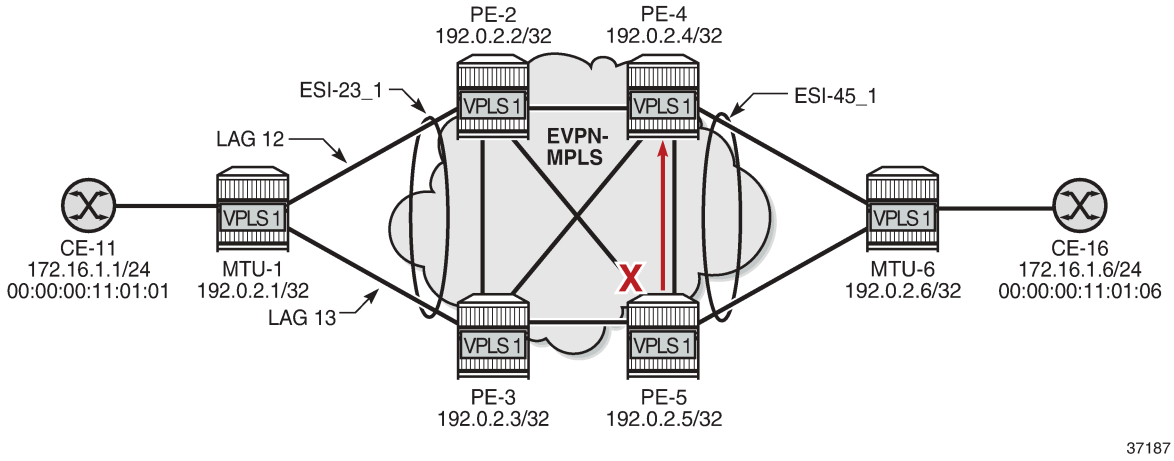

When the EVPN-VPLS service becomes isolated from the rest of the EVPN network (for example, all EVPN destinations are removed on DF PE-5), an operational group for EVPN destinations is required to trigger a DF switchover and bring the monitoring access SAP (or spoke SDP) down. EVPN single-active multi-homing PEs that are elected as NDF must notify their attached access nodes to prevent these from sending traffic to the NDF. Ethernet Connectivity Fault Management (ETH-CFM) is enabled on a down Maintenance Endpoint (MEP) configured on the SAP to detect SAP failure. After the remote MEP on MTU-6 detects the failure, MTU-6 redirects its traffic to PE-4. This avoids blackholes when PE-5 is disconnected from the EVPN core.

On PE-5, VPLS 1 is configured with operational group "vpls-1_45" in EVPN-MPLS and SAP 1/1/2:1 monitoring this operational group. The operational group configured under a BGP-EVPN instance cannot be configured under any other object, such as SAPs or SDP-bindings.

```
# on PE-5:
configure {
    service {
        oper-group "vpls-1_45" {
            hold-time {
                down 0
                up 0
            }
        }
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
            }
            bgp-evpn {
                evi 1
                routes {
                    mac-ip {
                        cfm-mac true
                    }
                }
                mpls 1 {
                    admin-state enable
                    oper-group "vpls-1_45"
                    auto-bind-tunnel {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        resolution any
                    }
                }
            }
            sap 1/1/2:1 {
                description "to MTU-6"
                monitor-oper-group "vpls-1_45"
                eth-cfm {
                    mep md-admin-name "domain-1" ma-admin-name "association-11" mep-id 56 {
                        admin-state enable
                        mac-address 00:00:00:00:56:05
                        fault-propagation suspend-ccm
                        ccm true
                    }
                }
            }
        }
```

Using operational groups in the EVPN service, it is possible to monitor if the PE is isolated and, if it is, trigger a Designated Forwarder switchover. The operational group associated with the EVPN-MPLS instance goes down in the following cases:

- bgp-evpn mpls is disabled
- VPLS is disabled
- all EVPN destinations associated with the instance are removed, for example, when:
  - no tunnels are available for auto-bind-tunnel resolution
  - the network ports facing the EVPN ports are down
  - the BGP sessions to the route reflector or PEs are down

Operational groups for Ethernet Segments (Port-active multi-homing)

Operational groups can be configured on single-active ESs that need to function as port-active multi-homing Ethernet Segments. 'Port-active' refers to a special single-active mode where the PE is DF or non-DF for all the services attached to the ES. The configuration of a port-active ES is as follows:

```
# on PE-2:
configure {
    service {
        oper-group "vpls-1_23" {
            hold-time {
                down 0
                up 0
            }
        }
        system {
            bgp {
                evpn {
                    ethernet-segment "ESI-23_1" {
                        admin-state enable
                        esi 01:23:00:00:00:00:01:00:00:00
                        multi-homing-mode single-active
                        oper-group "vpls-1_23"
                        ac-df-capability exclude
                        df-election {
                            es-activation-timer 3
                            service-carving-mode manual
                            manual {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                                preference {
                                    value 150              # on PE-3: value 100
                                }
                            }
                        }
                        association {
                            lag "lag-12" {                 # on PE-3: lag 13
                            }
                        }
                    }
                }
            }
        }
```

This ES operational group "vpls-1_23" can be monitored on the LAG:

```
# on PE-2:
configure {
    lag "lag-12" {
        admin-state enable
        description "to MTU-1"
        encap-type dot1q
        mode access
        standby-signaling lacp                      # default value
        monitor-oper-group "vpls-1_23"
        max-ports 64
        lacp {
            mode active
            system-id 00:00:00:01:02:01
            system-priority 1
            administrative-key 1
        }
        port 1/1/2 {
        }
    }
```

When the operational group is configured on the ES and monitored on the associated LAG:

- The status of the ES operational group is driven by the ES DF status.
  - When a node becomes NDF, the ES operational group goes down and all the SAPs in the ES go down.
- The ES operational group goes down when all the SAPs in the ES go down.
  - When all SAPs in the ES go down, the operational group goes down and the node becomes NDF.

The monitoring LAG goes down when the ES operational group is down. The LAG signals the LAG standby state to the access node. The LAG standby signaling can be configured as **lacp** or **power-off.**

```
*[ex:/configure lag "lag-12"]
A:admin@PE-2# standby-signaling ?

 standby-signaling <keyword>
 <keyword>  - (lacp|power-off)
Default    - lacp

    Way of signaling a member port to the remote side
```

- **standby-signaling lacp** signals LACP out-of-sync to the CE when the application layer instructs the LAG to become standby
- **standby-signaling power-off** brings the LAG members down, and hence the access SAPs down

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The ES and AD routes for the ES are not withdrawn because the router recognizes that the LAG becomes standby due to the ES operational group.

Some restrictions:

- Multi-chassis LAG and ES are mutually exclusive:

```
*[ex:/configure redundancy multi-chassis peer 192.0.2.2 mc-lag lag "lag-13"]
A:admin@PE-3# commit
MINOR: MGMT_CORE #4001: configure lag "lag-13" - invalid combination mc-lag <-> monitor-
oper-group
```

- LAG sub-groups are blocked:

```
*[ex:/configure lag "lag-13" port 1/1/1]
A:admin@PE-3# sub-group 2

*[ex:/configure lag "lag-13" port 1/1/1]
A:admin@PE-3# commit
MINOR: MGMT_CORE #4001: configure lag "lag-13" port 1/1/1 - invalid combination port sub-
group <-> monitor-oper-group - configure lag "lag-13" monitor-oper-group
```

- Only LAGs in access mode can monitor operational groups:

```
*[ex:/configure lag "lag-3"]
A:admin@PE-3# commit
MINOR: MGMT_CORE #3001: configure lag "lag-3" mode - monitor-oper-group not allowed when lag
 is not access
```

- Operational groups cannot be assigned to virtual ESs:

```
*[ex:/configure service system bgp evpn ethernet-segment "vESI-23_1" association lag "lag-5"
 virtual-ranges dot1q q-tag 1]
A:admin@PE-3# commit
MINOR: SVCMGR #12: configure service system bgp evpn ethernet-segment "vESI-23_1" oper-
group - Inconsistent Value error - ethernet-segment oper-group not supported with virtual
 ethernet-segment
```

- Operational groups cannot be assigned to all-active ESs:

```
*[ex:/configure service system bgp evpn]
A:admin@PE-3# commit
MINOR: SVCMGR #12: configure service system bgp evpn ethernet-segment "AA_ESI-23_1" oper-
group - Inconsistent Value error - all-active multi-homing not supported with ethernet-
segment oper-group
```

- Operational groups cannot be assigned to ESs with service-carving auto:

```
*[ex:/configure service system bgp evpn]
A:admin@PE-3# commit
MINOR: SVCMGR #12: configure service system bgp evpn ethernet-segment "ESI-23_auto" oper-
group - Inconsistent Value error - ethernet-segment oper-group not supported with service-
carving-mode auto
```

Link Loss Forwarding in EVPN-VPWS

Fault propagation in EVPN-VPWS services is supported using ETH-CFM. However, not all access nodes support ETH-CFM and, in that case, LAG **standby-signaling lacp** or **power-off** can be used instead.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Configuration

Figure 200: Sample topology with VPLS 1 shows the sample topology with VPLS 1 configured on all nodes.

Figure 200: Sample topology with VPLS 1

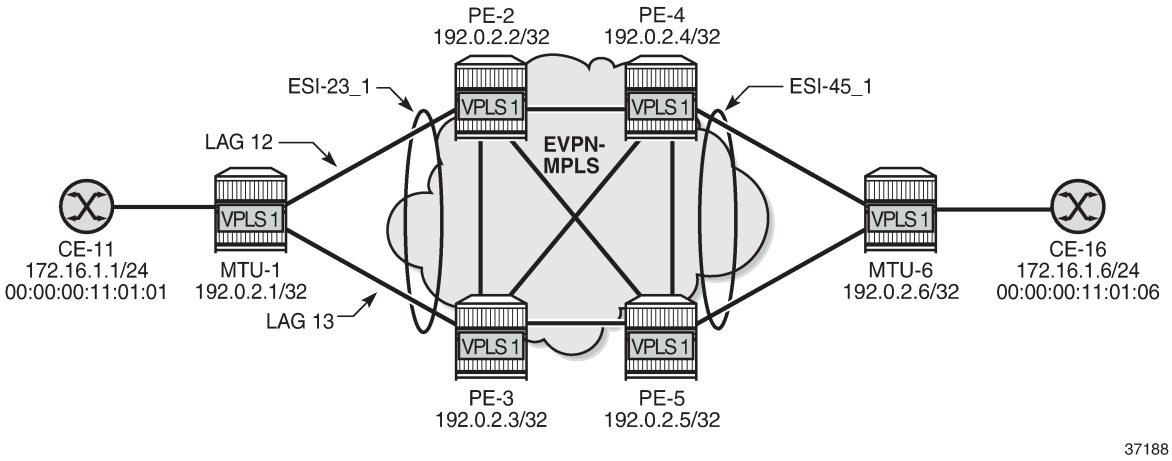

The initial configuration includes:

- Cards, MDAs, ports
- LAG 12 between PE-1 and PE-2; LAG 13 between PE-1 and PE-3
- Router interfaces between PE-2, PE-3, PE-4, and PE-5
- IS-IS on all router interfaces
- LDP between PE-2, PE-3, PE-4, and PE-5
- BGP between PE-2, PE-3, PE-4, and PE-5

For BGP, PE-2 acts as route reflector and the configuration is as follows:

```
# on PE-2:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            split-horizon true
            rapid-update {
                evpn true
            }
            group "internal" {
                peer-as 64500
                family {
                    evpn true
                }
                cluster {
                    cluster-id 192.0.2.2
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
            }
            neighbor "192.0.2.3" {
                group "internal"
            }
            neighbor "192.0.2.4" {
                group "internal"
            }
            neighbor "192.0.2.5" {
                group "internal"
            }
        }
```

Operational groups for EVPN destinations

On PE-4, single-active ES "ESI-45_1" is configured with service carving auto. Operational group "vpls-1_45" is associated with EVPN-MPLS in VPLS 1 and SAP 1/1/1:1 is monitoring that operational group. ETH-CFM is enabled on a down MEP configured on the SAP to detect SAP failures. The service configuration is as follows:

```
# on PE-4:
configure {
    service {
        oper-group "vpls-1_45" {
            hold-time {
                down 0
                up 0
            }
        }
        system {
            bgp {
                evpn {
                    ethernet-segment "ESI-45_1" {
                        admin-state enable
                        esi 01:45:00:00:00:00:01:00:00:00
                        multi-homing-mode single-active
                        df-election {
                            es-activation-timer 3
                        }
                        association {
                            port 1/1/1 {
                            }
                        }
                    }
                }
            }
        }
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
            }
            bgp-evpn {
                evi 1
                routes {
                    mac-ip {
                        cfm-mac true
                    }
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                mpls 1 {
                    admin-state enable
                    oper-group "vpls-1_45"
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            sap 1/1/1:1 {
                description "to MTU-6"
                monitor-oper-group "vpls-1_45"
                eth-cfm {
                    mep md-admin-name "domain-1" ma-admin-name "association-10" mep-id 46 {
                        admin-state enable
                        mac-address 00:00:00:00:46:04
                        ccm true
                    }
                }
            }
        }
```

The configuration on PE-5 is similar.

On MTU-6, VPLS 1 is configured with three SAPs: SAP 1/1/2:1 toward PE-4, SAP 1/1/1:1 toward PE-5, and SAP 1/2/1:1 toward CE-16. ETH-CFM MEPs are configured on SAP 1/1/1:1 and SAP 1/1/2:1. The service configuration is as follows:

```
# on MTU-6:
configure {
    service {
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            sap 1/1/1:1 {
                description "to PE-5"
                eth-cfm {
                    mep md-admin-name "domain-1" ma-admin-name "association-11" mep-id 65 {
                        admin-state enable
                        mac-address 00:00:00:00:65:06
                        ccm true
                    }
                }
            }
            sap 1/1/2:1 {
                description "to PE-4"
                eth-cfm {
                    mep md-admin-name "domain-1" ma-admin-name "association-10" mep-id 64 {
                        admin-state enable
                        mac-address 00:00:00:00:64:06
                        ccm true
                    }
                }
            }
            sap 1/2/1:1 {
                description "to CE-16"
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Initial situation without failure

On MTU-6, ETH-CFM MEP 65 receives Continuity Check (CC) messages from its remote peer 56 on PE-5:

```
[/]
A:admin@MTU-6# show eth-cfm mep 65 domain 1 association 11 all-remote-mepids

================================================================================
Eth-CFM Remote-Mep Table
================================================================================
R-mepId AD Rx CC RxRdi Port-Tlv If-Tlv Peer Mac Addr     CCM status since
--------------------------------------------------------------------------------
56         True  False Absent   Absent 00:00:00:00:56:05 12/23/2021 16:59:01
================================================================================
Entries marked with a 'T' under the 'AD' column have been auto-discovered.
```

The following command shows that PE-5 is DF for VPLS 1:

```
[/]
A:admin@PE-5# show service id 1 ethernet-segment

===============================================================================
SAP Ethernet-Segment Information
===============================================================================
SAP                   Eth-Seg                         Status
-------------------------------------------------------------------------------
1/1/2:1               ESI-45_1                        DF
===============================================================================
No sdp entries
No vxlan instance entries
```

PE-5 has full mesh with all EVPN destinations in VPLS 1:

```
[/]
A:admin@PE-5# show service id 1 evpn-mpls

===============================================================================
BGP EVPN-MPLS Dest
===============================================================================
TEP Address                   Egr Label     Num.   Mcast Last Change
                              Transport:Tnl MACs         Sup BCast Domain
-------------------------------------------------------------------------------
192.0.2.2                     524283        0      bum   12/23/2021 16:58:51
                              ldp:65539                  No
192.0.2.3                     524283        0      bum   12/23/2021 16:58:51
                              ldp:65538                  No
192.0.2.4                     524283        0      bum   12/23/2021 16:58:51
                              ldp:65537                  No
-------------------------------------------------------------------------------
Number of entries : 3
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN-MPLS Ethernet Segment Dest
===============================================================================
Eth SegId                      Num. Macs               Last Change
-------------------------------------------------------------------------------
01:23:00:00:00:00:01:00:00:00   1                      12/23/2021 16:59:37
-------------------------------------------------------------------------------
Number of entries: 1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
===============================================================================
```

Avoiding blackholes when EVPN destinations are removed

On PE-5, a failure is simulated by disabling LDP:

```
# on PE-5:
configure exclusive
    router "Base" {
        ldp {
            admin-state disable
            commit
```

With LDP disabled, PE-5 has no tunnels available for auto-bind-tunnel in VPLS 1 and all EVPN destinations are removed, as follows:

```
[/]
A:admin@PE-5# show service id 1 evpn-mpls

===============================================================================
BGP EVPN-MPLS Dest
===============================================================================
TEP Address                     Egr Label     Num.   Mcast Last Change
                                Transport:Tnl MACs         Sup BCast Domain
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================

===============================================================================
BGP EVPN-MPLS Ethernet Segment Dest
===============================================================================
Eth SegId                        Num. Macs              Last Change
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

Log 99 on PE-5 shows that the operational group "vpls-45_1" goes down and PE-5 becomes NDF in "ESI-45_1":

```
79 2021/12/23 17:01:15.697 CET MINOR: SVCMGR #2094 Base
"Ethernet Segment:ESI-45_1, EVI:1, Designated Forwarding state changed to:false"

78 2021/12/23 17:01:15.696 CET MINOR: SVCMGR #2542 Base
"Oper-group vpls-1_45 changed status to down"
```

The following command on PE-5 shows that the operational status of oper-group "vpls-45_1" is down, the EVPN-MPLS destinations are down, and the monitoring SAP 1/1/2:1 is down:

```
[/]
A:admin@PE-5# show service oper-group "vpls-1_45" detail

===============================================================================
Service Oper Group Information
===============================================================================
Oper Group       : vpls-1_45
Creation Origin  : manual                          Oper Status: down
Hold DownTime    : 0 secs                          Hold UpTime: 0 secs
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Members          : 1                                Monitoring : 1
===============================================================================

===============================================================================
Member BGP-EVPN for OperGroup: vpls-1_45
===============================================================================
SvcId:Instance (Type)                  Status
-------------------------------------------------------------------------------
1:1 (mpls)                             Inactive
-------------------------------------------------------------------------------
BGP-EVPN Entries found: 1
===============================================================================

===============================================================================
Monitoring SAPs for OperGroup: vpls-1_45
===============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.  Egr.   Adm  Opr
                                          QoS   Fltr    QoS   Fltr
-------------------------------------------------------------------------------
1/1/2:1                         1          1     none    1     none   Up   Down
-------------------------------------------------------------------------------
SAP Entries found: 1
===============================================================================
```

The following command shows that SAP 1/1/2:1 is operationally down with flags StandByForMHProtocol and OperGroupDown:

```
[/]
A:admin@PE-5# show service id 1 sap 1/1/2:1

===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 1
SAP                : 1/1/2:1                 Encap             : q-tag
Description        : to MTU-6
Admin State        : Up                      Oper State        : Down
Flags              : StandByForMHProtocol
                     OperGroupDown
Multi Svc Site     : None
Last Status Change : 12/23/2021 17:01:16
Last Mgmt Change   : 12/23/2021 16:58:49
===============================================================================
```

With ETH-CFM enabled, log 99 on MTU-6 shows that local MEP 65 did not receive a Continuity Check Message (CCM) from the remote MEP:

```
56 2021/12/23 17:01:19.288 CET MINOR: ETH_CFM #2001 Base
"MEP 1/11/65 highest defect is now defRemoteCCM"
```

PE-4 receives the following BGP-EVPN withdrawal messages:

```
33 2021/12/23 17:01:15.700 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 129
    Flag: 0x90 Type: 15 Len: 125 Multiprotocol Unreachable NLRI:
        Address Family EVPN
        Type: EVPN-ETH-SEG Len: 23 RD: 192.0.2.5:0
              ESI: 01:45:00:00:00:00:01:00:00:00, IP-Len: 4 Orig-IP-Addr: 192.0.2.5
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        Type: EVPN-AD Len: 25 RD: 192.0.2.5:1 ESI: 01:45:00:00:00:00:01:00:00:00,
              tag: 0 Label: 0 (Raw Label: 0x0) PathId:
        Type: EVPN-MAC Len: 33 RD: 192.0.2.5:1 ESI: ESI-0, tag: 0, mac len: 48
              mac: 00:00:00:11:01:06, IP len: 0, IP: NULL, label1: 0
        Type: EVPN-MAC Len: 33 RD: 192.0.2.5:1 ESI: ESI-0, tag: 0, mac len: 48
              mac: 00:00:00:00:65:06, IP len: 0, IP: NULL, label1: 0
"
```

The following command on PE-4 shows that PE-4 is the DF and the only DF candidate in "ESI-45_1" for VPLS 1:

```
[/]
A:admin@PE-4# show service system bgp-evpn ethernet-segment name "ESI-45_1"
                                                                        evi evi-1 1

===============================================================================
EVI DF and Candidate List
===============================================================================
EVI           SvcId         Actv Timer Rem      DF  DF Last Change
-------------------------------------------------------------------------------
1             1             0                   yes 12/23/2021 17:01:19
===============================================================================

===============================================================================
DF Candidates                          Time Added           Oper Pref  Do Not
                                                              Value    Preempt
-------------------------------------------------------------------------------
192.0.2.4                              12/23/2021 16:58:45  0          Disabl*
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

Finally, the failure is restored by re-enabling LDP on PE-5:

```
# on PE-5:
configure exclusive
    router "Base" {
        ldp {
            admin-state enable
            commit
```

Operational groups for ES (Port-Active Multi-Homing)

On PE-2 and PE-3, operational group vpls-1_23 is configured and associated with ES " ESI-23_1", but not configured or monitored in VPLS 1. The service configuration on PE-3 is as follows:

```
# on PE-3:
configure {
    service {
        oper-group "vpls-1_23" {
            hold-time {
                down 0
                up 0
            }
        }
        system {
            bgp {
                evpn {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    ethernet-segment "ESI-23_1" {
                        admin-state enable
                        esi 01:23:00:00:00:00:01:00:00:00
                        multi-homing-mode single-active
                        oper-group "vpls-1_23"
                        ac-df-capability exclude
                        df-election {
                            es-activation-timer 3
                            service-carving-mode manual
                            manual {
                                preference {
                                    value 100           # on PE-2: value 150
                                }
                            }
                        }
                        association {
                            lag "lag-13" {
                            }
                        }
                    }
                }
            }
        }
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
            }
            bgp-evpn {
                evi 1
                mpls 1 {
                    admin-state enable
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            sap lag-13:1 {                          # on PE-2: sap lag-12:1
                description "to MTU-1"
            }
        }
```

LAG 12 on PE-2 and LAG 13 on PE-3 monitor operational group "vpls-1_23". The **monitor-oper-group** command can be added to the LAG:

```
# on PE-3:
configure {
    lag "lag-13" {
        admin-state enable
        description "to MTU-1"
        encap-type dot1q
        mode access
        standby-signaling lacp      # default value
        monitor-oper-group "vpls-1_23"
        max-ports 64
        lacp {
            mode active
            system-id 00:00:00:01:03:01
            system-priority 1
            administrative-key 1
        }
        port 1/1/1 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
```

Note:
In this example, MTU-1 is connected to PE-2 and PE-3 through two different LAGs, however, this port-active multi-homing mode also supports the use of a single LAG on MTU-1. If a single LAG was used on MTU-1, the LAG ports on PE-2 and PE-3 must be configured with the same LACP parameters (administrative-key, system-id and system-priority) to ensure that PE-2 and PE-3 show themselves as a single system to MTU-1.

EVPN single-active multi-homing PEs that are elected as NDF must notify their attached access nodes to prevent these from sending traffic to the NDF. In this port-active multi-homing mode, ETH-CFM is not used, and other notification mechanisms are needed, such as LAG standby signaling (**lacp** or **power-off**). When the EVPN application layer instructs the LAG to become standby as a result of the NDF status, the behavior is as follows:

- the **lacp** option signals LACP out-of-sync to MTU-1
- the **power-off** option brings down the LAG ports connected to MTU-1

MTU-1 is connected to PE-2 and PE-3 using two different access LAGs with encapsulation dot1q and at least one port in each LAG. Any encapsulation type is supported in the LAGs. The LAG configuration is as follows:

```
# on MTU-1:
configure {
    lag "lag-12" {
        admin-state enable
        description "to PE-2"
        encap-type dot1q
        mode access
        max-ports 64
        lacp {
            mode active
            administrative-key 32768
        }
        port 1/1/1 {
        }
    }
    lag "lag-13" {
        admin-state enable
        description "to PE-3"
        encap-type dot1q
        mode access
        max-ports 64
        lacp {
            mode active
            administrative-key 32769
        }
        port 1/1/2 {
        }
    }
```

On MTU-1, VPLS 1 is configured as follows:

```
# on MTU-1:
configure {
    service {
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            customer "1"
            sap 1/2/1:1 {
                description "to CE-11"
            }
            sap lag-12:1 {
                description "to PE-2"
            }
            sap lag-13:1 {
                description "to PE-3"
            }
        }
```

Initial situation without failures

PE-2 is DF for VPLS 1:

```
[/]
A:admin@PE-2# show service id 1 ethernet-segment

===============================================================================
SAP Ethernet-Segment Information
===============================================================================
SAP                   Eth-Seg                         Status
-------------------------------------------------------------------------------
lag-12:1              ESI-23_1                        DF
===============================================================================
No sdp entries
No vxlan instance entries
```

```
[/]
A:admin@PE-3# show service id 1 ethernet-segment

===============================================================================
SAP Ethernet-Segment Information
===============================================================================
SAP                   Eth-Seg                         Status
-------------------------------------------------------------------------------
lag-13:1              ESI-23_1                        NDF
===============================================================================
No sdp entries
No vxlan instance entries
```

On NDF PE-3, operational group "vpls-1_23" is operationally down, which has an impact on the operational status of the monitoring LAG, as follows:

```
[/]
A:admin@PE-3# show service oper-group "vpls-1_23" detail

===============================================================================
Service Oper Group Information
===============================================================================
Oper Group       : vpls-1_23
Creation Origin  : manual                        Oper Status: down
Hold DownTime    : 0 secs                        Hold UpTime: 0 secs
Members          : 1                             Monitoring : 1
===============================================================================

===============================================================================
Member Ethernet-Segment for OperGroup: vpls-1_23
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Ethernet-Segment                        Status
-------------------------------------------------------------------------------
ESI-23_1                                Inactive
-------------------------------------------------------------------------------
Ethernet-Segment Entries found: 1
===============================================================================

===============================================================================
Monitoring LAG for OperGroup: vpls-1_23
===============================================================================
Lag-id         Adm       Opr       Weighted  Threshold Up-Count  Act/Stdby
    name
-------------------------------------------------------------------------------
13             up        down      No        0         0         N/A
    lag-13
-------------------------------------------------------------------------------
LAG Entries found: 1
===============================================================================
```

The following command shows that SAP lag-13:1 is operationally down on PE-3 with flags PortOperDown and StandByForMHProtocol:

```
[/]
A:admin@PE-3# show service id 1 sap lag-13:1

===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 1
SAP                : lag-13:1                Encap             : q-tag
Description        : to MTU-1
Admin State        : Up                      Oper State        : Down
Flags              : PortOperDown StandByForMHProtocol
Multi Svc Site     : None
Last Status Change : 12/23/2021 16:58:33
Last Mgmt Change   : 12/23/2021 16:58:33
===============================================================================
```

The following command on PE-3 shows that LAG 13 has LACP standby signaling enabled to the MTU-1. LAG 13 is operationally down because the operational group is down.

```
[/]
A:admin@PE-3# show lag 13 detail

===============================================================================
LAG Details
===============================================================================
Description        : N/A
-------------------------------------------------------------------------------
Details
-------------------------------------------------------------------------------
Lag-id              : 13                    Mode                : access
Lag-name            : lag-13
Adm                 : up                    Opr                 : down
---snip---

Standby Signaling   : lacp
---snip---

Monitor oper group  : vpls-1_23
Oper group status   : down
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Adaptive loadbal.   : disabled              Tolerance           : N/A

-------------------------------------------------------------------------------
Port-id        Adm     Act/Stdby Opr     Primary   Sub-group     Forced  Prio
-------------------------------------------------------------------------------
1/1/1          up      active    down    yes       1            -       32768

-------------------------------------------------------------------------------
Port-id        Role      Exp   Def   Dist  Col   Syn   Aggr  Timeout  Activity
-------------------------------------------------------------------------------
1/1/1          actor     No    No    No    No    No    Yes   Yes      Yes
1/1/1          partner   No    No    No    No    Yes   Yes   Yes      Yes
===============================================================================
```

DF switchover

To trigger a DF switchover, the preference value is modified on PE-2, as follows:

```
# on PE-2:
configure exclusive
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "ESI-23_1" {
                        df-election {
                            es-activation-timer 3
                            service-carving-mode manual
                            manual {
                                preference {
                                    value 50
                                }
                            }
                        }
```

Figure 201: DF switchover in single-active ESI-23_1 shows a DF switchover from PE-2 to PE-3. PE-2 becomes the NDF and LAG 12 is in standby.

Figure 201: DF switchover in single-active ESI-23_1

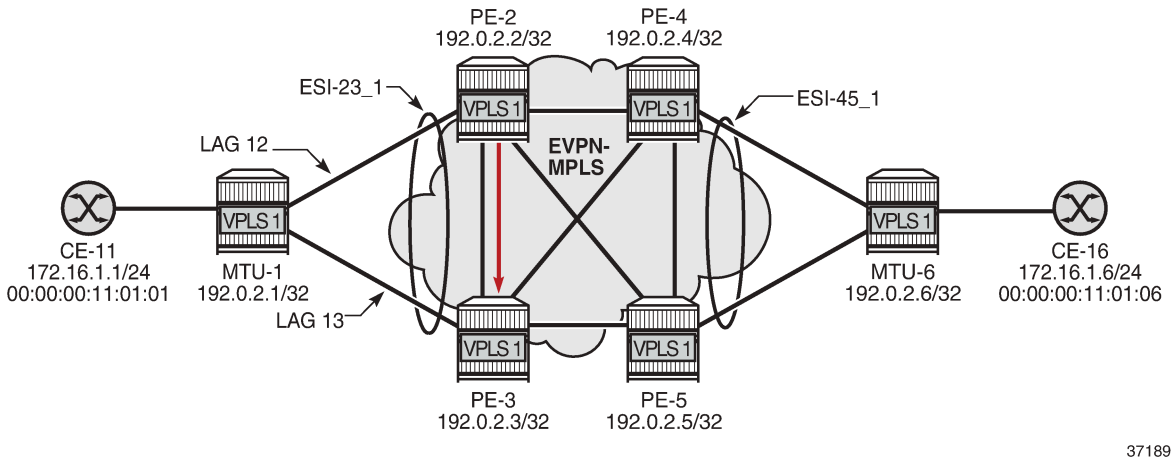

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Log 99 on PE-2 shows that SAP lag-12:1 goes down, the ES operational group goes down, the monitoring LAG 12 goes down, port 1/1/2 goes down, and subsequently an LACP out-of-sync message is sent:

```
110 2021/12/23 17:15:27.781 CET WARNING: LAG #2007 Base LAG
"LAG lag-12 : partner oper state bits changed on member 1/1/2 : [sync FALSE -> TRUE] [expired
 TRUE -> FALSE] [defaulted TRUE -> FALSE]"

109 2021/12/23 17:15:27.781 CET WARNING: LAG #2007 Base LAG
"LAG lag-12 : LACP RX state machine entered current state on member 1/1/2"

108 2021/12/23 17:15:27.777 CET MAJOR: SVCMGR #2210 Base
"Processing of an access port state change event is finished and the status of all affected
 SAPs on port lag-12 has been updated."

107 2021/12/23 17:15:27.777 CET WARNING: SNMP #2004 Base lag-12
"Interface lag-12 is not operational"

106 2021/12/23 17:15:27.777 CET MINOR: SVCMGR #2203 Base
"Status of SAP lag-12:1 in service 1 (customer 1) changed to admin=up oper=down flags=Mh
Standby"

105 2021/12/23 17:15:27.777 CET WARNING: SNMP #2004 Base 1/1/2
"Interface 1/1/2 is not operational"

104 2021/12/23 17:15:27.777 CET WARNING: LAG #2006 Base LAG
"LAG lag-12 : initializing LACP, all members will be brought down"

103 2021/12/23 17:15:27.777 CET MINOR: SVCMGR #2094 Base
"Ethernet Segment:ESI-23_1, EVI:1, Designated Forwarding state changed to:false"

102 2021/12/23 17:15:27.777 CET MINOR: SVCMGR #2542 Base
"Oper-group vpls-1_23 changed status to down"
```

On PE-3, log 99 shows that PE-3 becomes DF for "ESI-23_1" and operational group "vpls-1_23", interface 1/1/1, and LAG 13 are operationally up.

```
112 2021/12/23 17:15:31.753 CET WARNING: LAG #2007 Base LAG
"LAG lag-13 : partner oper state bits changed on member 1/1/1 :  [collecting FALSE -> TRUE]"

111 2021/12/23 17:15:31.734 CET MAJOR: SVCMGR #2210 Base
"Processing of an access port state change event is finished and the status of all affected
 SAPs on port lag-13 has been updated."

110 2021/12/23 17:15:31.733 CET WARNING: SNMP #2005 Base lag-13
"Interface lag-13 is operational"

109 2021/12/23 17:15:31.733 CET WARNING: SNMP #2005 Base 1/1/1
"Interface 1/1/1 is operational"

108 2021/12/23 17:15:30.831 CET MAJOR: SVCMGR #2210 Base
"Processing of an access port state change event is finished and the status of all affected
 SAPs on port lag-13 has been updated."

107 2021/12/23 17:15:30.811 CET MINOR: SVCMGR #2094 Base
"Ethernet Segment:ESI-23_1, EVI:1, Designated Forwarding state changed to:true"

106 2021/12/23 17:15:30.811 CET MINOR: SVCMGR #2542 Base
"Oper-group vpls-1_23 changed status to up"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Link Loss Forwarding in EVPN-VPWS

Fault propagation in EVPN-VPWS services is supported using ETH-CFM, but also using LAG **standby-signaling lacp** or **power-off**.

Figure 202: Sample topology with Epipe 2 shows the sample topology with Epipe 2.

Figure 202: Sample topology with Epipe 2

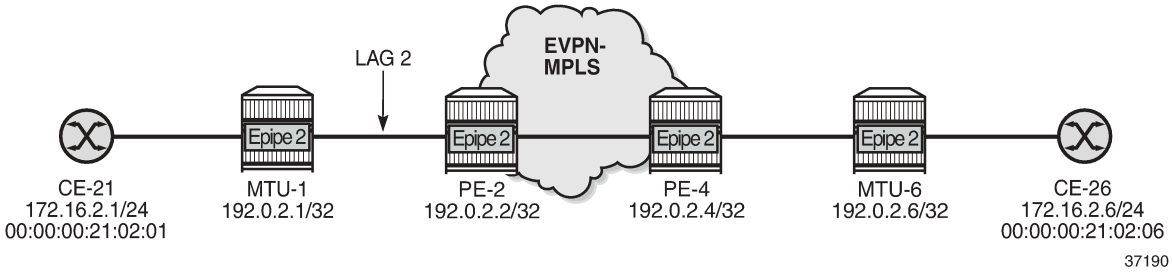

The configuration on MTU-1 is as follows:

```
# on MTU-1:
configure {
    lag "lag-2" {
        admin-state enable
        encap-type dot1q
        mode access
        max-ports 64
        lacp {
            administrative-key 32770
        }
        port 1/1/5 {
        }
    }
    service {
        epipe "Epipe 2" {
            admin-state enable
            service-id 2
            customer "1"
            sap 1/2/1:2 {
            }
            sap lag-2:2 {
            }
        }
```

On PE-2, operational group "llf-1" is configured and associated to EVPN-MPLS. LAG 2 monitors this operational group.

```
# on PE-2:
configure {
    lag "lag-2" {
        admin-state enable
        encap-type dot1q
        mode access
        standby-signaling lacp
        monitor-oper-group "llf-1"
        max-ports 64
        lacp {
            mode active
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            system-id 00:00:00:00:12:01
            system-priority 1
            administrative-key 2
        }
        port 1/1/5 {
        }
    }
    service {
        oper-group "llf-1" {
            hold-time {
                down 0
                up 0
            }
        }
        epipe "Epipe 2" {
            admin-state enable
            service-id 2
            customer "1"
            bgp 1 {
            }
            sap lag-2:2 {
            }
            bgp-evpn {
                evi 2
                local-attachment-circuit "ac-1_2" {
                    eth-tag 12
                }
                remote-attachment-circuit "ac-6_2" {
                    eth-tag 62
                }
                mpls 1 {
                    admin-state enable
                    oper-group "llf-1"
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
        }
```

The configuration on PE-4 is as follows:

```
# on PE-4:
configure {
    service {
        epipe "Epipe 2" {
            admin-state enable
            service-id 2
            customer "1"
            bgp 1 {
            }
            sap 1/1/5:2 {
            }
            bgp-evpn {
                evi 2
                local-attachment-circuit "ac-6_2" {
                    eth-tag 62
                }
                remote-attachment-circuit "ac-1_2" {
                    eth-tag 12
                }
                mpls 1 {
                    admin-state enable
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
        }
```

Figure 203: LLF in Epipe 2 - PE-4 failure shows when a failure occurs on PE-4.

Figure 203: LLF in Epipe 2 - PE-4 failure

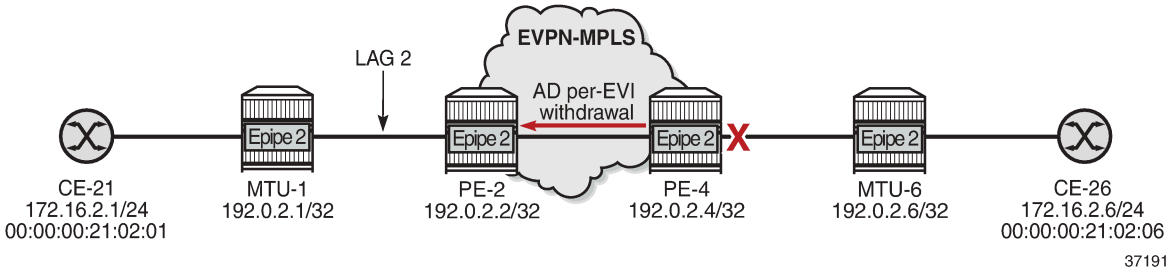

The failure is simulated on PE-4 by disabling port 1/1/5 toward MTU-6.

```
# on PE-4:
configure exclusive
    port 1/1/5 {
        admin-state disable
        commit
```

When the link between PE-4 and MTU-6 fails, PE-4 withdraws the AD per-EVI route for Epipe 2. PE-2 receives the following AD per-EVI withdrawal from PE-4:

```
155 2021/12/23 17:18:37.217 CET MINOR: DEBUG #2001 Base Peer 1: 192.0.2.4
"Peer 1: 192.0.2.4: UPDATE
Peer 1: 192.0.2.4 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 34
    Flag: 0x90 Type: 15 Len: 30 Multiprotocol Unreachable NLRI:
        Address Family EVPN
        Type: EVPN-AD Len: 25 RD: 192.0.2.4:2 ESI: ESI-0, tag: 62
              Label: 0 (Raw Label: 0x0) PathId:
"
```

Upon receiving this AD per-EVI route, Epipe 2 goes operationally down on PE-2:

```
[/]
A:admin@PE-2# show service id 2 base | match "Oper State"
Admin State       : Up                 Oper State        : Down
```

Operational group "llf-1" goes down when the Epipe is operationally down:

```
[/]
A:admin@PE-2# show lag 2 detail | match "per group"
Monitor oper group  : llf-1
Oper group status   : down
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

On PE-2, the detailed information for operational group "llf-1" shows that the operational group and the monitoring LAG are down.

```
[/]
A:admin@PE-2# show service oper-group "llf-1" detail

===============================================================================
Service Oper Group Information
===============================================================================
Oper Group       : llf-1
Creation Origin  : manual                          Oper Status: down
Hold DownTime    : 0 secs                          Hold UpTime: 0 secs
Members          : 1                               Monitoring : 1
===============================================================================

===============================================================================
Member BGP-EVPN for OperGroup: llf-1
===============================================================================
SvcId:Instance (Type)                  Status
-------------------------------------------------------------------------------
2:1 (mpls)                             Inactive
-------------------------------------------------------------------------------
BGP-EVPN Entries found: 1
===============================================================================

===============================================================================
Monitoring LAG for OperGroup: llf-1
===============================================================================
Lag-id         Adm       Opr       Weighted  Threshold Up-Count  Act/Stdby
    name
-------------------------------------------------------------------------------
2              up        down      No        0         0         N/A
    lag-2
-------------------------------------------------------------------------------
LAG Entries found: 1
===============================================================================
```

PE-2 signals the fault based on the configuration of the LAG standby signaling:

- If the LAG standby signaling is power-off, PE-2 brings down the ports in the LAG.
- If the LACP standby signaling is configured, PE-2 signals an LACP out-of-sync on the LAG ports.

In either case, MTU-1 stops forwarding traffic to PE-2.

The following debug message in log 99 on MTU-1 shows that MTU-1 received an LACP out-of-sync message for port 1/1/5 of LAG 2:

```
154 2021/12/23 17:18:37.216 CET WARNING: LAG #2007 Base LAG
"LAG lag-2 : partner oper state bits changed on member 1/1/5 : [sync TRUE -> FALSE] [collecting
 TRUE -> FALSE]"
```

The following debug messages in log 99 on MTU-1 show that LAG 2 and interface 1/1/5 are not operational:

```
156 2021/12/23 17:18:37.217 CET WARNING: SNMP #2004 Base lag-2
"Interface lag-2 is not operational"

155 2021/12/23 17:18:37.216 CET WARNING: SNMP #2004 Base 1/1/5
"Interface 1/1/5 is not operational"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

On MTU-1, LAG 2 is operationally down:

```
[/]
A:admin@MTU-1# show lag 2

===============================================================================
Lag Data
===============================================================================
Lag-id         Adm     Opr     Weighted Threshold Up-Count MC Act/Stdby
    name
-------------------------------------------------------------------------------
2              up      down    No       0         0        N/A
    lag-2
===============================================================================
```

Conclusion

Operational groups can be useful in EVPN services to avoid blackholes when a PE is disconnected from the EVPN core. Failures can be propagated by the PEs to access nodes, either by ETH-CFM or LAG standby signaling.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

P2MP mLDP FEC Resolution for BGP-LU in EVPN

This chapter provides information about P2MP mLDP FEC Resolution for BGP-LU in EVPN.

Topics in this chapter include:

Applicability

This chapter was initially written for SR OS Release 16.0.R3, but the MD-CLI in the current edition is based on SR OS Release 21.5.R1. Recursive and non-recursive multicast Label Distribution Protocol (mLDP) Forwarding Equivalence Class (FEC) resolution for BGP Labeled Unicast (BGP-LU) is supported in SR OS Release 15.0.R1 or later; see the P2MP mLDP Inter-AS Model C for EVPN-MPLS Services chapter.

In SR OS Release 15.0.R4, and later, a leaf node in an MVPN can generate non-recursive mLDP mapping messages even if the root IP address is resolved using BGP-LU, without the need to leak BGP routes to IGP and LDP. In SR OS Release 16.0.R1, this is also supported for EVPN-MPLS services.

Overview

In inter-AS and intra-AS scenarios, recursive and non-recursive FEC label mapping messages can be used to set up the mLDP tree. In the P2MP mLDP Inter-AS Model C for EVPN-MPLS Services chapter, recursive and non-recursive mLDP FEC resolution is documented for inter-AS model C.

This chapter describes the following use cases for recursive and non-recursive mLDP FEC resolution for BGP-LU:

- P2MP mLDP FEC resolution for inter-AS model C
- P2MP mLDP FEC resolution for seamless MPLS

Some routers do not support recursive mLDP FEC, so basic non-recursive mLDP FEC is used instead. The non-recursive mLDP FEC resolution does not require the root IP address to be leaked from BGP to IGP and LDP. This is different from the configuration in the P2MP mLDP Inter-AS Model C for EVPN-MPLS Services chapter.

Figure 204: Example topology for inter-AS model C shows the example topology for inter-AS model C with the configured protocols (IGP, LDP, BGP). Root node PE-1 is situated in AS 64501 and leaf node PE-4 in AS 64502. P-2 and P-3 are AS Border Routers (ASBRs) that are configured with next-hop-self (NHS). VPLS 1 is configured on root node PE-1 and leaf node PE-4, and is EVPN-MPLS enabled. The example topology for seamless MPLS is similar, but P-2 and P-3 will then act as Area Border Routers (ABRs) and IGP instance 0 is configured between them.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 204: Example topology for inter-AS model C

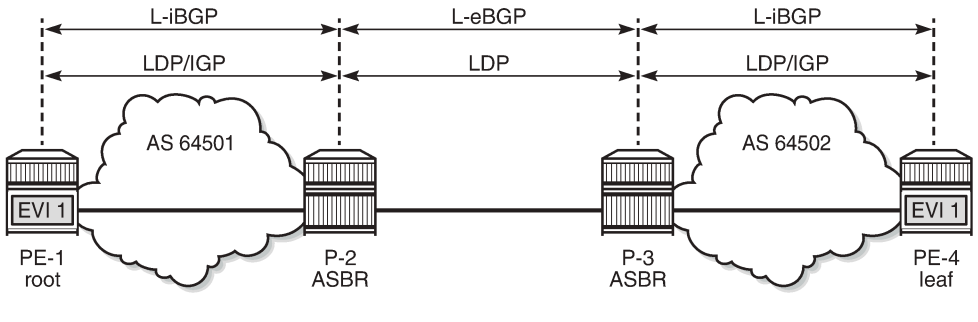

Recursive mLDP FEC resolution requires the nodes in a remote AS (or remote area in case of seamless MPLS) to support GRT recursive FEC type 7 to join the root node.

- PE-4 has a labeled BGP route to root PE-1 with next-hop P-3 in its route table. If PE-4 supports it, it sends a GRT recursive FEC type 7 label mapping message with inner root PE-1 and root P-3.
- P-3 has a labeled BGP route to PE-1 with next-hop P-2. When P-3 receives the mLDP label mapping message from PE-4, it generates its own GRT recursive FEC type 7 message with inner root PE-1 and root P-2.
- P-2 has an IGP route to root PE-1. When P-2 receives the mLDP label mapping message from P-3, it generates a non-recursive FEC type 1 message with root PE-1.

Figure 205: mLDP FEC label mapping messages for inter-AS model C shows the mLDP label mapping messages for inter-AS model C.

Figure 205: mLDP FEC label mapping messages for inter-AS model C

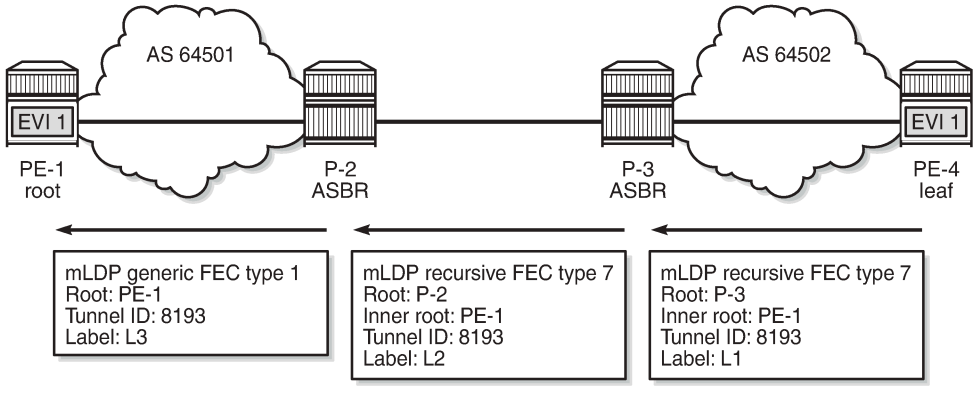

However, if the leaf node PE-4 does not support GRT recursive FEC type 7, it is possible to generate a non-recursive FEC type 1 label mapping message with root PE-1 to the local ASBR that supports GRT recursive FEC type 7. The following command generates only generic FEC type 1 label mapping messages with PE-1 as the root, on the leaf node PE-4:

```
# on PE-4:
configure {
    router "Base" {
        ldp {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            generate-basic-fec-only true
```

Note:
SR OS always generates a recursive FEC if the root node is resolved via BGP; if the root node is resolved via IGP, basic FEC is generated instead. The only way to not generate a recursive FEC when the root is resolved via BGP is by configuring the **generate-basic-fec-only** command.

Figure 206: Non-recursive mLDP FEC for inter-AS model C shows the non-recursive mLDP label mapping messages for inter-AS model C.

Figure 206: Non-recursive mLDP FEC for inter-AS model C

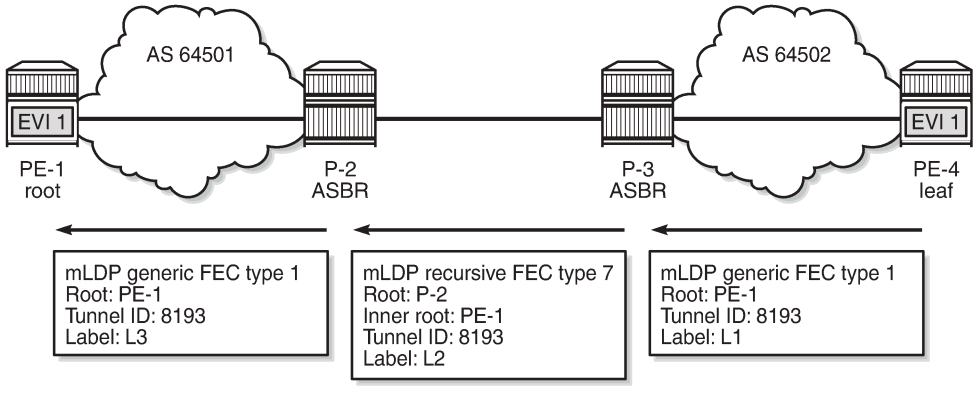

It is also possible that the ASBR routers do not support GRT recursive FEC either. The same **generate-basic-fec-only** command can be configured on all these nodes, which will then generate basic FEC type 1 label mapping messages with root address 192.0.2.1 to the next-hop.

Configuration

Figure 207: Example topology shows the example topology with four nodes.

Figure 207: Example topology

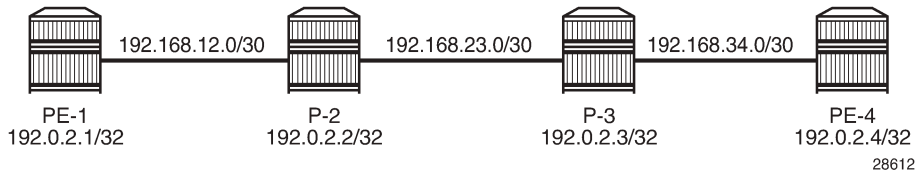

The initial configuration includes the following:

- Cards, MDAs, ports
- Router interfaces

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Inter-AS model C

Figure 204: Example topology for inter-AS model C showed the example topology for inter-AS model C. The following is configured for that topology. For a detailed explanation of the configuration, see the P2MP mLDP Inter-AS Model C for EVPN-MPLS Services chapter.

- Within each AS, OSPF is configured as IGP (alternatively, IS-IS can be used).
- LDP is enabled within each AS.
- LDP is enabled between the ASBRs using the interface IP addresses 192.168.23.x.
- On the ASBRs, a static route 192.168.23.y/32 for the interface IP address on the ASBR peer is configured (with mask /32 instead of /30). When a label mapping message is received for an LDP FEC prefix, the next-hop for a FEC prefix is resolved using the routing table. The FEC is installed in the Label Information Base (LIB) if the next-hop matches a /32 route entry.
- BGP is configured on all nodes for the labeled IPv4 address family. An export policy exports the system IP addresses of the root and leaf nodes PE-1 and PE-4.
- A multi-hop BGP session is established between PE-1 and PE-4 for the EVPN address family, allowing inclusive multicast EVPN routes to be exchanged.
- EVPN-MPLS VPLS 1 is configured on PE-1 and PE-4 with mLDP enabled. PE-1 is configured as root node.

The BGP configuration on PE-1 is as follows. The BGP configuration on PE-4 is similar, but with different neighbors and AS numbers. The export policy is identical.

```
# on PE-1:
configure {
    policy-options {
        prefix-list "sysPE" {
            prefix 192.0.2.0/29 type longer {
            }
        }
        policy-statement "PE-sys-to-labeled-BGP" {
            entry 10 {
                from {
                    prefix-list ["sysPE"]
                    protocol {
                        name [direct]
                    }
                }
                to {
                    protocol {
                        name [bgp-label]
                    }
                }
                action {
                    action-type accept
                }
            }
        }
    }
    router "Base" {
        autonomous-system 64501
        bgp {
            split-horizon true
            group "eBGP" {
                multihop 10
                type external
                peer-as 64502
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                family {
                    evpn true
                }
                ebgp-default-reject-policy {
                    import false
                    export false
                }
            }
            group "iBGP" {
                type internal
            }
            neighbor 192.0.2.2 {
                group "iBGP"
                family {
                    label-ipv4 true
                }
                export {
                    policy ["PE-sys-to-labeled-BGP"]
                }
            }
            neighbor "192.0.2.4" {
                group "eBGP"
            }
```

On PE-1, VPLS 1 is configured as follows. The service configuration on PE-4 is similar, but with different RT values and without the **root-and-leaf** parameter.

```
# on PE-1:
configure {
    service {
        vpls "EVI-1" {
            admin-state enable
            service-id 1
            customer 1
            bgp 1 {
                route-target {
                    export target:64501:1
                    import target:64502:1
                }
            }
            bgp-evpn {
                evi 1
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            sap 1/2/1:1 {
            }
            provider-tunnel {
                inclusive {
                    admin-state enable
                    owner bgp-evpn-mpls
                    root-and-leaf true     # PE-1 is configured as root node
                    mldp
                }
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

On P-2, the following static route with mask /32 is configured for the interface IP address of the peer ASBR. The configuration on P-3 is similar.

```
# on P-2:
configure {
    router "Base" {
        static-routes {
            route 192.168.23.2/32 route-type unicast {
                next-hop 192.168.23.2 {
                    admin-state enable
                }
            }
```

On P-2, the BGP and LDP configuration is as follows. The configuration on P-3 is similar.

```
# on P-2:
configure {
    policy-options {
        community "64501:0" {
            member "64501:0" { }
        }
        community "64502:0" {
            member "64502:0" { }
        }
        policy-statement "export-bgp" {
            entry 10 {
                from {
                    protocol {
                        name [bgp-label]
                    }
                }
                action {
                    action-type accept
                    origin igp
                    community {
                        add ["64501:0"]
                    }
                }
            }
        }
        policy-statement "import-bgp" {
            entry 10 {
                from {
                    community {
                        name "64502:0"
                    }
                }
                action {
                    action-type accept
                }
            }
        }
    }
    router "Base" {
        autonomous-system 64501
        bgp {
            split-horizon true
            group "eBGP" {
                type external
                import {
                    policy ["import-bgp"]
                }
                export {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    policy ["export-bgp"]
                }
            }
            group "iBGP" {
                type internal
            }
            neighbor "192.0.2.1" {
                group "iBGP"
                family {
                    label-ipv4 true
                }
                cluster {
                    cluster-id 192.0.2.2
                }
            }
            neighbor "192.168.23.2" {
                advertise-inactive true
                group "eBGP"
                next-hop-self true
                peer-as 64502
                family {
                    label-ipv4 true
                }
            }
        }
        ldp {
            interface-parameters {
                interface "int-P-2-P-3" {
                    ipv4 {
                        admin-state enable
                        local-lsr-id {
                            interface-name "int-P-2-P-3"
                        }
                    }
                }
                interface "int-P-2-PE-1" {
                    ipv4 {
                        admin-state enable
                    }
                }
            }
```

Leaf node PE-4 has a labeled BGP route toward root node PE-1 using next-hop 192.0.2.3, as follows:

```
[/]
A:admin@PE-4# show router route-table

===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
192.0.2.1/32                                  Remote  BGP_LABEL 00h00m08s  170
       192.0.2.3 (tunneled)                                        10
192.0.2.3/32                                  Remote  OSPF      00h02m33s  10
       192.168.34.1                                               10
192.0.2.4/32                                  Local   Local     00h02m34s  0
       system                                                     0
192.168.34.0/30                               Local   Local     00h02m34s  0
       int-PE-4-P-3                                               0
-------------------------------------------------------------------------------
No. of Routes: 4
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

Likewise, ASBR P-3 has a labeled BGP route toward root node PE-1 using next-hop 192.168.23.1, as follows:

```
[/]
A:admin@P-3# show router route-table 192.0.2.1/32

===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
192.0.2.1/32                                  Remote  BGP_LABEL 00h01m35s  170
       192.168.23.1                                                 0
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

P-2 has an IGP route toward root node PE-1, as follows:

```
[/]
A:admin@P-2# show router route-table 192.0.2.1/32

===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
192.0.2.1/32                                  Remote  OSPF      00h03m49s  10
       192.168.12.1                                                 10
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

Recursive mLDP FEC resolution for inter-AS model C

With the preceding configuration, leaf node PE-4 sends a recursive mLDP FEC label mapping message with PE-1 as inner root and P-3 as root. On PE-4, the number of GRT recursive mLDP bindings is 1, as follows:

```
[/]
A:admin@PE-4# show router ldp bindings active p2mp summary ipv4
 No. of Generic IPv4 P2MP Active Bindings: 0
 No. of In-Band-SSM IPv4 P2MP Active Bindings: 0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
 No. of In-Band-VPN-SSM IPv4 P2MP Active Bindings: 0
 No. of In-Band-SSM IPv4 P2MP Active Bindings: 0
 No. of VPN Recursive with Generic IPv4 P2MP Active Bindings: 0
 No. of GRT Recursive with Generic IPv4 P2MP Active Bindings: 1
```

```
[/]
A:admin@PE-4# show router ldp bindings p2mp opaque-type grt-recursive ipv4 detail

===============================================================================
LDP Bindings (IPv4 LSR ID 192.0.2.4)
             (IPv6 LSR ID ::)
===============================================================================
Label Status:
        U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
        e - Label ELC
FEC Flags:
        LF - Lower FEC, UF - Upper FEC, M - Community Mismatch,
        BA - ASBR Backup FEC
===============================================================================
LDP GRT Recursive with Generic IPv4 P2MP Bindings
===============================================================================
-------------------------------------------------------------------------------
P2MP Type      : 7                 P2MP-Id    : 8193
Root-Addr      : 192.0.2.3
InnerRoot-Addr : 192.0.2.1
-------------------------------------------------------------------------------
Peer           : 192.0.2.3:0
Ing Lbl        : 524282U
Egr Lbl        :   --
Egr Int/LspId  :   --
EgrNextHop     :   --
Egr. Flags     : None              Ing. Flags : None
===============================================================================
No. of GRT Recursive with Generic IPv4 P2MP Bindings: 1
===============================================================================
```

On P-3, there are two GRT recursive mLDP bindings with PE-1 as inner root, as follows:

```
[/]
A:admin@P-3# show router ldp bindings active p2mp summary ipv4
 No. of Generic IPv4 P2MP Active Bindings: 0
 No. of In-Band-SSM IPv4 P2MP Active Bindings: 0
 No. of In-Band-VPN-SSM IPv4 P2MP Active Bindings: 0
 No. of In-Band-SSM IPv4 P2MP Active Bindings: 0
 No. of VPN Recursive with Generic IPv4 P2MP Active Bindings: 0
 No. of GRT Recursive with Generic IPv4 P2MP Active Bindings: 2
```

The first GRT recursive mLDP binding has root 192.0.2.3 (P-3), which is the Lower FEC (LF) toward its peer PE-4; the second GRT recursive mLDP binding has root 192.168.23.1 (P-2), which is the Upper FEC (UF) toward the inner root PE-1, as follows:

```
[/]
A:admin@P-3# show router ldp bindings p2mp opaque-type grt-recursive ipv4 detail

===============================================================================
LDP Bindings (IPv4 LSR ID 192.0.2.3)
             (IPv6 LSR ID ::)
===============================================================================
Label Status:
        U - Label In Use, N - Label Not In Use, W - Label Withdrawn
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
        e - Label ELC
FEC Flags:
        LF - Lower FEC, UF - Upper FEC, M - Community Mismatch,
        BA - ASBR Backup FEC
===============================================================================
LDP GRT Recursive with Generic IPv4 P2MP Bindings
===============================================================================
-------------------------------------------------------------------------------
P2MP Type      : 7                  P2MP-Id    : 8193
Root-Addr      : 192.0.2.3 (LF)
InnerRoot-Addr : 192.0.2.1
-------------------------------------------------------------------------------
Peer           : 192.0.2.4:0
Ing Lbl        :   --
Egr Lbl        : 524282
Egr Int/LspId  : 1/1/1
EgrNextHop     : 192.168.34.2
Egr. Flags     : None               Ing. Flags : None
Egr If Name    : int-P-3-PE-4
Metric         : 1                  Mtu        : 8986
-------------------------------------------------------------------------------
P2MP Type      : 7                  P2MP-Id    : 8193
Root-Addr      : 192.168.23.1 (UF)
InnerRoot-Addr : 192.0.2.1
-------------------------------------------------------------------------------
Peer           : 192.168.23.1:0
Ing Lbl        : 524281U
Egr Lbl        :   --
Egr Int/LspId  :   --
EgrNextHop     :   --
Egr. Flags     : None               Ing. Flags : None
===============================================================================
No. of GRT Recursive with Generic IPv4 P2MP Bindings: 2
===============================================================================
```

On P-2, there is one GRT recursive mLDP binding with PE-1 as inner root and a non-recursive mLDP binding with root PE-1, as follows:

```
[/]
A:admin@P-2# show router ldp bindings active p2mp summary ipv4
 No. of Generic IPv4 P2MP Active Bindings: 1
 No. of In-Band-SSM IPv4 P2MP Active Bindings: 0
 No. of In-Band-VPN-SSM IPv4 P2MP Active Bindings: 0
 No. of In-Band-SSM IPv4 P2MP Active Bindings: 0
 No. of VPN Recursive with Generic IPv4 P2MP Active Bindings: 0
 No. of GRT Recursive with Generic IPv4 P2MP Active Bindings: 1
```

On P-2, the following GRT recursive mLDP binding with PE-1 as inner root has LF 192.168.23.1, which is an interface address of P-2. The peer is 192.168.23.2, which is an interface address of P-3.

```
[/]
A:admin@P-2# show router ldp bindings p2mp opaque-type grt-recursive ipv4 detail

===============================================================================
LDP Bindings (IPv4 LSR ID 192.0.2.2)
             (IPv6 LSR ID ::)
===============================================================================
Label Status:
        U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
        e - Label ELC
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
FEC Flags:
        LF - Lower FEC, UF - Upper FEC, M - Community Mismatch,
        BA - ASBR Backup FEC
===============================================================================
LDP GRT Recursive with Generic IPv4 P2MP Bindings
===============================================================================
-------------------------------------------------------------------------------
P2MP Type      : 7                  P2MP-Id    : 8193
Root-Addr      : 192.168.23.1 (LF)
InnerRoot-Addr : 192.0.2.1
-------------------------------------------------------------------------------
Peer           : 192.168.23.2:0
Ing Lbl        :   --
Egr Lbl        : 524281
Egr Int/LspId  : 1/1/1
EgrNextHop     : 192.168.23.2
Egr. Flags     : None               Ing. Flags : None
Egr If Name    : int-P-2-P-3
Metric         : 1                  Mtu        : 8986
===============================================================================
No. of GRT Recursive with Generic IPv4 P2MP Bindings: 1
===============================================================================
```

On P-2, the following non-recursive mLDP binding to root PE-1 has root address 192.0.2.1 as UF:

```
[/]
A:admin@P-2# show router ldp bindings p2mp opaque-type generic ipv4 detail

===============================================================================
LDP Bindings (IPv4 LSR ID 192.0.2.2)
             (IPv6 LSR ID ::)
===============================================================================
Label Status:
        U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
        e - Label ELC
FEC Flags:
        LF - Lower FEC, UF - Upper FEC, M - Community Mismatch,
        BA - ASBR Backup FEC
===============================================================================
LDP Generic IPv4 P2MP Bindings
===============================================================================
-------------------------------------------------------------------------------
P2MP Type      : 1                  P2MP-Id    : 8193
Root-Addr      : 192.0.2.1 (UF)
-------------------------------------------------------------------------------
Peer           : 192.0.2.1:0
Ing Lbl        : 524281U
Egr Lbl        :   --
Egr Int/LspId  :   --
EgrNextHop     :   --
Egr. Flags     : None               Ing. Flags : None
===============================================================================
No. of Generic IPv4 P2MP Bindings: 1
===============================================================================
```

On PE-1, there is only a non-recursive mLDP binding with root PE-1, as follows:

```
[/]
A:admin@PE-1# show router ldp bindings active p2mp summary ipv4
 No. of Generic IPv4 P2MP Active Bindings: 1
 No. of In-Band-SSM IPv4 P2MP Active Bindings: 0
 No. of In-Band-VPN-SSM IPv4 P2MP Active Bindings: 0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
 No. of In-Band-SSM IPv4 P2MP Active Bindings: 0
 No. of VPN Recursive with Generic IPv4 P2MP Active Bindings: 0
 No. of GRT Recursive with Generic IPv4 P2MP Active Bindings: 0
```

On PE-1, the following non-recursive mLDP binding with root PE-1 has peer 192.0.2.2 (P-2):

```
[/]
A:admin@PE-1# show router ldp bindings p2mp opaque-type generic ipv4 detail
===============================================================================
LDP Bindings (IPv4 LSR ID 192.0.2.1)
             (IPv6 LSR ID ::)
===============================================================================
Label Status:
        U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
        e - Label ELC
FEC Flags:
        LF - Lower FEC, UF - Upper FEC, M - Community Mismatch,
        BA - ASBR Backup FEC
===============================================================================
LDP Generic IPv4 P2MP Bindings
===============================================================================
-------------------------------------------------------------------------------
P2MP Type      : 1                 P2MP-Id    : 8193
Root-Addr      : 192.0.2.1
-------------------------------------------------------------------------------
Peer           : 192.0.2.2:0
Ing Lbl        :   --
Egr Lbl        : 524281
Egr Int/LspId  : 1/1/1
EgrNextHop     : 192.168.12.2
Egr. Flags     : None              Ing. Flags : None
Egr If Name    : int-PE-1-P-2
Metric         : 1                 Mtu        : 8986
===============================================================================
No. of Generic IPv4 P2MP Bindings: 1
===============================================================================
```

Figure 208: Recursive mLDP FEC for inter-AS model C shows the mLDP label mapping messages with the corresponding labels: label 524281 is used between PE-1 and P-2; label 524281 is used between P-2 and P-3; label 524282 is used between P-3 and PE-4.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 208: Recursive mLDP FEC for inter-AS model C

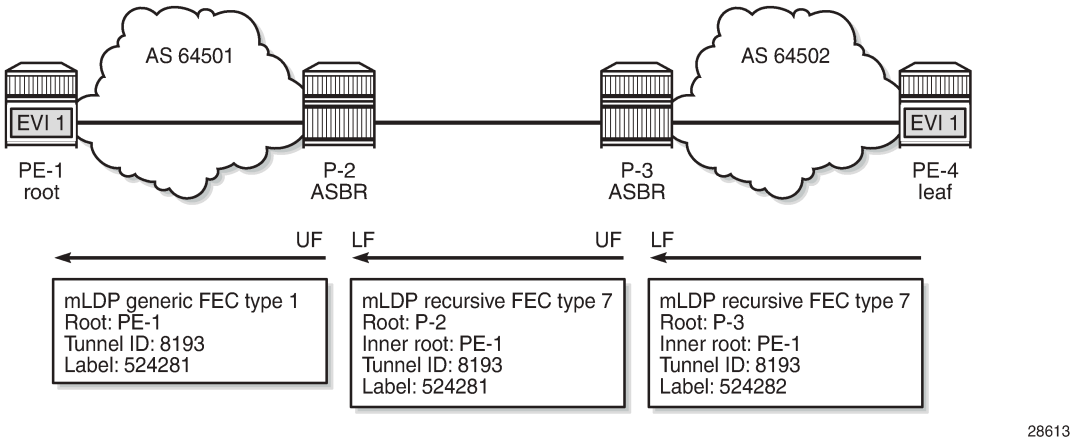

Non-recursive mLDP FEC resolution for inter-AS model C

Some routers may not support GRT recursive FEC type 7. In that case, the router generates a non-recursive FEC type 1 with root PE-1 to the next-hop P-3. In this example, leaf node PE-4 does not support GRT recursive FEC type 7 and is configured to only send basic FEC type 1 messages. ASBR P-3 supports GRT recursive type 7 and sends similar messages as in the preceding scenario. However, it is possible that none of the routers supports GRT recursive FEC type 7. In that case, the **generate-basic-fec-only** command is configured on all nodes.

The following command is configured on leaf node PE-4 to make the system send only basic FEC type 1 messages:

```
# on PE-4:
configure {
    router "Base" {
        ldp {
            generate-basic-fec-only true
```

When PE-4 is configured to only generate basic FEC type 1, PE-4 withdraws the GRT recursive type 7 (T:7) label mapping message with PE-1 as inner root and P-3 as root and sends a non-recursive generic type 1 (T:1) label mapping message with PE-1 as root instead. When debugging is enabled on PE-4 for LDP label mapping messages between P-3 and PE-4, the following messages are logged:

```
# on PE-4:
1 2021/06/17 08:44:52.342 UTC MINOR: DEBUG #2001 Base LDP
"LDP: LDP
Send Label Withdraw packet (msgId 96) to 192.0.2.3:0
Protocol version = 1
Label 524282 withdrawn for the following FECs
P2MP: root = 192.0.2.3, T: 7, L: 17 (InnerRoot: 192.0.2.1 T: 1, L: 4, TunnelId: 8193)
"

2 2021/06/17 08:44:52.342 UTC MINOR: DEBUG #2001 Base LDP
"LDP: LDP
Send Label Mapping packet (msgId 97) to 192.0.2.3:0
Protocol version = 1
Label 524281 advertised for the following FECs
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
P2MP: root = 192.0.2.1, T: 1, L: 4, TunnelId: 8193
"

3 2021/06/17 08:44:52.344 UTC MINOR: DEBUG #2001 Base LDP
"LDP: LDP
Recv Label Release packet (msgId 95) from 192.0.2.3:0
Protocol version = 1
Label 524282 released for the following FECs
P2MP: root = 192.0.2.3, T: 7, L: 17 (InnerRoot: 192.0.2.1 T: 1, L: 4, TunnelId: 8193)
"
```

On PE-4, there is one non-recursive generic mLDP binding, as follows:

```
[/]
A:admin@PE-4# show router ldp bindings p2mp summary ipv4
 No. of Generic IPv4 P2MP Bindings: 1
 No. of In-Band-SSM IPv4 P2MP Bindings: 0
 No. of In-Band-VPN-SSM IPv4 P2MP Bindings: 0
 No. of Recursive with In-Band-SSM IPv4 P2MP Bindings: 0
 No. of VPN Recursive with Generic IPv4 P2MP Bindings: 0
 No. of GRT Recursive with Generic IPv4 P2MP Bindings: 0
```

On PE-4, the following non-recursive generic mLDP binding has root PE-1 and peer P-3:

```
[/]
A:admin@PE-4# show router ldp bindings p2mp opaque-type generic detail ipv4

===============================================================================
LDP Bindings (IPv4 LSR ID 192.0.2.4)
             (IPv6 LSR ID ::)
===============================================================================
Label Status:
        U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
        e - Label ELC
FEC Flags:
        LF - Lower FEC, UF - Upper FEC, M - Community Mismatch,
        BA - ASBR Backup FEC
===============================================================================
LDP Generic IPv4 P2MP Bindings
===============================================================================
-------------------------------------------------------------------------------
P2MP Type      : 1                  P2MP-Id    : 8193
Root-Addr      : 192.0.2.1
-------------------------------------------------------------------------------
Peer           : 192.0.2.3:0
Ing Lbl        : 524281U
Egr Lbl        :   --
Egr Int/LspId  :   --
EgrNextHop     :   --
Egr. Flags     : None               Ing. Flags : None
===============================================================================
No. of Generic IPv4 P2MP Bindings: 1
===============================================================================
```

On P-3, there is one generic mLDP binding and one recursive mLDP binding, as follows:

```
[/]
A:admin@P-3# show router ldp bindings p2mp summary ipv4
 No. of Generic IPv4 P2MP Bindings: 1
 No. of In-Band-SSM IPv4 P2MP Bindings: 0
 No. of In-Band-VPN-SSM IPv4 P2MP Bindings: 0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
 No. of Recursive with In-Band-SSM IPv4 P2MP Bindings: 0
 No. of VPN Recursive with Generic IPv4 P2MP Bindings: 0
 No. of GRT Recursive with Generic IPv4 P2MP Bindings: 1
```

```
[/]
A:admin@P-3# show router ldp bindings p2mp opaque-type generic detail ipv4

===============================================================================
LDP Bindings (IPv4 LSR ID 192.0.2.3)
             (IPv6 LSR ID ::)
===============================================================================
Label Status:
        U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
        e - Label ELC
FEC Flags:
        LF - Lower FEC, UF - Upper FEC, M - Community Mismatch,
        BA - ASBR Backup FEC
===============================================================================
LDP Generic IPv4 P2MP Bindings
===============================================================================
-------------------------------------------------------------------------------
P2MP Type      : 1                  P2MP-Id    : 8193
Root-Addr      : 192.0.2.1 (LF)
-------------------------------------------------------------------------------
Peer           : 192.0.2.4:0
Ing Lbl        :   --
Egr Lbl        : 524281
Egr Int/LspId  : 1/1/1
EgrNextHop     : 192.168.34.2
Egr. Flags     : None               Ing. Flags : None
Egr If Name    : int-P-3-PE-4
Metric         : 1                  Mtu        : 8986
===============================================================================
No. of Generic IPv4 P2MP Bindings: 1
===============================================================================
```

```
[/]
A:admin@P-3# show router ldp bindings p2mp opaque-type grt-recursive detail ipv4

===============================================================================
LDP Bindings (IPv4 LSR ID 192.0.2.3)
             (IPv6 LSR ID ::)
===============================================================================
Label Status:
        U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
        e - Label ELC
FEC Flags:
        LF - Lower FEC, UF - Upper FEC, M - Community Mismatch,
        BA - ASBR Backup FEC
===============================================================================
LDP GRT Recursive with Generic IPv4 P2MP Bindings
===============================================================================
-------------------------------------------------------------------------------
P2MP Type      : 7                  P2MP-Id    : 8193
Root-Addr      : 192.168.23.1 (UF)
InnerRoot-Addr : 192.0.2.1
-------------------------------------------------------------------------------
Peer           : 192.168.23.1:0
Ing Lbl        : 524280U
Egr Lbl        :   --
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Egr Int/LspId  :   --
EgrNextHop     :   --
Egr. Flags     : None              Ing. Flags : None
===============================================================================
No. of GRT Recursive with Generic IPv4 P2MP Bindings: 1
===============================================================================
```

On P-2 and PE-1, the mLDP bindings are similar to the preceding scenario, but the labels are different. Figure 209: Non-recursive mLDP FEC for inter-AS model C shows the label mapping messages with label 524280 between PE-1 and P-2 and label 524280 between P-2 and P-3; label 524281 is used between P-3 and PE-4.

Figure 209: Non-recursive mLDP FEC for inter-AS model C

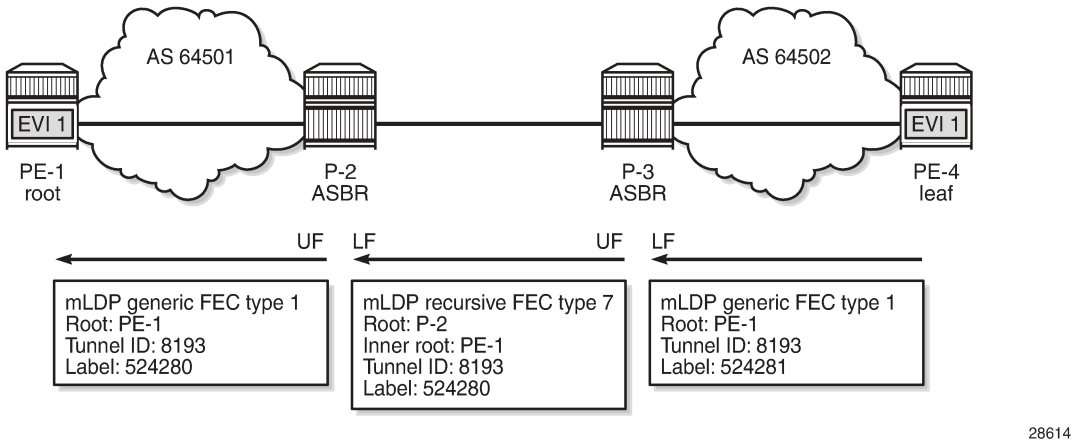

Seamless MPLS

Figure 210: Example topology for seamless MPLS shows the example topology for seamless MPLS.

Figure 210: Example topology for seamless MPLS

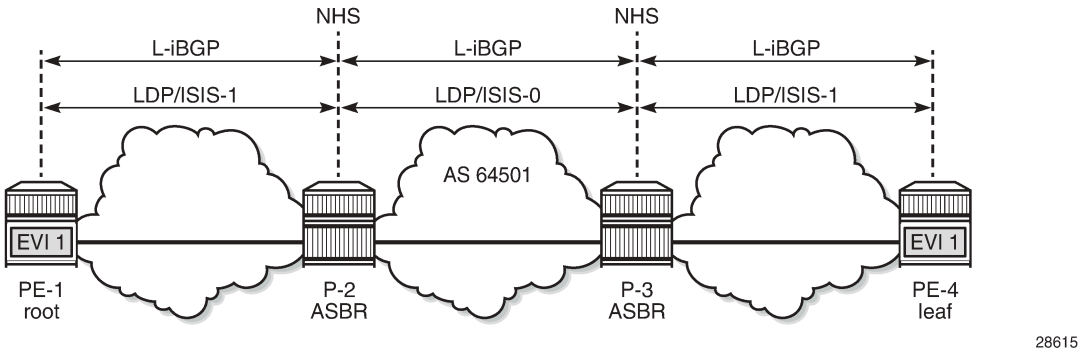

The configuration is according to the *Seamless MPLS: Isolated IGP/LDP Domains and Labeled BGP* chapter in the MPLS volume of *7450 ESS, 7750 SR, and 7950 XRS Advanced Configuration Guide — Part I*.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

IS-IS is configured as IGP. IS-IS instance 0 is configured between P-2 and P-3, whereas IS-IS instance 1 is configured between P-2 and PE-1 and between P-3 and PE-4. On P-2, IS-IS is configured as follows:

```
# on P-2:
configure {
    router "Base" {
        isis 0 {
            admin-state enable
            level-capability 2
            area-address [49.0001]
            interface "int-P-2-P-3" {
                interface-type point-to-point
            }
            interface "system" {
            }
        }
        isis 1 {
            admin-state enable
            level-capability 2
            area-address [49.0001]
            interface "int-P-2-PE-1" {
                interface-type point-to-point
            }
            interface "system" {
            }
        }
```

Other characteristics of this example are as follows:

- Unlike the preceding use case for inter-AS model C, no static route is required between P-2 and P-3.
- LDP is configured on all interfaces.
- VPLS 1 is configured as before, but the route target is identical for import and export, and equal to 64501:1.
- All nodes are in AS 64501, so only iBGP is configured.

On PE-1, BGP is configured as follows, using the same policy as for inter-AS model C. The BGP configuration on PE-4 is similar, but the neighbors are different.

```
# on PE-1:
configure {
    router "Base" {
        autonomous-system 64501
        bgp {
            split-horizon true
            group "iBGP" {
                type internal
            }
            neighbor "192.0.2.2" {
                group "iBGP"
                family {
                    label-ipv4 true
                }
                export {
                    policy ["PE-sys-to-labeled-BGP"]
                }
            }
            neighbor "192.0.2.4" {
                group "iBGP"
                family {
                    evpn true
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
            }
        }
```

On P-2, the BGP configuration is as follows. The ABRs are configured with **next-hop-self** in both directions. The BGP configuration is similar on P-3.

```
# on P-2:
configure {
    router "Base" {
        autonomous-system 64501
        bgp {
            split-horizon true
            group "iBGP" {
                type internal
            }
            neighbor "192.0.2.1" {
                group "iBGP"
                next-hop-self true
                family {
                    label-ipv4 true
                }
                cluster {
                    cluster-id 192.0.2.2
                }
            }
            neighbor "192.0.2.3" {
                advertise-inactive true
                group "iBGP"
                next-hop-self true
                family {
                    label-ipv4 true
                }
            }
        }
```

The following route table on PE-4 shows a labeled BGP route to root node PE-1 with P-3 as the next-hop:

```
[/]
A:admin@PE-4# show router route-table 192.0.2.1

===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
192.0.2.1/32                                  Remote  BGP_LABEL 00h00m47s  170
       192.0.2.3 (tunneled)                                       10
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

Likewise, P-3 has a labeled BGP route to root node PE-1 with P-2 as the next-hop, as follows:

```
[/]
A:admin@P-3# show router route-table 192.0.2.1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
192.0.2.1/32                                  Remote  BGP_LABEL 00h01m00s  170
       192.0.2.2 (tunneled)                                       10
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

P-2 has an IS-IS route to PE-1, using IS-IS instance 1, as follows:

```
[/]
A:admin@P-2# show router route-table 192.0.2.1

===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
192.0.2.1/32                                  Remote  ISIS(1)   00h01m51s  18
       192.168.12.1                                               10
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

Recursive mLDP FEC resolution for seamless MPLS

When the leaf node PE-4 supports GRT recursive FEC type 7, it generates one GRT recursive FEC label mapping message with PE-1 as inner root and P-3 as root, as follows:

```
[/]
A:admin@PE-4# show router ldp bindings p2mp summary ipv4
 No. of Generic IPv4 P2MP Bindings: 0
 No. of In-Band-SSM IPv4 P2MP Bindings: 0
 No. of In-Band-VPN-SSM IPv4 P2MP Bindings: 0
 No. of Recursive with In-Band-SSM IPv4 P2MP Bindings: 0
 No. of VPN Recursive with Generic IPv4 P2MP Bindings: 0
 No. of GRT Recursive with Generic IPv4 P2MP Bindings: 1
```

```
[/]
A:admin@PE-4# show router ldp bindings p2mp opaque-type grt-recursive detail ipv4

===============================================================================
LDP Bindings (IPv4 LSR ID 192.0.2.4)
             (IPv6 LSR ID ::)
===============================================================================
Label Status:
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
        e - Label ELC
FEC Flags:
        LF - Lower FEC, UF - Upper FEC, M - Community Mismatch,
        BA - ASBR Backup FEC
===============================================================================
LDP GRT Recursive with Generic IPv4 P2MP Bindings
===============================================================================
-------------------------------------------------------------------------------
P2MP Type      : 7                  P2MP-Id    : 8193
Root-Addr      : 192.0.2.3
InnerRoot-Addr : 192.0.2.1
-------------------------------------------------------------------------------
Peer           : 192.0.2.3:0
Ing Lbl        : 524281U
Egr Lbl        :   --
Egr Int/LspId  :   --
EgrNextHop     :   --
Egr. Flags     : None               Ing. Flags : None
===============================================================================
No. of GRT Recursive with Generic IPv4 P2MP Bindings: 1
===============================================================================
```

P-3 has two GRT recursive FEC bindings with inner root 192.0.2.1: one with UF 192.0.2.2 and another with LF 192.0.2.3, as follows:

```
[/]
A:admin@P-3# show router ldp bindings p2mp opaque-type grt-recursive detail ipv4

===============================================================================
LDP Bindings (IPv4 LSR ID 192.0.2.3)
             (IPv6 LSR ID ::)
===============================================================================
Label Status:
        U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
        e - Label ELC
FEC Flags:
        LF - Lower FEC, UF - Upper FEC, M - Community Mismatch,
        BA - ASBR Backup FEC
===============================================================================
LDP GRT Recursive with Generic IPv4 P2MP Bindings
===============================================================================
-------------------------------------------------------------------------------
P2MP Type      : 7                  P2MP-Id    : 8193
Root-Addr      : 192.0.2.2 (UF)
InnerRoot-Addr : 192.0.2.1
-------------------------------------------------------------------------------
Peer           : 192.0.2.2:0
Ing Lbl        : 524281U
Egr Lbl        :   --
Egr Int/LspId  :   --
EgrNextHop     :   --
Egr. Flags     : None               Ing. Flags : None
-------------------------------------------------------------------------------
P2MP Type      : 7                  P2MP-Id    : 8193
Root-Addr      : 192.0.2.3 (LF)
InnerRoot-Addr : 192.0.2.1
-------------------------------------------------------------------------------
Peer           : 192.0.2.4:0
Ing Lbl        :   --
Egr Lbl        : 524281
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Egr Int/LspId  : 1/1/1
EgrNextHop     : 192.168.34.2
Egr. Flags     : None               Ing. Flags : None
Egr If Name    : int-P-3-PE-4
Metric         : 1                  Mtu        : 8986
===============================================================================
No. of GRT Recursive with Generic IPv4 P2MP Bindings: 2
===============================================================================
```

P-2 has one GRT recursive FEC binding with inner root PE-1 and root P-2 (LF). P-2 also has one non-recursive FEC binding with root PE-1 (UF). PE-1 only has a non-recursive FEC binding with root PE-1. Figure 211: Recursive mLDP FEC for seamless MPLS shows the mLDP label mapping messages that all have label 524281 in this example.

Figure 211: Recursive mLDP FEC for seamless MPLS

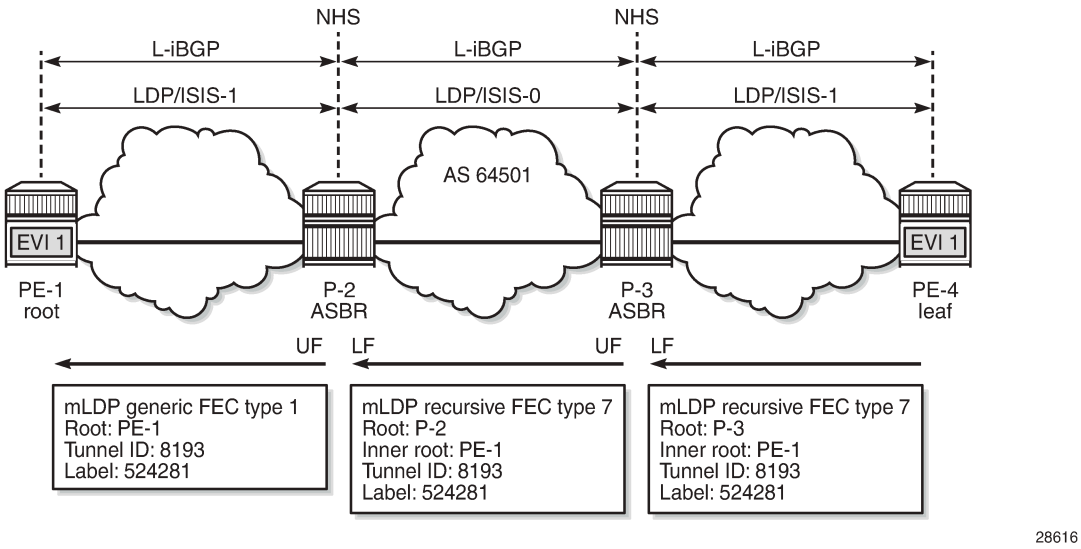

Non-recursive mLDP FEC resolution for seamless MPLS

For nodes that do not support GRT recursive mLDP FEC type 7, the following command ensures that only non-recursive mLDP type 1 label mapping messages will be sent. In this example, it is assumed that only PE-4 does not support GRT recursive mLDP FEC type 7.

```
# on PE-4:
configure {
    router "Base" {
        ldp {
            generate-basic-fec-only true
```

PE-4 sends a non-recursive mLDP label mapping message with PE-1 as the root to its peer P-3, as follows:

```
[/]
A:admin@PE-4# show router ldp bindings p2mp opaque-type generic detail ipv4

===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
LDP Bindings (IPv4 LSR ID 192.0.2.4)
             (IPv6 LSR ID ::)
===============================================================================
Label Status:
        U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
        e - Label ELC
FEC Flags:
        LF - Lower FEC, UF - Upper FEC, M - Community Mismatch,
        BA - ASBR Backup FEC
===============================================================================
LDP Generic IPv4 P2MP Bindings
===============================================================================
-------------------------------------------------------------------------------
P2MP Type      : 1                  P2MP-Id    : 8193
Root-Addr      : 192.0.2.1
-------------------------------------------------------------------------------
Peer           : 192.0.2.3:0
Ing Lbl        : 524280U
Egr Lbl        :   --
Egr Int/LspId  :   --
EgrNextHop     :   --
Egr. Flags     : None               Ing. Flags : None
===============================================================================
No. of Generic IPv4 P2MP Bindings: 1
===============================================================================
```

Figure 212: Leaf node sends basic FEC in seamless MPLS shows the label mapping messages when leaf node PE-4 only generates basic FEC type 1 messages.

Figure 212: Leaf node sends basic FEC in seamless MPLS

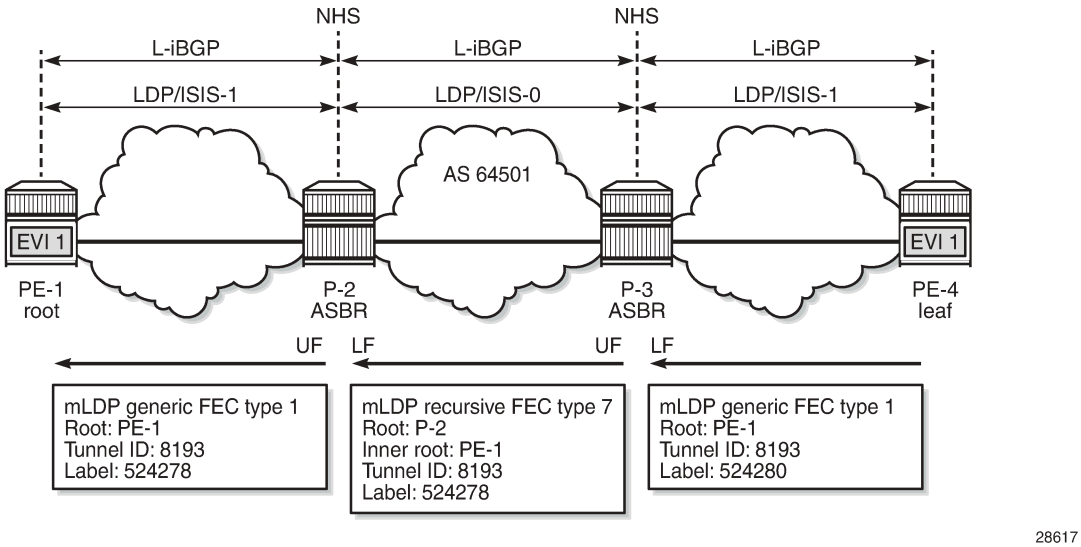

It is possible that ABR routers do not support GRT recursive either. The same command is configured on P-2 and P-3, as follows:

```
# on P-2, P-3:
configure {
    router "Base" {
        ldp {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            generate-basic-fec-only true
```

When **generate-basic-fec-only** is enabled in the ABRs, P-2 and P-3 will only generate basic FEC messages. On P-3, there are no GRT recursive mLDP bindings anymore, as follows:

```
[/]
A:admin@P-3# show router ldp bindings p2mp summary ipv4
 No. of Generic IPv4 P2MP Bindings: 2
 No. of In-Band-SSM IPv4 P2MP Bindings: 0
 No. of In-Band-VPN-SSM IPv4 P2MP Bindings: 0
 No. of Recursive with In-Band-SSM IPv4 P2MP Bindings: 0
 No. of VPN Recursive with Generic IPv4 P2MP Bindings: 0
 No. of GRT Recursive with Generic IPv4 P2MP Bindings: 0
```

The two generic mLDP bindings on P-3 have root address 192.0.2.1, as follows. There is no UF or LF.

```
[/]
A:admin@P-3# show router ldp bindings p2mp opaque-type generic detail ipv4

===============================================================================
LDP Bindings (IPv4 LSR ID 192.0.2.3)
             (IPv6 LSR ID ::)
===============================================================================
Label Status:
        U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
        e - Label ELC
FEC Flags:
        LF - Lower FEC, UF - Upper FEC, M - Community Mismatch,
        BA - ASBR Backup FEC
===============================================================================
LDP Generic IPv4 P2MP Bindings
===============================================================================
-------------------------------------------------------------------------------
P2MP Type      : 1                  P2MP-Id    : 8193
Root-Addr      : 192.0.2.1
-------------------------------------------------------------------------------
Peer           : 192.0.2.2:0
Ing Lbl        : 524281U
Egr Lbl        :   --
Egr Int/LspId  :   --
EgrNextHop     :   --
Egr. Flags     : None               Ing. Flags : None
-------------------------------------------------------------------------------
P2MP Type      : 1                  P2MP-Id    : 8193
Root-Addr      : 192.0.2.1
-------------------------------------------------------------------------------
Peer           : 192.0.2.4:0
Ing Lbl        :   --
Egr Lbl        : 524280
Egr Int/LspId  : 1/1/1
EgrNextHop     : 192.168.34.2
Egr. Flags     : None               Ing. Flags : None
Egr If Name    : int-P-3-PE-4
Metric         : 1                  Mtu        : 8986
===============================================================================
No. of Generic IPv4 P2MP Bindings: 2
===============================================================================
```

The output on P-2 is similar. Figure 213: ABRs and leaf node send basic FEC in seamless MPLS shows the label mapping messages when all nodes only generate basic FEC type 1 messages.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 213: ABRs and leaf node send basic FEC in seamless MPLS

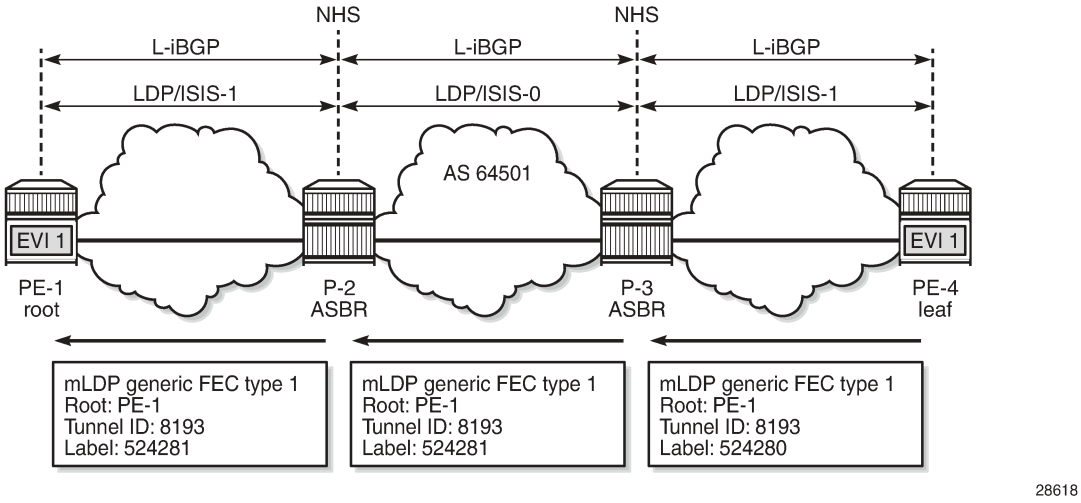

Conclusion

In inter-AS and intra-AS scenarios, mLDP trees can be set up using recursive or non-recursive label mapping messages. Routers not supporting recursive FEC can generate only non-recursive FEC, even if the system address of the root node is resolved via BGP. This feature is supported in MVPN and in EVPN.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

P2MP mLDP Inter-AS Model C for EVPN-MPLS Services

This chapter provides information about P2MP mLDP Inter-AS Model C for EVPN-MPLS Services.

Topics in this chapter include:

Applicability

This chapter was initially written for SR OS Release 15.0.R5, but the MD-CLI in the current edition is based on SR OS Release 21.5.R1.

Point-to-Multipoint Multicast Label Distribution Protocol (P2MP mLDP) for Broadcast, Unknown Unicast, and Multicast (BUM) traffic in EVPN-MPLS networks is supported in SR OS Release 14.0.R1, and later. EVPN with P2MP mLDP LSPs is supported in a seamless MPLS or inter-AS model C scenario in SR OS Release 15.0.R1, and later. This chapter describes the inter-AS model C scenario, but the configuration for seamless MPLS is similar.

Overview

Chapter "P2MP mLDP Tunnels for BUM Traffic in EVPN-MPLS Services" in the Layer 2 Services and EVPN volume in the *7450 ESS, 7750 SR, and 7950 XRS Advanced Configuration Guide — Part II* describes P2MP mLDP within an Autonomous System (AS). PEs configured with **root-and-leaf true** can send BUM traffic over P2MP mLDP tunnels; PEs configured with root-and-leaf false (that is, leaf-only) can only send BUM traffic over Ingress Replication (IR) tunnels. Both types of PEs (root-and-leaf and leaf-only) can receive BUM traffic over either P2MP mLDP tunnels or IR tunnels.

When **provider-tunnel inclusive mldp** is enabled in an EVPN-MPLS service in combination with **root-and-leaf true** and **bgp-evpn>mpls>ingress-replication-bum-label true**, the system will send an Inclusive Multicast Ethernet Tag (IMET) route with a composite tunnel type (IMET-P2MP-IR) in the provider tunnel attributed.

Inter-AS VPN model C is described in chapters Inter-AS VPRN Model C and Inter-AS Model C for VLL. Labeled IPv4 unicast BGP is used to provide inter-AS connectivity. The system IP addresses within each AS are exported by the Autonomous System Border Routers (ASBRs) and a multi-hop BGP session is established between root node and leaf node for address family EVPN. The root node advertises a composite IMET-P2MP-IR route to the leaf nodes and the leaf nodes advertise an IMET-IR route to the root node. Figure 214: Inter-AS Model C for P2MP mLDP shows an example topology with root node PE-1 in AS 64501 and leaf node PE-4 in AS 64502. P-2 and P-3 are ASBRs.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 214: Inter-AS Model C for P2MP mLDP

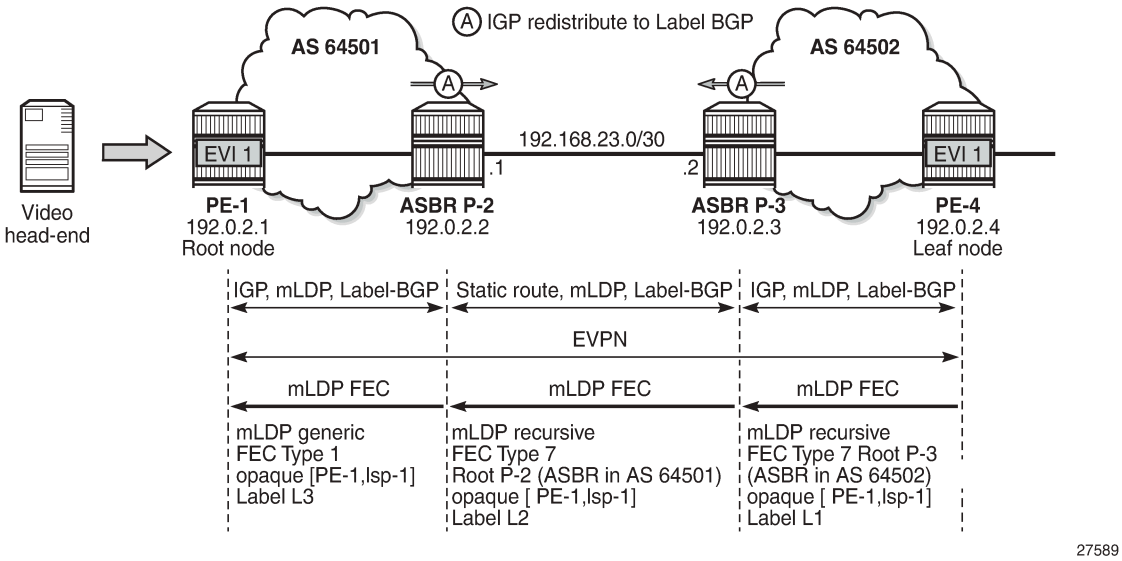

The composite IMET-P2MP-IR route received by leaf node PE-4 contains the root node (192.0.2.1) and the LSP ID (0x2001) that will be used by the nodes to set up a P2MP mLDP tree toward the root.

```
# on PE-4:
3 2021/06/02 08:31:13.913 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.1
"Peer 1: 192.0.2.1: UPDATE
Peer 1: 192.0.2.1 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 92
    Flag: 0x90 Type: 14 Len: 28 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.1
        Type: EVPN-INCL-MCAST Len: 17 RD: 192.0.2.1:1, tag: 0, orig_addr len: 32,
                              orig_addr: 192.0.2.1
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 6 AS Path:
        Type: 2 Len: 1 < 64501 >
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64501:1
        bgp-tunnel-encap:MPLS
    Flag: 0xc0 Type: 22 Len: 25 PMSI:
        Tunnel-type Composite LDP P2MP IR (130)
        Flags: (0x0)[Type: None BM: 0 U: 0 Leaf: not required]
        MPLS Label1 Ag  0
        MPLS Label2 IR 8388544
        Root-Node 192.0.2.1, LSP-ID 0x2001
"
```

The Provider Multicast Service Interface (PMSI) tunnel attribute for tunnel type 130 (composite tunnel) has two MPLS labels, of which MPLS label 1 always equals zero in SR OS Release 21.5.R1, because SR OS does not support aggregated P2MP tunnels. MPLS label 2 is used by the downstream nodes to set up the EVPN-MPLS destination to the root node and add it to the default multicast list. The actual MPLS label only uses the high-order 20 bits out of the 24 bits advertised in the MPLS label. Therefore, the value 8388544 needs to be divided by 16 to get the MPLS label value: 8388544/16 = 524284. This is due to the debug message being shown before the router can parse the label field and see whether it corresponds to an

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

MPLS label (20 bits) or a VXLAN VNI (24 bits). The following command on PE-4 shows the EVPN-MPLS destination 192.0.2.1 with MPLS label 524284 using a BGP transport tunnel:

```
[/]
A:admin@PE-4# show service id 1 evpn-mpls

===============================================================================
BGP EVPN-MPLS Dest
===============================================================================
TEP Address     Egr Label     Num. MACs   Mcast           Last Change
                Transport:Tnl                             Sup BCast Domain
-------------------------------------------------------------------------------
192.0.2.1       524284        0           bum             06/02/2021 08:31:14
                bgp:262146                                No
-------------------------------------------------------------------------------
Number of entries : 1
-------------------------------------------------------------------------------
===============================================================================
---snip---
```

The use of mLDP with recursive opaque values is specified in RFC 6512.

When the leaf node PE-4 receives the composite IMET-P2MP-IR route from the root node PE-1, a P2MP mLDP tree needs to be established from the leaf node to the root node. Leaf node PE-4 resolves the IP address of PE-1 to a labeled BGP route with next-hop ASBR P-3. PE-4 then sends an mLDP FEC with root node ASBR P-3 and an opaque value containing the root PE-1 and an LSP ID that was advertised in the IMET-P2MP-IR route, as follows:

```
# on PE-4:
4 2021/06/02 08:31:13.915 UTC MINOR: DEBUG #2001 Base LDP
"LDP: LDP
Send Label Mapping packet (msgId 40) to 192.0.2.3:0
Protocol version = 1
Label 524283 advertised for the following FECs
P2MP: root = 192.0.2.3, T: 7, L: 17 (InnerRoot: 192.0.2.1 T: 1, L: 4, TunnelId: 8193)
"
```

T: 7 indicates the MLDP recursive FEC type 7. The tunnel ID 8193 corresponds to the hexadecimal value 0x2001 sent by the root node PE-1, which is the inner root 192.0.2.1 in the recursive opaque value.

When ASBR P-3 receives this mLDP FEC, it identifies itself as root node and resolves the recursive opaque value (PE-1, LSP ID) and creates a new mLDP FEC element with root node ASBR P-2 and an identical opaque value (PE-1, LSP ID). The following mLDP FEC is sent to ASBR P-2:

```
# on P-3:
4 2021/06/02 08:32:36.794 UTC MINOR: DEBUG #2001 Base LDP
"LDP: LDP
Send Label Mapping packet (msgId 34) to 192.168.23.1:0
Protocol version = 1
Label 524279 advertised for the following FECs
P2MP: root = 192.168.23.1, T: 7, L: 17 (InnerRoot: 192.0.2.1 T: 1, L: 4, TunnelId: 8193)
"
```

ASBR P-2 receives the mLDP FEC and finds that it is the root node. P-2 creates a new mLDP FEC, but no recursion is required because P-2 knows the IP address of PE-1 through the IGP. P-2 sends the following mLDP FEC with root node PE-1, LSP ID 8193, and mLDP FEC type 1.

```
# on P-2:
6 2021/06/02 08:32:36.814 UTC MINOR: DEBUG #2001 Base LDP
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
"LDP: LDP
Send Label Mapping packet (msgId 50) to 192.0.2.1:0
Protocol version = 1
Label 524279 advertised for the following FECs
P2MP: root = 192.0.2.1, T: 1, L: 4, TunnelId: 8193
"
```

Configuration

The example topology was already shown in Figure 214: Inter-AS Model C for P2MP mLDP. The initial configuration includes the following:

- Cards, MDAs, ports
- Router interfaces
- OSPF as IGP within each AS (alternatively, IS-IS can be used)
- LDP enabled within each AS

The following two scenarios are configured:

- Inter-AS model C for mLDP
- Optimized inter-AS model C for mLDP

Inter-AS Model C for mLDP

The initial BGP configuration on the PEs only includes a label-IPv4 peering with the ASBRs. The BGP configuration on PE-1 is as follows:

```
# on PE-1:
configure {
    router "Base" {
        bgp {
            split-horizon true
            group "iBGP" {
                type internal
            }
            neighbor "192.0.2.2" {
                group "iBGP"
                family {
                    label-ipv4 true
                }
            }
```

On the ASBRs, BGP is configured for address family label-IPv4, both internal to PE-1 and external to the peer ASBR. A policy exports the system prefixes as label-IPv4 routes to the eBGP peer. The BGP configuration on P-2 is as follows:

```
# on P-2:
configure {
    policy-options {
        prefix-list "sysPE" {
            prefix 192.0.2.0/24 type longer {
            }
        }
        policy-statement "PE-sys-to-labeled-BGP" {
            entry 10 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                from {
                    prefix-list ["sysPE"]
                }
                to {
                    protocol {
                        name [bgp-label]
                    }
                }
                action {
                    action-type accept
                }
            }
        }
    }
    router "Base" {
        bgp {
            ebgp-default-reject-policy {
                import false
            }
            group "eBGP" {
                type external
            }
            group "iBGP" {
                type internal
            }
            neighbor "192.0.2.1" {
                group "iBGP"
                family {
                    label-ipv4 true
                }
            }
            neighbor "192.168.23.2" {
                split-horizon true
                group "eBGP"
                peer-as 64502
                family {
                    label-ipv4 true
                }
                local-as {
                    as-number 64501
                }
                export {
                    policy ["PE-sys-to-labeled-BGP"]
                }
            }
```

The BGP configuration on ASBR P-3 is similar, but the IP addresses are different and the local AS and peer AS are swapped. The export policy is identical on both ASBRs P-2 and P-3.

When a P2MP mLDP tree must be established across ASs, LDP needs to be enabled on the interface between the ASBRs with **local-lsr-id>interface-name <..>** instead of the default value "system". The LDP configuration on P-2 is as follows:

```
# on P-2:
configure {
    router "Base" {
        ldp {
            interface-parameters {
                interface "int-P-2-P-3" {
                    ipv4 {
                        local-lsr-id {
                            interface-name "int-P-2-P-3"
                        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
            }
```

With this LDP configuration, a link adjacency will be established toward the interface IP address instead of the system address, as follows:

```
[/]
A:admin@P-2# show router ldp session ipv4

===============================================================================
LDP IPv4 Sessions
===============================================================================
Peer LDP Id         Adj Type  State         Msg Sent  Msg Recv  Up Time
-------------------------------------------------------------------------------
192.0.2.1:0         Link      Established   100       101       0d 00:04:05
192.168.23.2:0      Link      Established   16        18        0d 00:00:20
-------------------------------------------------------------------------------
No. of IPv4 Sessions: 2
===============================================================================
```

However, this LDP configuration is insufficient for the resolution of mLDP FEC as link LSR ID. LDP needs a /32 route instead of a /30 route, so the following /32 static route is configured on P-2:

```
# on P-2:
configure {
    router "Base" {
        static-routes {
            route 192.168.23.2/32 route-type unicast {
                next-hop "192.168.23.2" {
                    admin-state enable
                }
            }
```

The configuration on ASBR P-3 is similar for static route 192.168.23.1/32. When this static route is not configured, no mLDP label mapping message will be sent from P-3 to P-2, so the mLDP P2MP tree cannot be established.

On PE-1, VPLS 1 is configured with mLDP root-and-leaf, as follows:

```
# on PE-1:
configure {
    service {
        vpls "EVI-1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
                route-target {
                    export "target:64501:1"
                    import "target:64502:1"
                }
            }
            bgp-evpn {
                evi 1
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    auto-bind-tunnel {
                        resolution any
                    }
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
            sap 1/2/1:1 {
            }
            provider-tunnel {
                inclusive {
                    admin-state enable
                    owner bgp-evpn-mpls
                    root-and-leaf true
                    mldp
                }
            }
        }
```

On PE-4, VPLS 1 is configured with mLDP leaf-only (no root-and-leaf, which is default), as follows:

```
# on PE-4:
configure {
    service {
        vpls "EVI-1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
                route-target {
                    export "target:64502:1"
                    import "target:64501:1"
                }
            }
            bgp-evpn {
                evi 1
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            sap 1/2/1:1 {
            }
            provider-tunnel {
                inclusive {
                    admin-state enable
                    owner bgp-evpn-mpls
                    mldp
                }
            }
        }
```

The Route Distinguisher (RD) is auto-derived from EVI 1, but the route target (RT) should not be auto-derived, because the export RT on PE-1 must match the import RT on PE-4, and vice versa. It is an option to configure an identical RT on all PEs, such as 1:1, but in this example, the export RT on PE-1 is 64501:1, which equals the import RT on PE-4. When the RTs do not match, the BGP routes will be received at the PE in the peer AS, but they will not become active and no mLDP P2MP tree can be established.

Multi-hop BGP peering is configured between PE-1 and PE-4 for address family EVPN. The external BGP configuration on PE-1 is as follows:

```
# on PE-1:
configure {
    router "Base" {
        bgp {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            split-horizon true
            ebgp-default-reject-policy {
                import false
                export false
            }
            group "eBGP" {
                multihop 10
                peer-as 64502
                family {
                    evpn true
                }
                local-as {
                    as-number 64501
                }
            }
            neighbor "192.0.2.4" {
                group "eBGP"
            }
```

The external BGP configuration on PE-4 is similar, but the local AS and peer AS are swapped, and the neighbor IP address is different.

Inter-AS Model C for mLDP - Verification

The following BGP summary shows that P-2 has sent and received two prefixes with its eBGP peer P-3 and has advertised two prefixes to its iBGP peer PE-1:

```
[/]
A:admin@P-2# show router bgp summary all

===============================================================================
BGP Summary
===============================================================================
Legend : D - Dynamic Neighbor
===============================================================================
Neighbor
Description
ServiceId          AS PktRcvd InQ  Up/Down   State|Rcv/Act/Sent (Addr Family)
                      PktSent OutQ
-------------------------------------------------------------------------------
192.0.2.1
Def. Instance  64501       11    0 00h04m03s 0/0/2 (Lbl-IPv4)
                           13    0
192.168.23.2
Def. Instance  64502       12    0 00h03m54s 2/2/2 (Lbl-IPv4)
                           12    0

-------------------------------------------------------------------------------
```

ASBR P-2 advertised the following prefixes from AS 64501 to its neighbor P-3:

```
[/]
A:admin@P-2# show router bgp neighbor 192.168.23.2 advertised-routes label-ipv4
===============================================================================
 BGP Router ID:192.0.2.2        AS:64501       Local AS:64501
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
BGP Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop (Router)                                 Path-Id     IGP Cost
      As-Path                                                      Label
-------------------------------------------------------------------------------
i     192.0.2.1/32                                     n/a         10
      192.168.23.1                                     None        n/a
      64501                                                        524284
i     192.0.2.2/32                                     n/a         None
      192.168.23.1                                     None        n/a
      64501                                                        524285
-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

ASBR P-2 received the following prefixes from AS 64502 from its neighbor P-3. Both routes are used.

```
[/]
A:admin@P-2# show router bgp neighbor 192.168.23.2 received-routes label-ipv4
===============================================================================
 BGP Router ID:192.0.2.2        AS:64501       Local AS:64501
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop (Router)                                 Path-Id     IGP Cost
      As-Path                                                      Label
-------------------------------------------------------------------------------
u*>i  192.0.2.3/32                                     n/a         None
      192.168.23.2                                     None        0
      64502                                                        524285
u*>i  192.0.2.4/32                                     n/a         10
      192.168.23.2                                     None        0
      64502                                                        524284
-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

These routes are advertised by P-2 to its iBGP neighbor PE-1, so PE-1 will have the same label-IPv4 routes. The following command shows the route table on PE-1 that includes tunneled routes to P-3 and PE-4 in AS 64502.

```
[/]
A:admin@PE-1# show router route-table

===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                             Type    Proto     Age        Pref
      Next Hop[Interface Name]                                     Metric
-------------------------------------------------------------------------------
192.0.2.1/32                                   Local   Local     00h06m21s  0
       system                                                         0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
192.0.2.2/32                                    Remote  OSPF      00h06m13s  10
       192.168.12.2                                                 10
192.0.2.3/32                                    Remote  BGP_LABEL 00h03m20s  170
       192.0.2.2 (tunneled)                                         10
192.0.2.4/32                                    Remote  BGP_LABEL 00h03m20s  170
       192.0.2.2 (tunneled)                                         10
192.168.12.0/30                                 Local   Local     00h06m21s  0
       int-PE-1-P-2                                                 0
-------------------------------------------------------------------------------
No. of Routes: 5
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

The following command shows the tunnel table on PE-1:

```
[/]
A:admin@PE-1# show router tunnel-table

===============================================================================
IPv4 Tunnel Table (Router: Base)
===============================================================================
Destination           Owner     Encap TunnelId  Pref   Nexthop        Metric
   Color
-------------------------------------------------------------------------------
127.0.128.0/32        sdp       MPLS  32767     5      127.0.128.0    0
192.0.2.2/32          ldp       MPLS  65537     9      192.168.12.2   10
192.0.2.3/32          bgp       MPLS  262145    12     192.0.2.2      1000
192.0.2.4/32          bgp       MPLS  262146    12     192.0.2.2      1000
-------------------------------------------------------------------------------
Flags: B = BGP or MPLS backup hop available
       L = Loop-Free Alternate (LFA) hop available
       E = Inactive best-external BGP route
       k = RIB-API or Forwarding Policy backup hop
===============================================================================
```

The tunnels toward P-3 and PE-4 are BGP tunnels. The SDP in the list is auto-created on the root node by mLDP. The output of these show commands on PE-4 is similar, but no SDP will be created on a leaf-only node.

The route-table on ASBR P-2 includes tunneled routes toward P-3 and PE-4 and a static route to 192.168.23.2/32, as follows:

```
[/]
A:admin@P-2# show router route-table

===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                     Metric
-------------------------------------------------------------------------------
192.0.2.1/32                                    Remote  OSPF      00h06m24s  10
       192.168.12.1                                                 10
192.0.2.2/32                                    Local   Local     00h06m25s  0
       system                                                       0
192.0.2.3/32                                    Remote  BGP_LABEL 00h03m50s  170
       192.168.23.2                                                 0
192.0.2.4/32                                    Remote  BGP_LABEL 00h03m50s  170
       192.168.23.2                                                 0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
192.168.12.0/30                                 Local   Local     00h06m25s  0
       int-P-2-PE-1                                                    0
192.168.23.0/30                                 Local   Local     00h06m25s  0
       int-P-2-P-3                                                     0
192.168.23.2/32                                 Remote  Static    00h00m28s  5
       192.168.23.2                                                    1
-------------------------------------------------------------------------------
No. of Routes: 7
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

The tunnel table on P-2 has an LDP tunnel toward PE-1 and a BGP tunnel toward P-3 and PE-4, as follows:

```
[/]
A:admin@P-2# show router tunnel-table

===============================================================================
IPv4 Tunnel Table (Router: Base)
===============================================================================
Destination           Owner     Encap TunnelId  Pref   Nexthop        Metric
   Color
-------------------------------------------------------------------------------
192.0.2.1/32          ldp       MPLS  65537     9      192.168.12.1   10
192.0.2.3/32          bgp       MPLS  262146    12     192.168.23.2   1000
192.0.2.4/32          bgp       MPLS  262145    12     192.168.23.2   1000
-------------------------------------------------------------------------------
Flags: B = BGP or MPLS backup hop available
       L = Loop-Free Alternate (LFA) hop available
       E = Inactive best-external BGP route
       k = RIB-API or Forwarding Policy backup hop
===============================================================================
```

One BGP-EVPN IMET route is received and used on PE-1:

```
[/]
A:admin@PE-1# show router bgp routes evpn incl-mcast
===============================================================================
 BGP Router ID:192.0.2.1        AS:64501       Local AS:64501
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Inclusive-Mcast Routes
===============================================================================
Flag  Route Dist.         OrigAddr
      Tag                 NextHop
-------------------------------------------------------------------------------
u*>i  192.0.2.4:1         192.0.2.4
      0                   192.0.2.4

-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The preceding route is an IMET-IR route received from node PE-4, as follows:

```
[/]
A:admin@PE-1# show router bgp routes evpn incl-mcast detail
===============================================================================
 BGP Router ID:192.0.2.1        AS:64501       Local AS:64501
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Inclusive-Mcast Routes
===============================================================================
Original Attributes

Network        : n/a
Nexthop        : 192.0.2.4
From           : 192.0.2.4
Res. Nexthop   : n/a
Local Pref.    : n/a                   Interface Name : NotAvailable
Aggregator AS  : None                  Aggregator     : None
Atomic Aggr.   : Not Atomic            MED            : None
AIGP Metric    : None                  IGP Cost       : 0
Connector      : None
Community      : target:64502:1 bgp-tunnel-encap:MPLS
Cluster        : No Cluster Members
Originator Id  : None                  Peer Router Id : 192.0.2.4
Flags          : Used Valid Best IGP
Route Source   : External
AS-Path        : 64502
EVPN type      : INCL-MCAST
Tag            : 0
Originator IP  : 192.0.2.4
Route Dist.    : 192.0.2.4:1
Route Tag      : 0
Neighbor-AS    : 64502
Orig Validation: N/A
Source Class   : 0                     Dest Class     : 0
Add Paths Send : Default
Last Modified  : 00h01m57s
-------------------------------------------------------------------------------
PMSI Tunnel Attributes :
Tunnel-type    : Ingress Replication
Flags          : Type: RNVE(0) BM: 0 U: 0 Leaf: not required
MPLS Label     : LABEL 524284
Tunnel-Endpoint: 192.0.2.4
-------------------------------------------------------------------------------
---snip---
```

PE-4 has received an IMET-P2MP-IR route sent by root node PE-1, as follows:

```
[/]
A:admin@PE-4# show router bgp routes evpn incl-mcast detail
===============================================================================
 BGP Router ID:192.0.2.4        AS:64502       Local AS:64502
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
BGP EVPN Inclusive-Mcast Routes
===============================================================================
Original Attributes

Network        : n/a
Nexthop        : 192.0.2.1
From           : 192.0.2.1
Res. Nexthop   : n/a
Local Pref.    : n/a                    Interface Name : NotAvailable
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : 0
Connector      : None
Community      : target:64501:1 bgp-tunnel-encap:MPLS
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.1
Flags          : Used Valid Best IGP
Route Source   : External
AS-Path        : 64501
EVPN type      : INCL-MCAST
Tag            : 0
Originator IP  : 192.0.2.1
Route Dist.    : 192.0.2.1:1
Route Tag      : 0
Neighbor-AS    : 64501
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
Add Paths Send : Default
Last Modified  : 00h01m59s
-------------------------------------------------------------------------------
PMSI Tunnel Attributes :
Tunnel-type    : Composite LDP P2MP IR
Flags          : Type: RNVE(0) BM: 0 U: 0 Leaf: not required
MPLS Label1 Ag : LABEL 0
MPLS Label2 IR : LABEL 524284
Root-Node      : 192.0.2.1              LSP-ID         : 8193
-------------------------------------------------------------------------------
---snip---
```

When leaf node PE-4 receives this IMET-P2MP-IR route, a provider tunnel is established toward the root. One P2MP LDP binding of opaque type GRT recursive is active on PE-4:

```
[/]
A:admin@PE-4# show router ldp bindings active p2mp summary ipv4
 No. of Generic IPv4 P2MP Active Bindings: 0
 No. of In-Band-SSM IPv4 P2MP Active Bindings: 0
 No. of In-Band-VPN-SSM IPv4 P2MP Active Bindings: 0
 No. of In-Band-SSM IPv4 P2MP Active Bindings: 0
 No. of VPN Recursive with Generic IPv4 P2MP Active Bindings: 0
 No. of GRT Recursive with Generic IPv4 P2MP Active Bindings: 1
```

The following GRT recursive P2MP LDP binding with root P-3 and inner root PE-1 is active on PE-4:

```
[/]
A:admin@PE-4# show router ldp bindings active p2mp opaque-type grt-recursive ipv4

===============================================================================
LDP Bindings (IPv4 LSR ID 192.0.2.4)
             (IPv6 LSR ID ::)
===============================================================================
Label Status:
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
        e - Label ELC
FEC Flags:
        LF - Lower FEC, UF - Upper FEC, M - Community Mismatch,
        BA - ASBR Backup FEC
===============================================================================
LDP GRT Recursive with Generic IPv4 P2MP Bindings (Active)
===============================================================================
P2MP-Id
InnerRootAddr                           Interface
RootAddr                                Op
IngLbl                                  EgrLbl
EgrNH                                   EgrIf/LspId
-------------------------------------------------------------------------------
8193
192.0.2.1                               73728
192.0.2.3                               Pop
524283                                    --
  --                                      --

-------------------------------------------------------------------------------
No. of GRT Recursive with Generic IPv4 P2MP Active Bindings: 1
===============================================================================
```

The following detailed output shows that the P2MP type is 7:

```
[/]
A:admin@PE-4# show router ldp bindings active p2mp opaque-type grt-recursive ipv4 detail

===============================================================================
LDP Bindings (IPv4 LSR ID 192.0.2.4)
             (IPv6 LSR ID ::)
===============================================================================
Label Status:
        U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
        e - Label ELC
FEC Flags:
        LF - Lower FEC, UF - Upper FEC, M - Community Mismatch,
        BA - ASBR Backup FEC
===============================================================================
LDP GRT Recursive with Generic IPv4 P2MP Bindings (Active)
===============================================================================
-------------------------------------------------------------------------------
P2MP Type      : 7                  P2MP-Id    : 8193
Root-Addr      : 192.0.2.3
InnerRoot-Addr : 192.0.2.1
-------------------------------------------------------------------------------
Op             : Pop
Ing Lbl        : 524283
Egr Lbl        :   --
Egr Int/LspId  :   --
EgrNextHop     :   --
Egr. Flags     : None               Ing. Flags : None
===============================================================================
No. of GRT Recursive with Generic IPv4 P2MP Active Bindings: 1
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

P-3 has two P2MP LDP bindings active: one toward the—downstream—lower FEC (LF) PE-4 and another to the—upstream—upper FEC (UF) P-2, as follows. Both P2MP LDP bindings have inner root 192.0.2.1 and they are stitched to each other.

```
[/]
A:admin@P-3# show router ldp bindings active p2mp opaque-type grt-recursive ipv4

===============================================================================
LDP Bindings (IPv4 LSR ID 192.0.2.3)
             (IPv6 LSR ID ::)
===============================================================================
Label Status:
        U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
        e - Label ELC
FEC Flags:
        LF - Lower FEC, UF - Upper FEC, M - Community Mismatch,
        BA - ASBR Backup FEC
===============================================================================
LDP GRT Recursive with Generic IPv4 P2MP Bindings (Active)
===============================================================================
P2MP-Id
InnerRootAddr                           Interface
RootAddr                                Op
IngLbl                                  EgrLbl
EgrNH                                   EgrIf/LspId
-------------------------------------------------------------------------------
8193
192.0.2.1                               Unknw
192.0.2.3 (LF)                          Push
  --                                    524283
192.168.34.2                            1/1/1

8193
192.0.2.1                               Unknw
192.168.23.1 (UF)                       Swap
524279                                  Stitched
  --                                      --

-------------------------------------------------------------------------------
No. of GRT Recursive with Generic IPv4 P2MP Active Bindings: 2
===============================================================================
```

P-2 has two P2MP LDP bindings active: one GRT recursive (type 7) and one generic (type 1), as follows:

```
[/]
A:admin@P-2# show router ldp bindings active p2mp summary ipv4
 No. of Generic IPv4 P2MP Active Bindings: 1
 No. of In-Band-SSM IPv4 P2MP Active Bindings: 0
 No. of In-Band-VPN-SSM IPv4 P2MP Active Bindings: 0
 No. of In-Band-SSM IPv4 P2MP Active Bindings: 0
 No. of VPN Recursive with Generic IPv4 P2MP Active Bindings: 0
 No. of GRT Recursive with Generic IPv4 P2MP Active Bindings: 1
```

On P-2, the GRT recursive P2MP LDP binding with inner root 192.0.2.1 is toward LF P-3, as follows:

```
[/]
A:admin@P-2# show router ldp bindings active p2mp opaque-type grt-recursive ipv4

===============================================================================
LDP Bindings (IPv4 LSR ID 192.0.2.2)
             (IPv6 LSR ID ::)
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Label Status:
        U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
        e - Label ELC
FEC Flags:
        LF - Lower FEC, UF - Upper FEC, M - Community Mismatch,
        BA - ASBR Backup FEC
===============================================================================
LDP GRT Recursive with Generic IPv4 P2MP Bindings (Active)
===============================================================================
P2MP-Id
InnerRootAddr                           Interface
RootAddr                                Op
IngLbl                                  EgrLbl
EgrNH                                   EgrIf/LspId
-------------------------------------------------------------------------------
8193
192.0.2.1                               Unknw
192.168.23.1 (LF)                       Push
  --                                    524279
192.168.23.2                            1/1/1

-------------------------------------------------------------------------------
No. of GRT Recursive with Generic IPv4 P2MP Active Bindings: 1
===============================================================================
```

On P-2, the generic P2MP LDP binding is toward UF PE-1, as follows. The UF has root address 192.0.2.1 and is stitched to the LF with inner root address 192.0.2.1.

```
[/]
A:admin@P-2# show router ldp bindings active p2mp opaque-type generic ipv4

===============================================================================
LDP Bindings (IPv4 LSR ID 192.0.2.2)
             (IPv6 LSR ID ::)
===============================================================================
Label Status:
        U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
        e - Label ELC
FEC Flags:
        LF - Lower FEC, UF - Upper FEC, M - Community Mismatch,
        BA - ASBR Backup FEC
===============================================================================
LDP Generic IPv4 P2MP Bindings (Active)
===============================================================================
P2MP-Id                                 Interface
RootAddr                                Op
IngLbl                                  EgrLbl
EgrNH                                   EgrIf/LspId
-------------------------------------------------------------------------------
8193                                    Unknw
192.0.2.1 (UF)                          Swap
524279                                  Stitched
  --                                      --

-------------------------------------------------------------------------------
No. of Generic IPv4 P2MP Active Bindings: 1
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

PE-1 has one P2MP LDP active binding toward LF P-2 (type 1- generic):

```
[/]
A:admin@PE-1# show router ldp bindings active p2mp opaque-type generic ipv4

===============================================================================
LDP Bindings (IPv4 LSR ID 192.0.2.1)
             (IPv6 LSR ID ::)
===============================================================================
Label Status:
        U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
        e - Label ELC
FEC Flags:
        LF - Lower FEC, UF - Upper FEC, M - Community Mismatch,
        BA - ASBR Backup FEC
===============================================================================
LDP Generic IPv4 P2MP Bindings (Active)
===============================================================================
P2MP-Id                                     Interface
RootAddr                                    Op
IngLbl                                      EgrLbl
EgrNH                                       EgrIf/LspId
-------------------------------------------------------------------------------
8193                                        73728
192.0.2.1                                   Push
  --                                        524279
192.168.12.2                                1/1/1

-------------------------------------------------------------------------------
No. of Generic IPv4 P2MP Active Bindings: 1
===============================================================================
```

The EVPN BUM traffic is forwarded from the root node PE-1 to the leaf node PE-4 over the P2MP tree. The following command on root node PE-1 shows that an EVPN destination (that uses a BGP tunnel) toward leaf node PE-4 is established, and can carry multicast traffic (BUM):

```
[/]
A:admin@PE-1# show service id 1 evpn-mpls

===============================================================================
BGP EVPN-MPLS Dest
===============================================================================
TEP Address     Egr Label     Num. MACs   Mcast           Last Change
                Transport:Tnl                             Sup BCast Domain
-------------------------------------------------------------------------------
192.0.2.4       524284        0           bum             06/02/2021 08:31:14
                bgp:262146                                No
-------------------------------------------------------------------------------
Number of entries : 1
-------------------------------------------------------------------------------
===============================================================================
---snip---
```

The provider tunnel in VPLS 1 is established using LDP and the operational state is up, as follows. The router will always use the provider tunnel and not the EVPN-MPLS destination, as long as the provider tunnel Oper State is up:

```
[/]
A:admin@PE-1# show service id 1 provider-tunnel
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Service Provider Tunnel Information
===============================================================================
Type               : inclusive          Root and Leaf      : enabled
Admin State        : enabled            Data Delay Intvl   : 15 secs
PMSI Type          : ldp                LSP Template       :
Remain Delay Intvl : 0 secs             LSP Name used      : 8193
PMSI Owner         : bgpEvpnMpls
Oper State         : up                 Root Bind Id       : 32767
===============================================================================
```

The following SDP of type VplsPmsi is auto-created in VPLS 1 on root node PE-1:

```
[/]
A:admin@PE-1# show service id 1 sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
32767:4294967294 VplsPmsi not applicable  Up      Up        None      3
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
```

The following **tools dump** command shows the originating provider tunnels for VPLS 1 on root node PE-1:

```
[/]
A:admin@PE-1# tools dump service id 1 provider-tunnels type originating

===============================================================================
VPLS 1 Inclusive Provider Tunnels Originating
===============================================================================
ipmsi (LDP)                                     P2MP-ID  Root-Addr
-------------------------------------------------------------------------------
8193                                            8193    192.0.2.1

-------------------------------------------------------------------------------
```

The following command shows the terminating provider tunnels for VPLS 1 on leaf node PE-4:

```
[/]
A:admin@PE-4# tools dump service id 1 provider-tunnels type terminating

===============================================================================
VPLS 1 Inclusive Provider Tunnels Terminating
===============================================================================
ipmsi (LDP)                                     P2MP-ID  Root-Addr
-------------------------------------------------------------------------------
                                                8193    192.0.2.1

-------------------------------------------------------------------------------
```

Optimized Inter-AS Model C for mLDP

When some leaf nodes do not support labeled BGP routes or recursive opaque mLDP label mapping, the ASBR in the AS where the leaf nodes are situated needs to leak the root IP address into the leaf PE

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

IGP, which allows the leaf node PE-4 to send a generic FEC type 1 to join the root. The recursive opaque functionality is pushed to the local ASBR P-3.

Figure 215: Example topology for optimized Inter-AS Model C for mLDP shows the example topology for the optimized inter-AS model C for mLDP.

Figure 215: Example topology for optimized Inter-AS Model C for mLDP

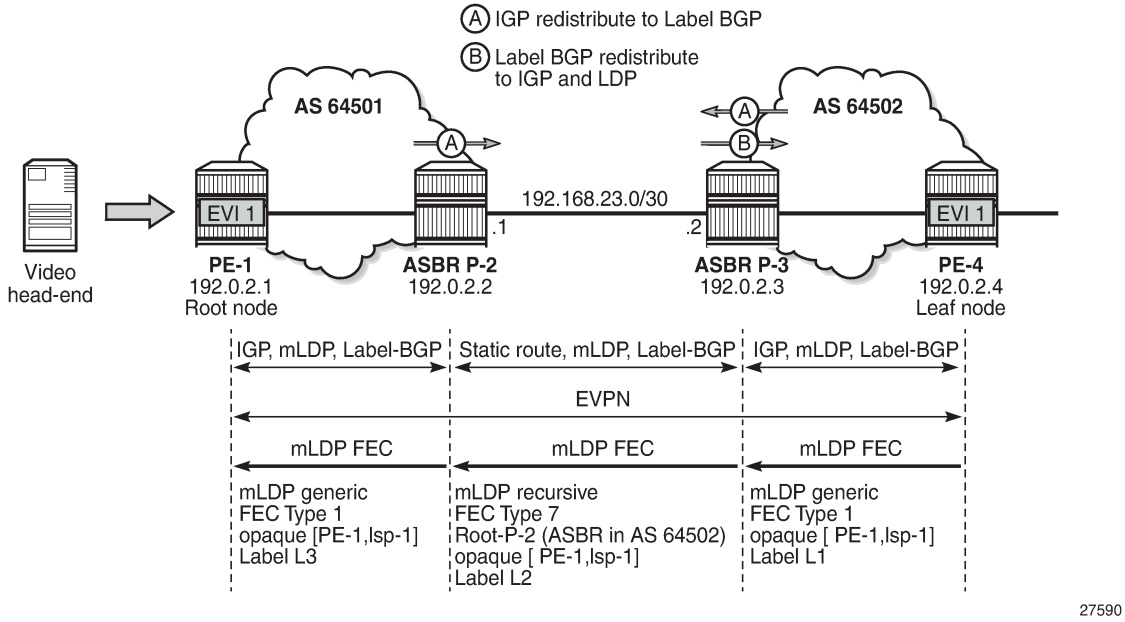

The configuration starts with the configuration in the preceding section Inter-AS Model C for mLDP. The policy to export system prefixes from the ASs to labeled BGP is already configured and applied on both ASBRs. The following additional policies are defined on ASBR P-3 in the AS of the leaf node to export labeled BGP routes to OSPF and to LDP.

```
# on ASBR P-3:
configure {
    policy-options {
        policy-statement "bgpToLdp" {
            entry 10 {
                from {
                    protocol {
                        name [bgp-label]
                    }
                }
                to {
                    protocol {
                        name [ldp]
                    }
                }
                action {
                    action-type accept
                }
            }
        }
        policy-statement "bgpToOspf" {
            entry 10 {
                from {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    protocol {
                        name [bgp-label]
                    }
                }
                to {
                    protocol {
                        name [ospf]
                    }
                }
                action {
                    action-type accept
                }
            }
        }
```

Policy "bgpToOspf" is configured in the OSPF context and policy "bgpToLdp" in the **ldp** context, as follows:

```
# on ASBR P-3:
configure {
    router "Base" {
        ldp {
            export-tunnel-table ["bgpToLdp"]
        }
        ospf 0 {
            export-policy ["bgpToOspf"]
        }
```

Optimized Inter-AS Model C for mLDP - Verification

The prefixes from AS 64501 are now exported to OSPF and LDP in AS 64502; therefore, leaf node PE-4 will no longer use the labeled BGP routes to a node in AS 64501.

```
[/]
A:admin@PE-4# show router bgp routes label-ipv4
===============================================================================
 BGP Router ID:192.0.2.4        AS:64502       Local AS:64502
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop (Router)                                 Path-Id     IGP Cost
      As-Path                                                      Label
-------------------------------------------------------------------------------
*i    192.0.2.1/32                                     100         10
      192.0.2.3                                        None        10
      64501                                                        524283
*i    192.0.2.2/32                                     100         None
      192.0.2.3                                        None        10
      64501                                                        524282
-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The following route table in PE-4 shows that an OSPF route exists toward prefix 192.0.2.1:

```
[/]
A:admin@PE-4# show router route-table 192.0.2.1

===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
192.0.2.1/32                                  Remote  OSPF      00h00m21s  150
       192.168.34.1                                                 10
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

On PE-4, all tunnels are LDP tunnels; no BGP tunnels are established from PE-4 to PE-1 and P-2, as follows:

```
[/]
A:admin@PE-4# show router tunnel-table

===============================================================================
IPv4 Tunnel Table (Router: Base)
===============================================================================
Destination           Owner     Encap TunnelId  Pref   Nexthop        Metric
   Color
-------------------------------------------------------------------------------
192.0.2.1/32          ldp       MPLS  65538     9      192.168.34.1   10
192.0.2.2/32          ldp       MPLS  65539     9      192.168.34.1   1
192.0.2.3/32          ldp       MPLS  65537     9      192.168.34.1   10
-------------------------------------------------------------------------------
Flags: B = BGP or MPLS backup hop available
       L = Loop-Free Alternate (LFA) hop available
       E = Inactive best-external BGP route
       k = RIB-API or Forwarding Policy backup hop
===============================================================================
```

On all other nodes, the route table and tunnel table are the same as in the non-optimized scenario. The route table and the tunnel table for ASBR P-3 are as follows:

```
[/]
A:admin@P-3# show router route-table protocol bgp-label

===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
192.0.2.1/32                                  Remote  BGP_LABEL 00h10m48s  170
       192.168.23.1                                                 0
192.0.2.2/32                                  Remote  BGP_LABEL 00h10m48s  170
       192.168.23.1                                                 0
-------------------------------------------------------------------------------
No. of Routes: 2
Flags: n = Number of times nexthop is repeated
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

```
[/]
A:admin@P-3# show router tunnel-table

===============================================================================
IPv4 Tunnel Table (Router: Base)
===============================================================================
Destination           Owner     Encap TunnelId  Pref   Nexthop        Metric
   Color
-------------------------------------------------------------------------------
192.0.2.1/32          bgp       MPLS  262145    12     192.168.23.1   1000
192.0.2.2/32          bgp       MPLS  262146    12     192.168.23.1   1000
192.0.2.4/32          ldp       MPLS  65537     9      192.168.34.2   10
-------------------------------------------------------------------------------
Flags: B = BGP or MPLS backup hop available
       L = Loop-Free Alternate (LFA) hop available
       E = Inactive best-external BGP route
       k = RIB-API or Forwarding Policy backup hop
===============================================================================
```

Root node PE-1 will send an IMET-P2MP-IR route to leaf node PE-4. PE-4 will send an mLDP label mapping message type 1 instead of type 7, because there is an LDP tunnel toward PE-1 instead of a BGP tunnel. The only P2MP mLDP binding on leaf node PE-4 is a generic P2MP binding, as follows:

```
[/]
A:admin@PE-4# show router ldp bindings p2mp summary ipv4
 No. of Generic IPv4 P2MP Bindings: 1
 No. of In-Band-SSM IPv4 P2MP Bindings: 0
 No. of In-Band-VPN-SSM IPv4 P2MP Bindings: 0
 No. of Recursive with In-Band-SSM IPv4 P2MP Bindings: 0
 No. of VPN Recursive with Generic IPv4 P2MP Bindings: 0
 No. of GRT Recursive with Generic IPv4 P2MP Bindings: 0
```

PE-4 sends the following mLDP label mapping message type 1 with root address 192.0.2.1 (PE-1) to its peer P-3.

```
15 2021/06/02 08:39:21.702 UTC MINOR: DEBUG #2001 Base LDP
"LDP: LDP
Send Label Mapping packet (msgId 100) to 192.0.2.3:0
Protocol version = 1
Label 524279 advertised for the following FECs
P2MP: root = 192.0.2.1, T: 1, L: 4, TunnelId: 8193
"
```

The following generic P2MP mLDP binding for root address 192.0.2.1 is seen on PE-4:

```
[/]
A:admin@PE-4# show router ldp bindings p2mp opaque-type generic detail ipv4

===============================================================================
LDP Bindings (IPv4 LSR ID 192.0.2.4)
             (IPv6 LSR ID ::)
===============================================================================
Label Status:
        U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        e - Label ELC
FEC Flags:
        LF - Lower FEC, UF - Upper FEC, M - Community Mismatch,
        BA - ASBR Backup FEC
===============================================================================
LDP Generic IPv4 P2MP Bindings
===============================================================================
-------------------------------------------------------------------------------
P2MP Type      : 1                 P2MP-Id    : 8193
Root-Addr      : 192.0.2.1
-------------------------------------------------------------------------------
Peer           : 192.0.2.3:0
Ing Lbl        : 524279U
Egr Lbl        :   --
Egr Int/LspId  :   --
EgrNextHop     :   --
Egr. Flags     : None              Ing. Flags : None
===============================================================================
No. of Generic IPv4 P2MP Bindings: 1
===============================================================================
```

ASBR P-3 receives the generic P2MP mLDP label mapping message from PE-4 (T: 1) and resolves the root node 192.0.2.1 to next-hop P-2. P-3 sends a GRT recursive P2MP mLDP label mapping message (T: 7) with inner root 192.0.2.1 to its peer P-2 (root 192.168.23.1) in AS 64501:

```
17 2021/06/02 08:39:21.696 UTC MINOR: DEBUG #2001 Base LDP
"LDP: LDP
Recv Label Mapping packet (msgId 100) from 192.0.2.4:0
Protocol version = 1
Label 524279 advertised for the following FECs
P2MP: root = 192.0.2.1, T: 1, L: 4, TunnelId: 8193
"
```

```
18 2021/06/02 08:39:21.696 UTC MINOR: DEBUG #2001 Base LDP
"LDP: LDP
Send Label Mapping packet (msgId 81) to 192.168.23.1:0
Protocol version = 1
Label 524278 advertised for the following FECs
P2MP: root = 192.168.23.1, T: 7, L: 17 (InnerRoot: 192.0.2.1 T: 1, L: 4, TunnelId: 8193)
"
```

```
[/]
A:admin@P-3# show router ldp bindings p2mp opaque-type generic detail ipv4

===============================================================================
LDP Bindings (IPv4 LSR ID 192.0.2.3)
             (IPv6 LSR ID ::)
===============================================================================
Label Status:
        U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
        e - Label ELC
FEC Flags:
        LF - Lower FEC, UF - Upper FEC, M - Community Mismatch,
        BA - ASBR Backup FEC
===============================================================================
LDP Generic IPv4 P2MP Bindings
===============================================================================
-------------------------------------------------------------------------------
P2MP Type      : 1                 P2MP-Id    : 8193
Root-Addr      : 192.0.2.1 (LF)
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
Peer           : 192.0.2.4:0
Ing Lbl        :   --
Egr Lbl        : 524279
Egr Int/LspId  : 1/1/1
EgrNextHop     : 192.168.34.2
Egr. Flags     : None               Ing. Flags : None
Egr If Name    : int-P-3-PE-4
Metric         : 1                  Mtu        : 1564
===============================================================================
No. of Generic IPv4 P2MP Bindings: 1
===============================================================================
```

```
[/]
A:admin@P-3# show router ldp bindings p2mp opaque-type grt-recursive detail ipv4

===============================================================================
LDP Bindings (IPv4 LSR ID 192.0.2.3)
             (IPv6 LSR ID ::)
===============================================================================
Label Status:
        U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
        e - Label ELC
FEC Flags:
        LF - Lower FEC, UF - Upper FEC, M - Community Mismatch,
        BA - ASBR Backup FEC
===============================================================================
LDP GRT Recursive with Generic IPv4 P2MP Bindings
===============================================================================
-------------------------------------------------------------------------------
P2MP Type      : 7                  P2MP-Id    : 8193
Root-Addr      : 192.168.23.1 (UF)
InnerRoot-Addr : 192.0.2.1
-------------------------------------------------------------------------------
Peer           : 192.168.23.1:0
Ing Lbl        : 524278U
Egr Lbl        :   --
Egr Int/LspId  :   --
EgrNextHop     :   --
Egr. Flags     : None               Ing. Flags : None
===============================================================================
No. of GRT Recursive with Generic IPv4 P2MP Bindings: 1
===============================================================================
```

The P2MP mLDP bindings on P-2 and PE-1 are the same as in the previous non-optimized inter-AS model C for mLDP scenario. P-2 has one GRT recursive mLDP binding to P-3 and one generic mLDP binding to root node PE-1, whereas PE-1 only has a generic mLDP binding to P-2.

The following command on root node PE-1 shows that an EVPN-MPLS destination is created to the leaf node PE-4. This EVPN destination runs over a BGP tunnel and can transport multicast (BUM) traffic. However, as discussed in the preceding section, the EVPN destination is used for BUM traffic only in the case where the provider tunnel goes operationally down.

```
[/]
A:admin@PE-1# show service id 1 evpn-mpls

===============================================================================
BGP EVPN-MPLS Dest
===============================================================================
TEP Address     Egr Label     Num. MACs   Mcast           Last Change
                Transport:Tnl                             Sup BCast Domain
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
192.0.2.4       524284        0           bum             06/02/2021 08:31:14
                bgp:262146                                No
-------------------------------------------------------------------------------
Number of entries : 1
-------------------------------------------------------------------------------
===============================================================================
---snip---
```

The same command on the leaf node PE-4 shows an EVPN destination running on an LDP tunnel instead of a BGP tunnel. This destination is used whenever PE-4 needs to send BUM traffic to PE-1:

```
[/]
A:admin@PE-4# show service id 1 evpn-mpls

===============================================================================
BGP EVPN-MPLS Dest
===============================================================================
TEP Address     Egr Label     Num. MACs   Mcast           Last Change
                Transport:Tnl                             Sup BCast Domain
-------------------------------------------------------------------------------
192.0.2.1       524284        0           bum             06/02/2021 08:31:14
                ldp:65538                                 No
-------------------------------------------------------------------------------
Number of entries : 1
-------------------------------------------------------------------------------
===============================================================================
---snip---
```

The other **show** commands in the Inter-AS Model C for mLDP section have an identical output for both scenarios.

Conclusion

P2MP mLDP is supported in inter-AS model C for EVPN-MPLS services with or without optimization. Optimization in this chapter refers to the ability to set up an end-to-end mLDP tunnel without the need for recursive opaque mLDP FECs on the leaf nodes. A similar configuration is applied in the case of seamless MPLS across different areas.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

P2MP mLDP Tunnels for BUM Traffic in EVPN-MPLS Services

This chapter provides information about P2MP mLDP Tunnels for BUM Traffic in EVPN-MPLS Services.

Topics in this chapter include:

Applicability

This chapter was initially written for SR OS Release 14.0.R4, but the CLI in the current edition is based on SR OS Release 23.3.R3.

Point-to-Multipoint (P2MP) multicast Label Distribution Protocol (mLDP) tunnels for Broadcast, Unknown unicast, and Multicast (BUM) traffic in Ethernet Virtual Private Network Multiprotocol Label Switching (EVPN-MPLS) networks are supported in SR OS Release 14.0.R1, and later. Internet Group Management Protocol (IGMP) snooping support for EVPN-MPLS services is supported in SR OS Release 14.0.R4, and later.

Overview

Service providers are moving their existing VPN services to EVPN. Providers using P2MP LSPs for VPLS services expect the same capabilities in EVPN. Before SR OS Release 14.0.R1, only Ingress Replication (IR) was supported. This works well for broadcast and unknown unicast traffic, but it is inefficient for multicast. Ingress replication does not use a multicast mechanism. Instead, the parent node makes n individual copies and unicasts each copy through an MPLS or IP tunnel to each child node.

BUM traffic is sent from a root node to a number of leaf nodes, but leaf nodes are also allowed to send BUM traffic to root nodes. If most BUM traffic is flowing from a few root nodes to leaf nodes, it would be inefficient to promote all leaf nodes to root-and-leaf nodes because of the amount of P2MP tunnels that would need to be set up. Another solution is to use a combination of P2MP mLDP and ingress replication (IR) tunnels in the service. The root nodes send BUM traffic using P2MP tunnels while the leaf nodes use IR tunnels to send BUM traffic to the root nodes. This avoids the need to set up a P2MP tree from each leaf, while it still allows leaf nodes to send BUM traffic to the root nodes.

Figure 216: P2MP mLDP tree with root node PE-1 and leaf nodes PE-5, PE-6, and PE-7 shows a multicast mLDP tree with root node PE-1 and leaf nodes PE-5, PE-6, and PE-7.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 216: P2MP mLDP tree with root node PE-1 and leaf nodes PE-5, PE-6, and PE-7

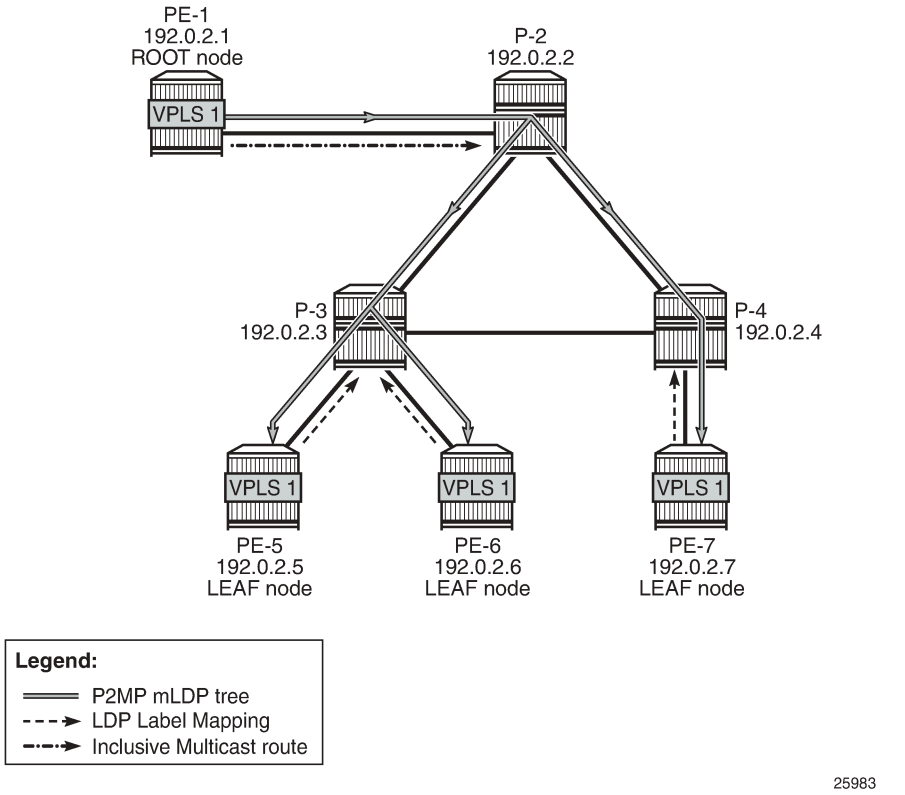

The Inclusive Multicast Ethernet Tag (IMET) route (EVPN route type 3) sent by root node PE-1 contains the required information to set up an mLDP tree, such as the root node IP address and an opaque value. As described in chapter "Multicast Label Distribution Protocol" in the MPLS volume of the *7450 ESS, 7750 SR, and 7950 XRS Advanced Configuration Guide - Part I*, the mLDP tree is set up from the leaf nodes toward the root.

The LDP label mapping message contains the root node address, an opaque value, and an MPLS label. The leaf nodes send an LDP label mapping message to their upstream next hop toward the root node of the tree. Each transit node that has received such LDP label mapping message generates a new LDP label mapping message to its upstream next hop toward the root. This is repeated until the root node receives an LDP label mapping message and the multicast tree is completed.

Figure 216: P2MP mLDP tree with root node PE-1 and leaf nodes PE-5, PE-6, and PE-7 shows a P2MP mLDP tree rooted in PE-1, which is optimal for multicast traffic. However, no P2MP mLDP tree needs to be rooted in PE-5, PE-6, and PE-7 for the reverse direction. These three PEs can use IR to send traffic to the root (and to the other leaf nodes if needed).

EVPN route type 3 is used for setting up the flooding tree for a specified VPLS service. EVPN route type 3 includes the Provider Multicast Service Interface (PMSI) Tunnel Attribute (PMSI Tunnel Attribute = PTA), which can have different formats depending on the tunnel type; see Figure 217: BGP-EVPN route type 3 with PTA.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 217: BGP-EVPN route type 3 with PTA

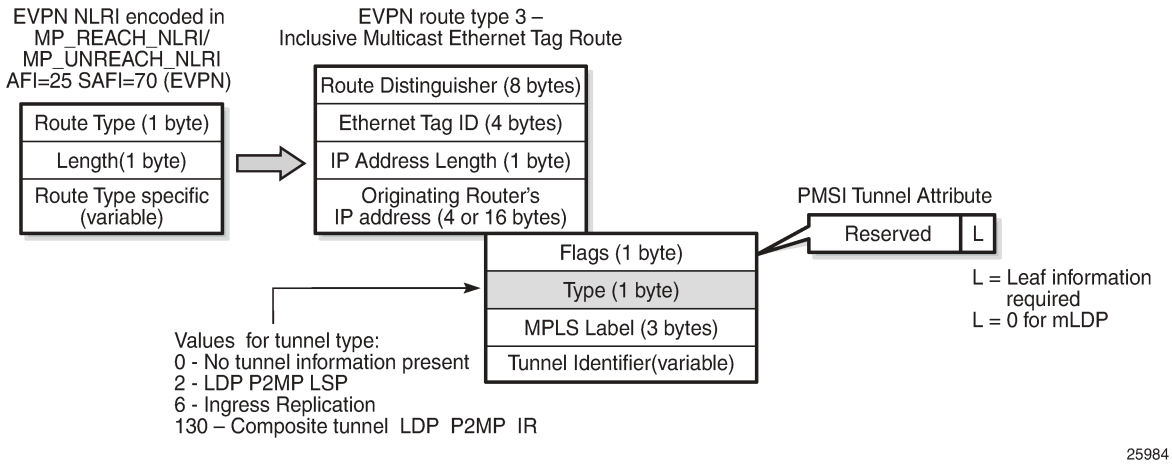

The following route values are used for EVPN-MPLS services:

- The route distinguisher (RD) is taken from the RD of the VPLS service, which can be configured in the BGP context or auto-derived from the BGP-EVPN EVPN Instance (EVI) value. In this case, the RD is auto-derived from the EVI, resulting in a value of 192.0.2.1:1 for VPLS 1 on PE-1.
- The Ethernet tag ID equals 0.
- The IP address length equals 32.
- The originating router's IP address carries the IPv4 system address.
- The PTA can have different formats depending on the tunnel type enabled in the service. The SR OS EVPN-MPLS implementation supports the following tunnel types (SR OS supports different tunnel types for EVPN-VXLAN):
  - Tunnel type 2 - P2MP mLDP
    - The route is referred to as an Inclusive Multicast Ethernet Tag Point-to-Multipoint (IMET-P2MP).
    - Flags: leaf not required.
    - The MPLS label is zero.
    - The tunnel identifier includes the root node address and an opaque number. This is the tunnel identifier that the leaf nodes use to join to the mLDP P2MP tree.
  - Tunnel type 6 - Ingress Replication (IR)
    - The route is referred to as an Inclusive Multicast Ethernet Tag Ingress Replication (IMET-IR).
    - Flags: leaf not required.
    - The MPLS label is a non-zero, downstream allocated label. This MPLS label is allocated to the service and is the same for unicast MAC/IP routes for the same service, unless **ingress-replication-bum-label** is configured in the service.
    - The tunnel identifier is the tunnel endpoint and is equal to the originating IP address.
  - Tunnel type 130 - Composite tunnel: Type: C-bit (composite) + type 2 (mLDP)
    - The route is referred to as an IMET-P2MP-IR.
    - Flags: leaf not required.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- MPLS label 1 equals zero.
- MPLS label 2 is a non-zero, downstream allocated label (as any other IR label). The leaf nodes use the label to set up an EVPN-MPLS binding to the root and add it to the default multicast list.
- The mLDP tunnel identifier is the root node address and an opaque number. This is the tunnel identifier that the leaf nodes use to join the mLDP P2MP tree.

Figure 218: PTA for composite tunnel IMET-P2MP-IR shows the PTA for tunnel type 130.

Figure 218: PTA for composite tunnel IMET-P2MP-IR

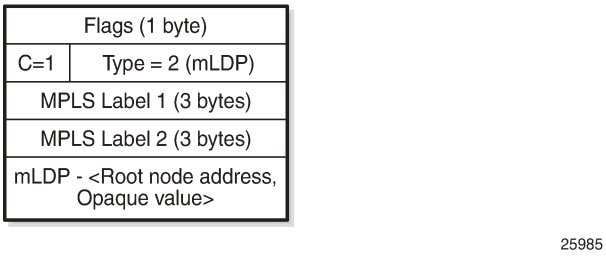

The composite bit C is set, indicating that the PTA identifies two tunnels: the transmit tunnel is a P2MP mLDP tunnel and the receive tunnel is an IR tunnel.

IMET-P2MP-IR routes

The composite tunnel type is an optimized solution that combines mLDP and IR within the same EVPN service so that each root node sends BUM traffic using the P2MP tunnel whereas each leaf-only node sends BUM traffic to the root node using IR.

- PEs configured with **root-and-leaf** can send all BUM traffic over P2MP mLDP tunnels while they receive BUM traffic either over P2MP mLDP tunnels (from other root-and-leaf nodes) or over ingress-replication tunnels (from leaf-only nodes).
- PEs configured without **root-and-leaf** (default setting) can use IR to send BUM traffic to root nodes and other leaf-only nodes, while receiving BUM traffic over either P2MP mLDP tunnels (from root nodes) or ingress-replication tunnels (from leaf-only nodes).

The root PEs signal an IMET-P2MP-IR route, indicating that they intend to transmit BUM traffic using an mLDP P2MP tunnel, while they can receive traffic over an IR EVPN-MPLS binding. Composite tunnels reduce the number of P2MP mLDP tunnels that the PE/P routers in the EVI need to handle, because no full mesh of P2MP tunnels among all the PEs in the EVI is required. This is important (in terms of scaling) in services where there are just a pair of root nodes sending BUM in P2MP tunnels and hundreds of leaf nodes that only need to send BUM traffic to the root nodes using IR tunnels.

Configuration

Initial configuration

The PE and P nodes have the following initial configuration:

- The ports between the routers are configured as network ports and have router interfaces configured.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- IS-IS is enabled on all the router interfaces.
- LDP is enabled on all the router interfaces.
- BGP is enabled on all PEs with route reflector (RR) P-2. The BGP configuration on RR P-2 is as follows:

```
# On P-2:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-import true
            vpn-apply-export true
            peer-ip-tracking true
            rapid-withdrawal true
            split-horizon true
            ebgp-default-reject-policy {
                import false
                export false
            }
            rapid-update {
                evpn true
            }
            group "internal" {
                peer-as 64500
                family {
                    evpn true
                }
                cluster {
                    cluster-id 1.1.1.1
                }
            }
            neighbor "192.0.2.1" {
                group "internal"
            }
            neighbor "192.0.2.5" {
                group "internal"
            }
            neighbor "192.0.2.6" {
                group "internal"
            }
            neighbor "192.0.2.7" {
                group "internal"
            }
        }
    }
```

Configure EVPN P2MP mLDP in VPLS Service

On the root node PE-1, VPLS 1 is configured as follows:

```
# On PE-1:
configure {
    service {
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            bgp-evpn {
                evi 1
                mpls 1 {
                    admin-state enable
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            provider-tunnel {
                inclusive {
                    admin-state enable
                    owner bgp-evpn-mpls
                    root-and-leaf true
                    mldp
                }
            }
            sap 1/2/c3/1 {    # sap for ingress traffic from STC
            }
        }
    }
```

The configuration options in the **bgp-evpn** context of the VPLS are as follows:

```
configure {
    service {
        vpls "VPLS 1" {
            bgp-evpn ?
*[ex:/configure service vpls "VPLS 1"]
A:admin@PE-1#             bgp-evpn ?

 bgp-evpn

 accept-ivpls-evpn-    - Accept and process non-zero ethernet-tag MAC routes
  flush
 apply-groups          - Apply a configuration group at this level
 apply-groups-exclude  - Exclude a configuration group at this level
 evi                   - EVPN ID
 ignore-mtu-mismatch   - Ignore MTU mismatch
 incl-mcast-orig-ip    - Originating IP address
 isid-route-target     + Enter the isid-route-target context
 mac-duplication       + Enter the mac-duplication context
 mpls                  + Enter the mpls list instance
 routes                + Enter the routes context
 segment-routing-v6    + Enter the segment-routing-v6 list instance
 vxlan                 + Enter the vxlan list instance
```

By default, the advertisement of the inclusive multicast route with IR is enabled (**ingress-repl-inc-mcast-advertisement**). However, if it is disabled, the router does not send the IMET-IR or IMET-P2MP-IR routes, regardless of the service being enabled for BGP EVPN-MPLS or BGP EVPN-VXLAN.

For information about the other parameters in the **bgp-evpn** context of the VPLS, see chapters EVPN for VXLAN Tunnels (Layer 2) and EVPN for MPLS Tunnels.

The configuration options in the **provider-tunnel inclusive** context are as follows:

```
configure {
    service {
        vpls "VPLS 1" {
            provider-tunnel {
                inclusive ?
*[ex:/configure service vpls "VPLS 1" provider-tunnel]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
A:admin@PE-1#                 inclusive ?

 inclusive

 admin-state           - Administrative state of P2MP LSP as the I-PMSI
 data-delay-interval   - I-PMSI data delay timer
 owner                 - Configure provider-tunnel owner
 root-and-leaf         - Configure whether the provider tunnel acts as a leaf or both a root
 and leaf

 Choice: ipmsi
 mldp                 :- Enable/Disable MLDP
 rsvp                 :+ Enter the rsvp context
```

- The **data-delay-interval** is configured in seconds in the range from 3 to 180 seconds. A node configured with **root-and-leaf** sends all BUM packets (data plane and control plane: ARP, CCMs, and so on) to its provider tunnel after the delay-data-interval has expired. This timer keeps the provider tunnel operationally down until its expiration, and, during that time, the router can use the EVPN-MPLS destinations typically used for IR.
- mLDP is enabled by adding the keyword **mldp** and enabling the provider tunnel (**admin-state enable**).
- The owner must be **bgp-evpn-mpls** if MPLS is enabled in the EVPN.

```
configure {
    service {
        vpls "VPLS 1" {
            provider-tunnel {
                inclusive {
                    owner ?
*[ex:/configure service vpls "VPLS 1" provider-tunnel inclusive]
A:admin@PE-1#                     owner ?

 owner <keyword>
 <keyword> - (bgp-ad|bgp-vpls|bgp-evpn-mpls)
```

Only one of the three possible owner protocols supports the provider tunnel in the service and needs to be set before the provider tunnel can be enabled. By default, no owner is configured. The following error is raised when a user wants to enable the provider tunnel without an owner:

```
*[ex:/configure service vpls "VPLS 1" provider-tunnel inclusive]
MINOR: SVCMGR #12: configure service vpls "VPLS 1" provider-tunnel inclusive admin-state
 - Inconsistent Value error
 - no owner configured for the provider tunnel
```

After the provider tunnel has an owner and is enabled, the owner can only be changed when the provider tunnel is disabled.

```
*[ex:/configure service vpls "VPLS 1" provider-tunnel inclusive]
MINOR: MGMT_CORE #4001: configure service vpls "VPLS 1" provider-tunnel inclusive owner
 - the enabled provider-tunnel cannot have owner set to bgp-vpls when evpn-mpls is enabled
 - configure service vpls "VPLS 1" bgp-evpn mpls 1 admin-state
```

After the owner is set, the corresponding protocol is checked to see if it is enabled in the service configuration.

```
*[ex:/configure service vpls "VPLS 1" provider-tunnel inclusive]
MINOR: SVCMGR #12: configure service vpls "VPLS 1" provider-tunnel inclusive admin-state
 - Inconsistent Value error
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
  - bgp-vpls must be configured when the provider tunnel with owner bgp-vpls is enabled
  - configure service vpls "VPLS 1"
```

- If **ingress-repl-inc-mcast-advertisement** is enabled and the PE is configured with **root-and-leaf**, the router sends an IMET-P2MP-IR route; if the PE is configured without **root-and-leaf** (default), the router sends an IMET-IR route. However, if **ingress-repl-inc-mcast-advertisement** is disabled and the PE is configured with **root-and-leaf**, the router only sends IMET-P2MP routes. Leaf-only nodes do not send any IMET routes at all in case no IR multicast advertisement is allowed.

  Root-and-leaf nodes only send BUM traffic to the P2MP tunnel as long as it is active. If the P2MP tunnel goes operationally down, it starts sending BUM traffic to IR tunnels (EVPN-MPLS destinations shown in the **show service id 1 evpn-mpls** command).
- If a provider tunnel is configured on a node, the router can join P2MP trees as a leaf, by generating an LDP label mapping message including the corresponding P2MP mLDP FEC. If no provider tunnel is configured, the node does not join P2MP mLDP trees, and can only use IR for BUM.
- If one node is configured as root, all other nodes must be configured with provider tunnels; otherwise, they do not receive BUM traffic sent on P2MP tunnels. The configuration of leaf-only node PE-5 is as follows, the main difference with the configuration for the root being the absence of **root-and-leaf** (default setting):

```
# On PE-5:
configure {
    service {
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
            }
            bgp-evpn {
                evi 1
                mpls 1 {
                    admin-state enable
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            provider-tunnel {
                inclusive {
                    admin-state enable
                    owner bgp-evpn-mpls
                    mldp
                }
            }
            sap 1/2/c1/1:1 {    # sap for egress traffic to VPLS 1
            }
        }
    }
```

As described, the tunnel types for BUM traffic are controlled by **ingress-repl-inc-mcast-advertisement** and the **provider-tunnel** context (**root-and-leaf**). The IMET route sending behavior is summarized in Table 11: IMET routes and Tunnel Types advertised based on the configuration .

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Table 11: IMET routes and Tunnel Types advertised based on the configuration

IMET route set	Root + Leaf PE	Leaf-only	No provider-tunnel
IR-mcast advertisement	Composite P2MP + IR	IR	IR
No IR-mcast advertisement	P2MP	-	-

Information about the provider tunnel can be retrieved as follows:

```
[/]
A:admin@PE-1# show service id 1 provider-tunnel

===============================================================================
Service Provider Tunnel Information
===============================================================================
Type               : inclusive          Root and Leaf      : enabled
Admin State        : enabled            Data Delay Intvl   : 15 secs
PMSI Type          : ldp                LSP Template       :
Remain Delay Intvl : 0 secs             LSP Name used      : 8193
PMSI Owner         : bgpEvpnMpls
Oper State         : up                 Root Bind Id       : 32767
-------------------------------------------------------------------------------
Type               : selective          Wildcard SPMSI     : disabled
Admin State        : disabled           Data Delay Intvl   : 3 secs
PMSI Type          : none               Max P2MP SPMSI     : 10
PMSI Owner         : none
===============================================================================
```

Note:
The same IMET-P2MP route cannot be imported by two services at the same time. If two VPLS services (where a provider tunnel is enabled) have the same import route-target, only one service joins the mLDP tree (whichever comes first).

EVPN P2MP mLDP operation

After the root node and leaf nodes are configured as shown, the root node sends BGP EVPN composite IMET-P2MP-IR routes, as follows:

```
# On PE-1:
2 2023/07/03 22:26:27.189 CEST MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 93
    Flag: 0x90 Type: 14 Len: 28 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.1
        Type: EVPN-INCL-MCAST Len: 17 RD: 192.0.2.1:1, tag: 0, orig_addr len: 32, orig_addr:
 192.0.2.1
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        bgp-tunnel-encap:MPLS
    Flag: 0xc0 Type: 22 Len: 25 PMSI:
        Tunnel-type Composite LDP P2MP IR (130)
        Flags: (0x0)[Type: None BM: 0 U: 0 Leaf: not required]
        MPLS Label1 Ag  0
        MPLS Label2 IR 8388480
        Root-Node 192.0.2.1, LSP-ID 0x2001
"
```

The PTA for tunnel type 130 (composite tunnel) has two MPLS labels, of which MPLS label 1 equals zero. MPLS label 2 is used by the downstream nodes to set up the EVPN-MPLS destination to the root node and add it to the default multicast list. The actual MPLS label only uses the high-order 20 bits out of the 24 bits advertised in the MPLS label. Therefore, the value 8388480 needs to be divided by 16 to have the MPLS label: 8388480/16 = 524280. This is because the debug message is shown before the router can parse the label field and see whether it corresponds to an MPLS label (20 bits) or a VXLAN VNI (24 bits).

The tunnel identifier field contains the root node address 192.0.2.1 and the opaque value 0x2001, which corresponds to decimal value 8193. With this tunnel identifier, the leaf nodes can join the mLDP multicast tree toward the root node by sending LDP label mapping messages that contain the root node IP address and the opaque value.

Note:
When static P2MP mLDP tunnels and dynamic P2MP mLDP tunnels used by BGP-EVPN coexist on the same router,Nokia recommends that the static tunnels use a tunnel ID less than 8193. If a tunnel ID is statically configured with a value equal to or greater than 8193, BGP-EVPN may attempt to use the same tunnel ID for services with an enabled provider tunnel and fail to set up an mLDP tunnel.

The root node PE-1 receives IMET-IR routes from all leaf nodes, as shown for the BGP update sent by leaf node PE-5 (via RR P-2):

```
# On PE-1:
3 2023/07/03 22:28:45.189 CEST MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 91
    Flag: 0x90 Type: 14 Len: 28 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.5
        Type: EVPN-INCL-MCAST Len: 17 RD: 192.0.2.5:1, tag: 0, orig_addr len: 32, orig_addr:
 192.0.2.5
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0x80 Type: 9 Len: 4 Originator ID: 192.0.2.5
    Flag: 0x80 Type: 10 Len: 4 Cluster ID:
        1.1.1.1
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:1
        bgp-tunnel-encap:MPLS
    Flag: 0xc0 Type: 22 Len: 9 PMSI:
        Tunnel-type Ingress Replication (6)
        Flags: (0x0)[Type: None BM: 0 U: 0 Leaf: not required]
        MPLS Label 8388480
        Tunnel-Endpoint 192.0.2.5
"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The PTA tunnel type 6 for IR has only one MPLS label, which corresponds to the MPLS label 524280 allocated for the service. The tunnel identifier is the tunnel endpoint 192.0.2.5, which is the system address of the originating leaf node.

On leaf node PE-5, three BGP EVPN inclusive multicast routes have been learned and are used, as follows:

```
[/]
A:admin@PE-5# show router bgp routes evpn incl-mcast
===============================================================================
 BGP Router ID:192.0.2.5        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Inclusive-Mcast Routes
===============================================================================
Flag  Route Dist.         OrigAddr
      Tag                 NextHop
-------------------------------------------------------------------------------
u*>i  192.0.2.1:1         192.0.2.1
      0                   192.0.2.1

u*>i  192.0.2.6:1         192.0.2.6
      0                   192.0.2.6

u*>i  192.0.2.7:1         192.0.2.7
      0                   192.0.2.7

-------------------------------------------------------------------------------
Routes : 3
===============================================================================
```

The details of the BGP EVPN inclusive multicast route sent by root node PE-1 to leaf node PE-5 are as follows:

```
[/]
A:admin@PE-5# show router bgp routes evpn incl-mcast rd 192.0.2.1:1 detail
===============================================================================
 BGP Router ID:192.0.2.5        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Inclusive-Mcast Routes
===============================================================================
Original Attributes

Network        : n/a
Nexthop        : 192.0.2.1
Path Id        : None
From           : 192.0.2.2
Res. Nexthop   : 192.168.35.1
Local Pref.    : 100                   Interface Name : int-PE-5-P-3
Aggregator AS  : None                  Aggregator     : None
Atomic Aggr.   : Not Atomic            MED            : None
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
AIGP Metric    : None                   IGP Cost       : 30
Connector      : None
Community      : target:64500:1 bgp-tunnel-encap:MPLS
Cluster        : 1.1.1.1
Originator Id  : 192.0.2.1              Peer Router Id : 192.0.2.2
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
EVPN type      : INCL-MCAST
Tag            : 0
Originator IP  : 192.0.2.1
Route Dist.    : 192.0.2.1:1
Route Tag      : 0
Neighbor-AS    : n/a
DB Orig Val    : N/A                    Final Orig Val : N/A
Source Class   : 0                      Dest Class     : 0
Add Paths Send : Default
Last Modified  : 00h01m30s
-------------------------------------------------------------------------------
PMSI Tunnel Attributes :
Tunnel-type    : Composite LDP P2MP IR
Flags          : Type: RNVE(0) BM: 0 U: 0 Leaf: not required
MPLS Label1 Ag : LABEL 0
MPLS Label2 IR : LABEL 524280
Root-Node      : 192.0.2.1              LSP-ID         : 8193
-------------------------------------------------------------------------------
---snip---
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

The MPLS label is 524280, as described. The LSP ID equals 8193, which corresponds to the hexadecimal value 0x2001 in the preceding BGP update message sent by the root node PE-1.

To set up the mLDP tree, leaf node PE-5 has generated an LDP label mapping message to the next hop router toward the root, P-3. The label mapping message includes the root address 192.0.2.1, the opaque value 8193, and MPLS label 524279, as follows:

```
[/]
A:admin@PE-5# show router ldp bindings active p2mp ipv4

===============================================================================
LDP Bindings (IPv4 LSR ID 192.0.2.5)
             (IPv6 LSR ID ::)
===============================================================================
Label Status:
        U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
        e - Label ELC
FEC Flags:
        LF - Lower FEC, UF - Upper FEC, M - Community Mismatch,
        BA - ASBR Backup FEC
===============================================================================
LDP Generic IPv4 P2MP Bindings (Active)
===============================================================================
P2MP-Id                                 Interface
RootAddr                                Op
IngLbl                                  EgrLbl
EgrNH                                   EgrIf/LspId
-------------------------------------------------------------------------------
8193                                    73728
192.0.2.1                               Pop
524279                                    --
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
  --                                          --

-------------------------------------------------------------------------------
No. of Generic IPv4 P2MP Active Bindings: 1
---snip---
===============================================================================
```

P-3 has received two label mapping messages: one from PE-5 and one from PE-6. P-3 has sent one label mapping message to its upstream next hop P-2 with label 524279, as follows:

```
[/]
A:admin@P-3# show router ldp bindings active p2mp ipv4 opaque-type generic

===============================================================================
LDP Bindings (IPv4 LSR ID 192.0.2.3)
             (IPv6 LSR ID ::)
===============================================================================
Label Status:
        U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
        e - Label ELC
FEC Flags:
        LF - Lower FEC, UF - Upper FEC, M - Community Mismatch,
        BA - ASBR Backup FEC
===============================================================================
LDP Generic IPv4 P2MP Bindings (Active)
===============================================================================
P2MP-Id                                     Interface
RootAddr                                    Op
IngLbl                                      EgrLbl
EgrNH                                       EgrIf/LspId
-------------------------------------------------------------------------------
8193                                        Unknw
192.0.2.1                                   Swap
524280                                      524279
192.168.35.2                                1/1/c3/1

8193                                        Unknw
192.0.2.1                                   Swap
524280                                      524279
192.168.36.2                                1/1/c4/1

-------------------------------------------------------------------------------
No. of Generic IPv4 P2MP Active Bindings: 2
===============================================================================
```

P-2 has received two label mapping messages: one from P-3 and one from P-4. P-2 has sent a label mapping message toward the root node PE-1 with label 524280, as follows:

```
[/]
A:admin@P-2# show router ldp bindings active p2mp ipv4 opaque-type generic

===============================================================================
LDP Bindings (IPv4 LSR ID 192.0.2.2)
             (IPv6 LSR ID ::)
===============================================================================
Label Status:
        U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
        e - Label ELC
FEC Flags:
        LF - Lower FEC, UF - Upper FEC, M - Community Mismatch,
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        BA - ASBR Backup FEC
===============================================================================
LDP Generic IPv4 P2MP Bindings (Active)
===============================================================================
P2MP-Id                                  Interface
RootAddr                                 Op
IngLbl                                   EgrLbl
EgrNH                                    EgrIf/LspId
-------------------------------------------------------------------------------
8193                                     Unknw
192.0.2.1                                Swap
524280                                   524280
192.168.23.2                             1/1/c2/1

8193                                     Unknw
192.0.2.1                                Swap
524280                                   524280
192.168.24.2                             1/1/c1/1

-------------------------------------------------------------------------------
No. of Generic IPv4 P2MP Active Bindings: 2
===============================================================================
```

When the LDP label reaches the root node PE-1, the mLDP tree is complete and it can be used for BUM traffic.

The following **tools** command shows the provider tunnels for VPLS 1 on root node and leaf nodes. On root node PE-1, there is one originating inclusive provider tunnel and there are no terminating inclusive provider tunnels, as follows:

```
[/]
A:admin@PE-1# tools dump service id 1 provider-tunnels

===============================================================================
VPLS 1 Inclusive Provider Tunnels Originating
===============================================================================
ipmsi (LDP)                                 P2MP-ID  Root-Addr
-------------------------------------------------------------------------------
8193                                        8193    192.0.2.1

-------------------------------------------------------------------------------

===============================================================================
VPLS 1 Inclusive Provider Tunnels Terminating
===============================================================================
ipmsi (LDP)                                 P2MP-ID  Root-Addr
-------------------------------------------------------------------------------

No Tunnels Found
-------------------------------------------------------------------------------
---snip---
```

On leaf node PE-5, no originating inclusive provider tunnels are established; only one terminating provider tunnel, as follows:

```
[/]
A:admin@PE-5# tools dump service id 1 provider-tunnels

===============================================================================
VPLS 1 Inclusive Provider Tunnels Originating
===============================================================================
ipmsi (LDP)                                 P2MP-ID  Root-Addr
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------

No Tunnels Found
-------------------------------------------------------------------------------

===============================================================================
VPLS 1 Inclusive Provider Tunnels Terminating
===============================================================================
ipmsi (LDP)                                     P2MP-ID  Root-Addr
-------------------------------------------------------------------------------
                                                8193    192.0.2.1

-------------------------------------------------------------------------------
---snip---
```

The inclusive provider tunnels are identified by the combination of the P2MP ID (opaque value) and the root address. These parameters are in every label mapping message and they are included in the PTA tunnel identifier for tunnel type 130 (IMET-P2MP-IR) and for tunnel type 2 (IMET-P2MP).

In VPLS 1 on root node PE-1, an SDP of type VplsPmsi is auto-created, as follows:

```
[/]
A:admin@PE-1# show service id 1 sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
32767:4294967294 VplsPmsi not applicable  Up      Up        None      3
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
```

The detailed information about this SDP includes the traffic statistics: ingress/egress and forwarding/dropped, as follows:

```
[/]
A:admin@PE-1# show service id 1 sdp detail

===============================================================================
Services: Service Destination Points Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 32767:4294967294  -(not applicable)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id             : 32767:4294967294         Type              : VplsPmsi
Split Horiz Grp    : (Not Specified)
Etree Root Leaf Tag: Disabled                 Etree Leaf AC     : Disabled
VC Type            : Ether                    VC Tag            : n/a
Admin Path MTU     : 9782                     Oper Path MTU     : 9782
Delivery           : MPLS
Far End            : not applicable           Tunnel Far End    : n/a
---snip---
PMSI Owner         : bgpEvpnMpls

Admin State        : Up                       Oper State        : Up
---snip---
Statistics           :
I. Fwd. Pkts.      : 0                        I. Dro. Pkts.     : 0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
I. Fwd. Octs.      : 0                        I. Dro. Octs.     : 0
E. Fwd. Pkts.      : 30437                    E. Fwd. Octets    : 45533752
---snip---
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
```

IGMP snooping

When IGMP snooping is disabled and a multicast stream enters VPLS 1 on the root node, this stream is sent to all the leaf nodes, even if no receivers join the multicast group on the leaf nodes. In this example, a receiver connected to PE-5 joins a multicast group, but there are no receivers for any multicast group on PE-6 and PE-7. By default, IGMP is disabled and the multicast stream is flooded to all leaf PEs, as can be verified with the following monitor command on PE-6 where no receivers have joined any multicast stream:

```
[/]
A:admin@PE-6# monitor port all-ethernet-rates repeat 15 interval 4

=======================================================================
Monitor statistics for all Ethernet Port Rates
=======================================================================
Port-Id          D                 Bits    Packets     Errors    Util
-----------------------------------------------------------------------
---snip---
-------------------------------------------------------------------------------
At time t = 12 sec (Mode: Rate)
-------------------------------------------------------------------------------
1/1/c1/1         I              2799472       231          0    0.00
                 O                  912         1          0    0.00

1/2/c1/1         I                    0         0          0    0.00
                 O              2773376       231          0    0.00

1/2/c2/1         I              2773376       231          0    0.00
                 O                    0         0          0    0.00

1/2/c3/1         I                    0         0          0    0.00
                 O                    0         0          0    0.00

-------------------------------------------------------------------------------
At time t = 16 sec (Mode: Rate)
-------------------------------------------------------------------------------
1/1/c1/1         I              2750616       227          0    0.00
                 O                  840         1          0    0.00

1/2/c1/1         I                    0         0          0    0.00
                 O              2562816       213          0    0.00

1/2/c2/1         I              2562816       213          0    0.00
                 O                    0         0          0    0.00

1/2/c3/1         I                    0         0          0    0.00
                 O                    0         0          0    0.00

---snip---
=======================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

This implies that bandwidth is wasted, which can be prevented by enabling IGMP snooping. IGMP snooping ensures that multicast traffic is only sent to the receivers that joined a multicast group. IGMP snooping can be enabled in VPLS 1 on all PEs, as follows:

```
configure {
    service {
        vpls "VPLS 1" {
            igmp-snooping {
                admin-state enable
            }
```

A receiver connected to PE-5 has sent an IGMP report whereas PE-6 has no receivers that joined a multicast group. The traffic counters are monitored on the outgoing port to the (potential) receivers. On PE-5, traffic is sent to the receiver, as follows:

```
[/]
A:admin@PE-5# monitor port all-ethernet-rates repeat 15 interval 4

=============================================================================
Monitor statistics for all Ethernet Port Rates
=============================================================================
Port-Id          D                 Bits   Packets     Errors    Util
-----------------------------------------------------------------------------
---snip---
-------------------------------------------------------------------------------
At time t = 12 sec (Mode: Rate)
-------------------------------------------------------------------------------
1/1/c1/1         I              9986792       824          0    0.01
                 O                 1048         2          0    0.00

1/2/c1/1         I                    0         0          0    0.00
                 O              9893312       822          0    0.01

1/2/c2/1         I              9893312       822          0    0.01
                 O                    0         0          0    0.00

1/2/c3/1         I                    0         0          0    0.00
                 O                    0         0          0    0.00

---snip--
=============================================================================
```

On PE-6, no traffic is sent to any receiver, as follows:

```
[/]
A:admin@PE-6# monitor port all-ethernet-rates repeat 15 interval 4

=============================================================================
Monitor statistics for all Ethernet Port Rates
=============================================================================
Port-Id          D                 Bits   Packets     Errors    Util
---snip---
-------------------------------------------------------------------------------
At time t = 12 sec (Mode: Rate)
-------------------------------------------------------------------------------
1/1/c1/1         I              9986456       824          0    0.01
                 O                 1488         2          0    0.00

1/2/c1/1         I                    0         0          0    0.00
                 O                    0         0          0    0.00
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
1/2/c2/1         I                   0         0          0    0.00
                 O                   0         0          0    0.00

1/2/c3/1         I                   0         0          0    0.00
                 O                   0         0          0    0.00

---snip---
===============================================================================
```

IGMP snooping can be enabled in EVPN-MPLS services with IR or provider-tunnel mLDP trees. When IGMP snooping is enabled on the VPLS, all the EVPN-MPLS destinations are added to the MFIB as a single router interface. IGMP queries and reports are properly forwarded to and from EVPN-MPLS destinations.

The following shows the EVPN-MPLS destinations as part of the MFIB when IGMP snooping is enabled:

```
[/]
A:admin@PE-5# show service id 1 mfib

===============================================================================
Multicast FIB, Service 1
===============================================================================
Source Address  Group Address         Port Id                      Svc Id   Fwd
                                                                          Blk
-------------------------------------------------------------------------------
*               *                     sap:1/2/c1/1:1               Local    Fwd
                                      mpls:192.0.2.1:524280        Local    Fwd
                                      mpls:192.0.2.6:524280        Local    Fwd
                                      mpls:192.0.2.7:524280        Local    Fwd
*               * (mac)               mpls:192.0.2.1:524280        Local    Fwd
                                      mpls:192.0.2.6:524280        Local    Fwd
                                      mpls:192.0.2.7:524280        Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 2
===============================================================================
```

Connected to SAP 1/2/c1/1:1, PE-5 has a receiver that joined the multicast stream. EVPN-MPLS is added as a single logical IGMP snooping interface and treated as an mrouter, also on the other leaf nodes, as follows:

```
[/]
A:admin@PE-5# show service id 1 igmp-snooping base

===============================================================================
IGMP Snooping Base info for service 1
===============================================================================
Admin State : Up
Querier     : 172.16.0.5 on SAP 1/2/c1/1:1
SBD service : N/A
Evpn-proxy  : Disabled
-------------------------------------------------------------------------------
Port                      Oper MRtr Pim  Send Max  Max  Max   MVR       Num
Id                        Stat Port Port Qrys Grps Srcs Grp   From-VPLS Grps
                                                      Srcs
-------------------------------------------------------------------------------
sap:1/2/c1/1:1            Up   Yes  No   No   None None None  Local     0
evpn-mpls
                          Up   Yes  No   N/A  N/A  N/A  N/A   N/A       N/A
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

On leaf node PE-5, the receiving host connected to SAP 1/2/c1/1:1 has IP address 172.16.0.5, as follows:

```
[/]
A:admin@PE-5# show service id 1 igmp-snooping mrouters

===============================================================================
IGMP Snooping Multicast Routers for service 1
===============================================================================
MRouter          Port Id                   Up Time        Expires   Version
-------------------------------------------------------------------------------
172.16.0.5       sap:1/2/c1/1:1            0d 00:02:39    131s      3
-------------------------------------------------------------------------------
Number of mrouters: 1
===============================================================================
```

On leaf node PE-6, SAP 1/2/c1/1:1 has no receiving host connected, but EVPN-MPLS is always added as an mrouter, as follows:

```
[/]
A:admin@PE-6# show service id 1 igmp-snooping base

===============================================================================
IGMP Snooping Base info for service 1
===============================================================================
Admin State : Up
Querier     : 172.16.0.5 on evpn-mpls
SBD service : N/A
Evpn-proxy  : Disabled
-------------------------------------------------------------------------------
Port                      Oper MRtr Pim  Send Max  Max  Max   MVR       Num
Id                        Stat Port Port Qrys Grps Srcs Grp   From-VPLS Grps
                                                  Srcs
-------------------------------------------------------------------------------
sap:1/2/c1/1:1            Up   No   No   No   None None None  Local     0
evpn-mpls
                          Up   Yes  No   N/A  N/A  N/A  N/A   N/A       N/A
===============================================================================
```

On PE-6, the only mrouter in the list is the receiving host connected to PE-5, with port ID EVPN-MPLS instead of a local SAP, as follows:

```
[/]
A:admin@PE-6# show service id 1 igmp-snooping mrouters

===============================================================================
IGMP Snooping Multicast Routers for service 1
===============================================================================
MRouter          Port Id                   Up Time        Expires   Version
-------------------------------------------------------------------------------
172.16.0.5       evpn-mpls                 0d 00:02:41    129s      3
-------------------------------------------------------------------------------
Number of mrouters: 1
===============================================================================
```

PBB-EVPN and P2MP mLDP

Provider Backbone Bridging (PBB) EVPN is described in chapter EVPN for PBB over MPLS (PBB-EVPN).

Figure 219: P2MP mLDP in PBB-EVPN shows the setup for P2MP mLDP in PBB-EVPN.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 219: P2MP mLDP in PBB-EVPN

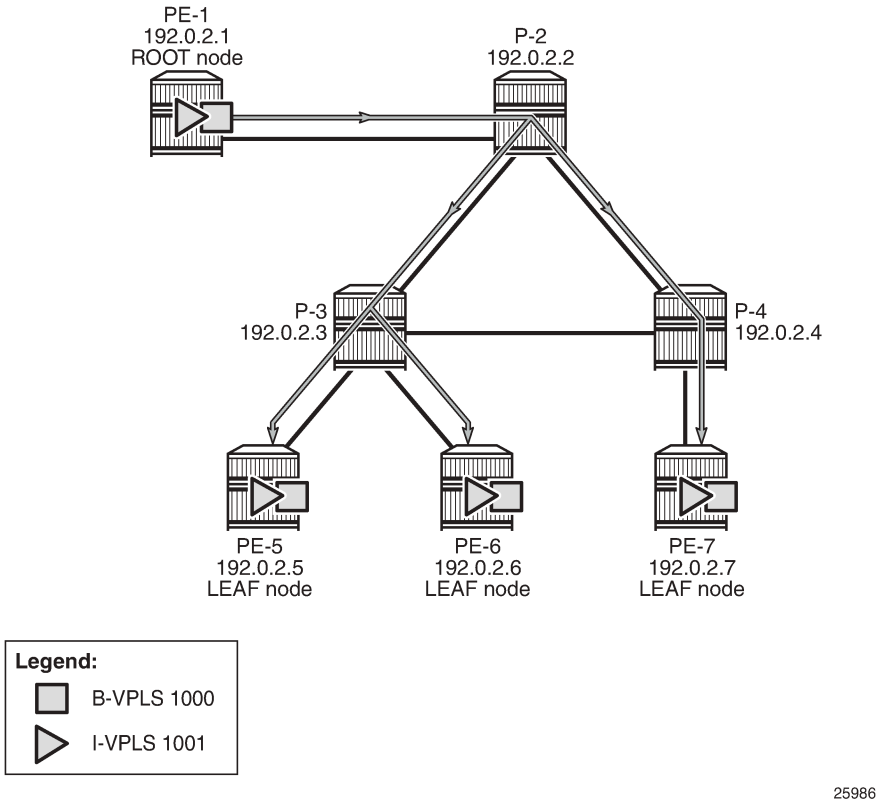

P2MP mLDP tunnels can also be used in PBB-EVPN services. In Release 14.0, the use of **provider-tunnel inclusive mldp** is only for the default multicast list; no per-ISID IMET-P2MP routes are supported.

The Backbone (B) -VPLS still uses Multicast Forwarding Information Bases (MFIBs) for ISIDs using IR.

If an ISID policy is configured in the B-VPLS, a range of ISIDs configured with **use-def-mcast** use the P2MP tree, and a range of ISIDs configured with **advertise-local** make the router advertise IMET-IR routes for the local ISIDs in the range.

PE-1 is configured with **root-and-leaf**. The configuration for B-VPLS and I-VPLS is as follows:

```
# On PE-1:
configure {
    service {
        vpls "B-VPLS 1000" {
            admin-state enable
            service-id 1000
            customer "1"
            service-mtu 2000
            pbb-type b-vpls
            pbb {
                source-bmac {
                    address 00:00:00:00:00:01
                }
            }
            bgp 1 {
            }
            bgp-evpn {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                evi 1000
                mpls 1 {
                    admin-state enable
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            provider-tunnel {
                inclusive {
                    admin-state enable
                    owner bgp-evpn-mpls
                    root-and-leaf true
                    mldp
                }
            }
            isid-policy {
                entry 10 {
                    advertise-local false
                    use-def-mcast true
                    range {
                        start 1001
                        end 2000
                    }
                }
            }
        }
        vpls "I-VPLS 1001" {
            admin-state enable
            service-id 1001
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS 1000" {
                    isid 1001
                }
            }
            sap 1/2/c3/1 {    # sap for ingress traffic from STC
            }
        }
    }
```

In this example, ISIDs in the range from 1001 to 2000 use the P2MP tree (**use-def-mcast**) and the router does not advertise the IMET-IR routes for the local ISIDs included in that range (**no advertise-local**). Any other local ISID advertises an IMET-IR and uses the MFIB to forward BUM packets to the remote EVPN-MPLS bindings created by IMET-IR routes.

The configuration on the leaf nodes PE-5, PE-6, and PE-7 is similar to the one for the root node, except for the absence of the **root-and-leaf** setting (which is default), as follows:

```
# On PE-5:
configure {
    service {
        vpls "B-VPLS 1000" {
            admin-state enable
            service-id 1000
            customer "1"
            service-mtu 2000
            pbb-type b-vpls
            pbb {
                source-bmac {
                    address 00:00:00:00:00:05
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 1000
                mpls 1 {
                    admin-state enable
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            provider-tunnel {
                inclusive {
                    admin-state enable
                    owner bgp-evpn-mpls
                    mldp
                }
            }
            isid-policy {
                entry 10 {
                    advertise-local false
                    use-def-mcast true
                    range {
                        start 1001
                        end 2000
                    }
                }
            }
        }
        vpls "I-VPLS 1001" {
            admin-state enable
            service-id 1001
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS 1000" {
                    isid 1001
                }
            }
            sap 1/2/c1/1:1001 {    # sap for egress traffic to VPLS 1001
            }
        }
    }
```

A VPLS-PMSI SDP is auto-created in the B-VPLS at the root node, as follows:

```
[/]
A:admin@PE-1# show service id 1000 sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
32767:4294967292 VplsPmsi not applicable  Up      Up        None      3
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The default multicast list for the B-VPLS 1000 can be retrieved on root node and leaf nodes, for instance for leaf node PE-5, as follows:

```
[/]
A:admin@PE-5# tools dump service id 1000 evpn-mpls default-multicast-list
-------------------------------------------------------------------------
TEP Address                           Egr Label
                                      Transport
-------------------------------------------------------------------------
192.0.2.1                             524283
                                      ldp
192.0.2.6                             524286
                                      ldp
192.0.2.7                             524281
                                      ldp
-------------------------------------------------------------------------
```

IGMP snooping can be enabled in the I-VPLS 1001 on all PEs, as follows:

```
configure {
    service {
        vpls "I-VPLS 1001" {
            igmp-snooping {
                admin-state enable
            }
```

After IGMP snooping is enabled, the multicast stream is not flooded anymore to any receivers until they send an IGMP report for the multicast stream.

On each PE, the logical interface B-EVPN-MPLS is added as a single IGMP snooping interface and treated as an mrouter, as follows:

```
[/]
A:admin@PE-5# show service id 1001 igmp-snooping base

===============================================================================
IGMP Snooping Base info for service 1001
===============================================================================
Admin State : Up
Querier     : 172.16.0.55 on SAP 1/2/c1/1:1001
SBD service : N/A
Evpn-proxy  : Disabled
-------------------------------------------------------------------------------
Port                      Oper MRtr Pim  Send Max  Max  Max   MVR       Num
Id                        Stat Port Port Qrys Grps Srcs Grp   From-VPLS Grps
                                                      Srcs
-------------------------------------------------------------------------------
b-evpn-mpls
                          Up   Yes  No   N/A  N/A  N/A  N/A   N/A       N/A
sap:1/2/c1/1:1001         Up   Yes  No   No   None None None  Local     0
===============================================================================
```

PE-5 has a receiver that sent an IGMP report for a multicast group in I-VPLS 1001 on SAP 1/2/c1/1:1001 and this SAP is an mrouter port. On PE-6, there is no receiver that sent IGMP reports; therefore, the only mrouter port corresponds to the B-EVPN-MPLS logical interface, as follows:

```
[/]
A:admin@PE-6# show service id 1001 igmp-snooping base

===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
IGMP Snooping Base info for service 1001
===============================================================================
Admin State : Up
Querier     : 172.16.0.55 on evpn-mpls
SBD service : N/A
Evpn-proxy  : Disabled
-------------------------------------------------------------------------------
Port                      Oper MRtr Pim  Send Max  Max  Max   MVR       Num
Id                        Stat Port Port Qrys Grps Srcs Grp   From-VPLS Grps
                                                  Srcs
-------------------------------------------------------------------------------
b-evpn-mpls
                          Up   Yes  No   N/A  N/A  N/A  N/A   N/A       N/A
sap:1/2/c1/1:1001         Up   No   No   No   None None None  Local     0
===============================================================================
```

PE-5 has a local mrouter 172.16.0.55 on SAP 1/2/c1/1:1001, as follows:

```
[/]
A:admin@PE-5# show service id 1001 igmp-snooping mrouters

===============================================================================
IGMP Snooping Multicast Routers for service 1001
===============================================================================
MRouter          Port Id                   Up Time        Expires   Version
-------------------------------------------------------------------------------
172.16.0.55      sap:1/2/c1/1:1001         0d 00:03:24    216s      3
-------------------------------------------------------------------------------
Number of mrouters: 1
===============================================================================
```

On PE-6, mrouter 172.16.0.55 is not local; therefore, the EVPN-MPLS logical interface is used, as follows:

```
[/]
A:admin@PE-6# show service id 1001 igmp-snooping mrouters

===============================================================================
IGMP Snooping Multicast Routers for service 1001
===============================================================================
MRouter          Port Id                   Up Time        Expires   Version
-------------------------------------------------------------------------------
172.16.0.55      evpn-mpls                 0d 00:03:25    215s      3
-------------------------------------------------------------------------------
Number of mrouters: 1
===============================================================================
```

Conclusion

Service providers are migrating their existing VPN services to EVPN and expect at least the same capabilities in EVPN, including the forwarding of BUM traffic. Ingress replication is a good mechanism for broadcast and unknown unicast traffic in EVPN networks, but not efficient for multicast applications. EVPN P2MP mLDP offers efficiency for multicast, using composite tunnels combining the benefits of P2MP mLDP and IR.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

PBB-Epipe

This chapter provides information about Provider Backbone Bridging (PBB) — Ethernet Virtual Leased Line in an MPLS-based network which is applicable to SR OS.

Topics in this chapter include:

Applicability

This chapter was initially written for SR OS Release 7.0.R5. The MD-CLI in the current edition corresponds to SR OS Release 20.10.R2. There are no specific prerequisites.

Overview

RFC 7041, *Extensions to VPLS PE model for Provider Backbone Bridging*, describes the PBB-VPLS model supported by SR OS. This model expands the VPLS PE model to support PBB as defined by the IEEE 802.1ah.

The PBB model is organized around a B-component (backbone instance) and an I-component (customer instance). In Nokia's implementation of the PBB model, the use of an Epipe as I-component is allowed for point-to-point services. Multiple I-VPLS and Epipe services can be all mapped to the same B-VPLS (backbone VPLS instance).

The use of Epipe scales the E-Line services because no MAC switching, learning, or replication is required in order to deliver the point-to-point service. All packets ingressing the customer SAP are PBB-encapsulated and unicasted through the B-VPLS tunnel using the backbone destination MAC of the remote PBB PE. All the packets egressing the B-VPLS destined for the Epipe are PBB de-encapsulated and forwarded to the customer SAP.

Some use cases for PBB-Epipe are:

- Get a more efficient and scalable solution for point-to-point services:
  - Up to 8K VPLS services per box are supported (including I-VPLS or B-VPLS) and using I-VPLS for point-to-point services takes VPLS resources as well as unnecessary customer MAC learning. A better solution is to connect a PBB-Epipe to a B-VPLS instance, where there is no customer MAC switching/learning.
- Take advantage of the pseudowire aggregation in the M:1 model:
  - Many Epipe services may use only a single service and set of pseudowires over the backbone.
- Have a uniform provisioning model for both point-to-point (Epipe) and multipoint (VPLS) services.
  - Using the PBB-Epipe, the core MPLS/pseudowire infrastructure does not need to be modified: the new Epipe inherits the existing pseudowire and MPLS structure already configured on the B-VPLS and there is no need for configuring new tunnels or pseudowire switching instances at the core.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Knowledge of the PBB-VPLS architecture and functionality on the service router family is assumed throughout this section. For additional information, see the relevant Nokia user documentation.

The following network setup will be used throughout the rest of the chapter.

Figure 220: Network topology

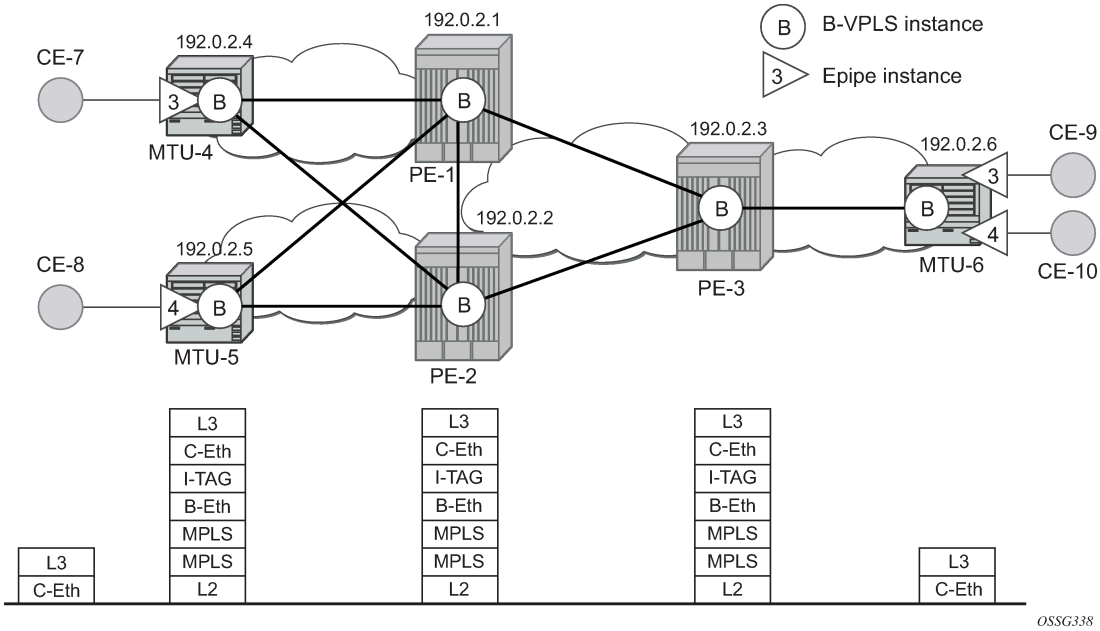

The setup consists of a three SR OS routers in the core (PE-1, PE-2, and PE-3) core and three Multi-Tenant Unit (MTU) nodes connected to the core. A backbone VPLS instance (B-VPLS 101) will be defined in all the six nodes, whereas two Epipe services will be defined as illustrated in Figure 220: Network topology (Epipe 3 in nodes MTU-4 and MTU-6, Epipe 4 in nodes MTU-5 and MTU-6). Those Epipe services will be multiplexed into the common B-VPLS 101, using the I-Service ID (ISID) field within the I-TAG as the demultiplexer field required at the egress MTU to differentiate each specific customer. I-VPLS and Epipe services can be mapped to the same B-VPLS.

The B-VPLS domain constitutes a H-VPLS network itself, with spoke-SDPs from the MTUs to the core PE layer. Active/standby (A/S) spoke-SDPs can be used from the MTUs to the PEs (like in the MTU-4 and MTU-5 cases) or single non-redundant spoke-SDPs (like MTU-6).

The protocol stack being used along the path between the CEs is represented in Figure 220: Network topology.

Configuration

This section describes all the relevant PBB-Epipe configuration tasks for the setup shown in Figure 220: Network topology. The appropriate B-VPLS and associated IP/MPLS configuration is out of the scope of this document. In this particular example, the following protocols will be configured beforehand in the core:

- ISIS-TE as IGP with all the interfaces being level-2. Alternatively, OSPF could have been used.
- RSVP-TE as the MPLS protocol to signal the transport tunnels.
- LSPs between core PEs will be fast re-route protected (facility bypass tunnels) whereas LSP tunnels between MTUs and PEs will not be protected.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- The protection between MTU-4, MTU-5 and PE-1, PE-2 will be based on the A/S pseudowire protection configured in the B-VPLS.
- BGP is configured for auto-discovery—BGP-AD (Layer 2 VPN family), because FEC 129 will be used to establish the pseudowires between PEs in the core (FEC 128 between MTU and PE nodes).

Once the IP/MPLS infrastructure is up and running, the service configuration tasks described in the following sections can be implemented.

PBB Epipe service configuration

In this particular example, the Epipes 3 and 4 are using the B-VPLS 101 in the core. The same B-VPLS which is multiplexing the Epipe services into a common service provider infrastructure can also be used to connect the I-VPLS instances existing in the network for multipoint services.

Figure 221: Setup detailed view

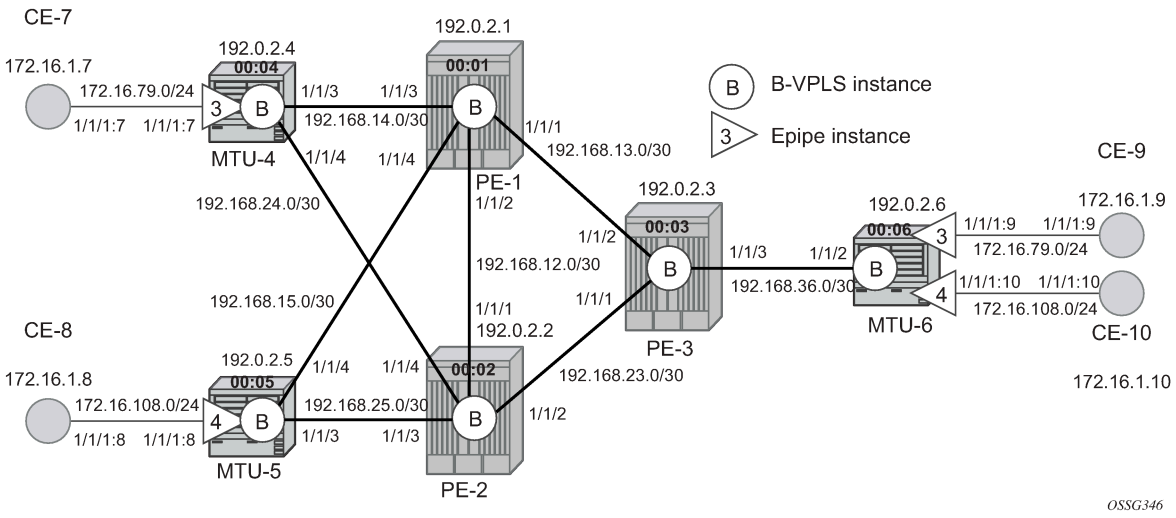

B-VPLS and PBB configuration

First, configure the B-VPLS instance that will carry the PBB traffic. There is no specific requirement on the B-VPLS to support Epipes. The following shows the B-VPLS configuration on MTU-4 and PE-1.

```
# on MTU-4:
configure {
    service {
        vpls "B-VPLS 101" {
            admin-state enable
            service-id 101
            customer "1"
            service-mtu 2000
            pbb-type b-vpls
            pbb {
                source-bmac {
                    address 00:04:04:04:04:04
                }
            }
            endpoint "core" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                suppress-standby-signaling false
            }
            spoke-sdp 41:101 {
                endpoint {
                    name "core"
                    precedence primary
                }
                stp {
                    admin-state disable
                }
            }
            spoke-sdp 42:101 {
                endpoint {
                    name "core"
                }
                stp {
                    admin-state disable
                }
            }
        }
```

```
# on PE-1:
configure {
    service {
        pw-template "PW1" {
            pw-template-id 1
            provisioned-sdp use
            split-horizon-group {
                name "CORE"
            }
        }
        vpls "B-VPLS 101" {
            admin-state enable
            service-id 101
            customer "1"
            service-mtu 2000
            pbb-type b-vpls
            pbb {
                source-bmac {
                    address 00:01:01:01:01:01
                }
            }
            bgp 1 {
                route-target {
                    export "target:65000:101"
                    import "target:65000:101"
                }
                pw-template-binding "PW1" {
                }
            }
            bgp-ad {
                admin-state enable
                vpls-id "65000:101"
            }
            spoke-sdp 14:101 {
            }
            spoke-sdp 15:101 {
            }
        }
```

- The B-VPLS service MTU must be at least 18 bytes greater than the Epipe MTU of the multiplexed instances. In this example, the I-VPLS instances will have the default service MTU (1514 bytes),

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

therefore, any MTU equal or greater than 1532 bytes must be configured. In this particular example, an MTU of 2000 bytes is configured in the B-VPLS instance throughout the network.

- The source B-MAC is the MAC that will be used as a source when the PBB traffic is originated from that node. It is possible to configure a source B-MAC per B-VPLS instance (if there are more than one B-VPLS) or a common source B-MAC that will be shared by all the B-VPLS instances in the node. A common B-MAC is configured as follows:

```
# on MTU-4:
configure {
    service {
        pbb {
            source-bmac {
                address 00:04:04:04:04:04
            }
        }
```

The following considerations will be taken into account when configuring the B-VPLS:

- B-VPLS SAPs:
  - Ethernet null, dot1q, and qinq encapsulations are supported.
  - Default SAP types are blocked in the CLI for the B-VPLS SAP.
- B-VPLS SDPs:
  - For MPLS, both mesh and spoke-SDPs with split-horizon groups are supported.
  - Similar to regular pseudowire, the outgoing PBB frame on an SDP (for example, Bpseudowire) contains a BVID q-tag only if the pseudowire type is Ethernet VLAN (vc-type=vlan). If the pseudowire type is Ethernet (vc-type=ether), the BVID q-tag is stripped before the frame goes out.
  - BGP-AD is supported in the B-VPLS, therefore, spoke-SDPs in the B-VPLS can be signaled using FEC 128 or FEC 129. In this example, BGP-AD and FEC 129 are used. A split-horizon group has been configured to emulate the behavior of mesh SDPs in the core.
- While Multiple MAC Registration Protocol (MMRP) is useful to optimize the flooding in the B-VPLS domain and build a flooding tree on a per I-VPLS basis, it does not have any effect for Epipes because the destination B-MAC used for Epipes is always the destination B-MAC configured in the Epipe and never the group B-MAC corresponding to the ISID.
- If a local Epipe instance is associated with the B-VPLS, local frames originated or terminated on local Epipe(s) are PBB encapsulated or de-encapsulated using the PBB Etype provisioned under the related port or SDP component.

By default, the PBB Etype is 0x88e7 (which is the standard one defined in the 802.1ah, indicating that there is an I-TAG in the payload) but this PBB Etype can be changed if required due to interoperability reasons. This is the way to change it at port and/or SDP level:

```
[ex:configure port 1/1/3 ethernet]
A:admin@MTU-4# pbb-etype ?

 pbb-etype <number>
 <number> - <0x600..0xffff>
 Default  - 35047

    Ethertype for PBB encapsulation on the Ethernet port
```

```
[ex:configure service sdp 41]
A:admin@MTU-4# pbb-etype ?
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
 pbb-etype <number>
 <number> - <0x600..0xffff>
 Default  - 0x88E7

    Ethertype used in frames sent out on this SDP when VC type is 'vlan' for
    Provider Backbone Bridging frames
    as 0xXXYY with range 0x0600-0xFFFF.
```

The following commands are useful to check the actual PBB Etype.

```
[]
A:admin@MTU-4# show service sdp 41 detail | match PBB
Bw BookingFactor     : 100                   PBB Etype          : 0x88e7
```

```
[]
A:admin@MTU-4# show port 1/1/3 | match PBB
PBB Ethertype      : 0x88e7
```

Before configuring the Epipe itself, the operator can optionally configure MAC names under the PBB context. MAC names will simplify the Epipe provisioning later on and in case of any change on the remote node MAC address, only one configuration modification is required as opposed as one change per affected Epipe (potentially thousands of Epipes which are terminated onto the same remote node). The MAC names are configured in the service PBB CLI context:

```
[ex:configure service pbb]
A:admin@MTU-4# mac ?

 [name] <string>
 <string> - <1..32 characters>

    MAC address name
```

The MAC names of the MTUs are configured on all nodes, as follows:

```
# on all nodes:
configure {
    service {
        pbb {
            mac "MTU-4" {
                address 00:04:04:04:04:04
            }
            mac "MTU-5" {
                address 00:05:05:05:05:05
            }
            mac "MTU-6" {
                address 00:06:06:06:06:06
            }
```

It is not required to configure a node with its own MAC address, so on MTU-4, the line defining the mac-name MTU-4 can be omitted.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Epipe configuration

Once the common B-VPLS is configured, the next step is the provisioning of the customer Epipe instances. For PBB-Epipes, the I-component or Epipe is composed of an I-SAP and a PBB tunnel endpoint which points to the backbone destination MAC address (B-DA).

The following outputs show the relevant CLI configuration for the two Epipe instances represented in Figure 221: Setup detailed view. The Epipe instances are configured on the MTU devices, whereas the core PEs are kept as customer-unaware nodes.

Epipes 3 and 4 are configured on MTU-6 as follows:

```
# on MTU-6:
configure {
    service {
        epipe "Epipe 3" {
            admin-state enable
            description "pbb epipe number 3"
            service-id 3
            customer "1"
            pbb {
                tunnel {
                    backbone-vpls-service-name "B-VPLS 101"
                    isid 3
                    backbone-dest-mac-name "MTU-4"
                }
            }
            sap 1/1/1:9 {
            }
        }
        epipe "Epipe 4" {
            admin-state enable
            description "pbb epipe number 4"
            service-id 4
            customer "1"
            pbb {
                tunnel {
                    backbone-vpls-service-name "B-VPLS 101"
                    isid 4
                    backbone-dest-mac-name "MTU-5"
                }
            }
            sap 1/1/1:10 {
            }
        }
```

The following shows the Epipe configuration on MTU-4 and MTU-5.

```
# on MTU-4:
configure {
    service {
        epipe "Epipe 3" {
            admin-state enable
            description "pbb epipe number 3"
            service-id 3
            customer "1"
            pbb {
                tunnel {
                    backbone-vpls-service-name "B-VPLS 101"
                    isid 3
                    backbone-dest-mac-name "MTU-6"
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
            sap 1/1/1:7 {
            }
        }
```

```
# on MTU-5:
configure {
    service {
        epipe "Epipe 4" {
            admin-state enable
            description "pbb epipe number 4"
            service-id 4
            customer "1"
            pbb {
                tunnel {
                    backbone-vpls-service-name "B-VPLS 101"
                    isid 4
                    backbone-dest-mac-name "MTU-6"
                }
            }
            sap 1/1/1:8 {
            }
        }
```

All Ethernet SAPs supported by a regular Epipe are also supported in the PBB Epipe. spoke-SDPs are not supported in PBB-Epipes, for example, no spoke-SDP is allowed when PBB tunnels are configured on the Epipe.

The PBB tunnel links the SAP configured to the B-VPLS 101 existing in the core. The following parameters are accepted in the PBB tunnel configuration:

```
[ex:configure service epipe "Epipe 4" pbb]
A:admin@MTU-5# tunnel ?

 tunnel

 Immutable fields     - backbone-vpls-service-name, isid

 apply-groups          - Apply a configuration group at this level
 apply-groups-exclude  - Exclude a configuration group at this level
 backbone-vpls-        ^ Backbone VPLS service
  service-name
 isid                  ^ Service instance ID

 Mandatory choice: backbone-mac
 backbone-dest-mac    :- Backbone Destination MAC address
 backbone-dest-mac-   :- Name for backbone Destination MAC address
  name
```

Where:

- The backbone VPLS service name matches the B-VPLS name, in this case, "B-VPLS 101".
- The backbone destination can be configured by a MAC name (as in the configuration example, with MAC name "MTU-6") or the MAC address itself. It is recommended to use MAC names, as explained in the previous section.
- The ISID corresponds to the Epipe service ID and must be specified.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Flood avoidance in PBB-Epipes

As already discussed in the previous section, when provisioning a PBB Epipe, the remote backbone destination MAC (MAC name or MAC address) must be explicitly configured on the PBB tunnel so that the ingress PBB node can build the 802.1ah encapsulation.

If the configured remote backbone destination MAC address is not known in the local FDB, the Epipe customer frames will be 802.1ah encapsulated and flooded into the B-VPLS until the MAC address is learned. As previously stated, MMRP does not help to minimize the flooding because the PBB Epipes always use the configured backbone destination MAC for flooding traffic as opposed to the group B-MAC derived from the ISID.

Flooding could be indefinably prolonged in the following cases:

- Configuration mistake of the backbone destination MAC (either MAC name or MAC address). The service will not work, but the operator will not detect the mistake, because the customer traffic is not dropped at the source node. Every single frame is turned into an unknown unicast PBB frame and therefore flooded into the B-VPLS domain.
- Change the backbone source MAC in the remote PE B-VPLS instance.
- There is only unidirectional traffic in the Epipe service. In this case, the backbone destination MAC address will never be learned in the local FIB and the frames will always be flooded into the B-VPLS domain.
- The remote node owning the backbone destination MAC simply goes down.

In any of those cases, the operator can easily check whether the PBB Epipe is flooding into the B-VPLS domain, just by looking at the flood flag in the following command output:

```
[]
A:admin@MTU-4# show service id 3 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 3                   Vpn Id            : 0
Service Type      : Epipe
---snip---

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                               Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/1:7                              q-tag        9000    9000    Up   Up
-------------------------------------------------------------------------------
PBB Tunnel Point
-------------------------------------------------------------------------------
B-vpls     Backbone-dest-MAC Isid     AdmMTU OperState Flood Oper-dest-MAC
-------------------------------------------------------------------------------
101        MTU-6             3        2000   Up        Yes   00:06:06:06:06:06
-------------------------------------------------------------------------------
Last Status Change: 01/11/2021 16:18:26
Last Mgmt Change  : 01/11/2021 16:18:26
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

In this particular example, the PBB Epipe 3 is flooding into the B-VPLS 101, as the flood flag indicates. The operator can also confirm that the operational destination B-MAC for the PBB tunnel, MTU-6, has not been learned in the B-VPLS FDB:

```
[]
A:admin@MTU-4# show service id 101 fdb pbb

===============================================================================
Forwarding Database, b-Vpls Service 101
===============================================================================
MAC               Source-Identifier     iVplsMACs  Epipes     Type/Age
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

In small B-VPLS environments (up to 20 B-VPLSs, each with 10 MC-LAGs), it is possible to configure the PBB V-VPLS MAC notification mechanism to send notification messages at regular intervals (using the renotify parameter), rather than being only event-driven. This can avoid flooding into the B-VPLS.

Flooding cases 1 and 2 — Wrong backbone destination MAC

Flooding cases 1 and 2 should be fixed after detecting the flooding (see previous commands) and checking the FDBs and PBB tunnel configurations.

Flooding case 3 — Unidirectional traffic: virtual MEP and CCM configuration

For flooding case 3 (unidirectional traffic), Nokia recommends the use of ETH-CFM (802.1ag/Y.1731 Connectivity Fault Management) virtual Maintenance End Points (MEPs). By defining a virtual MEP per node terminating a PBB Epipe, configuring the MEP MAC address to be the source BMAC value and activating continuity check messages (CCM), a twofold effect is achieved:

- The PBB tunnel backbone destination MAC address will always be learned at the local FDB, as long as the remote virtual MEP is active and sending CC messages. As a result, there will not be flooding even if we have unidirectional traffic.
- An automatic proactive OAM mechanism exists to detect failures on remote nodes, which ultimately cause unnecessary flooding in the B-VPLS domain.

In the following network example, the virtual MEPs in B-VPLS 101: MEP4, MEP5, and MEP6 are configured.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 222: Virtual MEPs for flooding avoidance

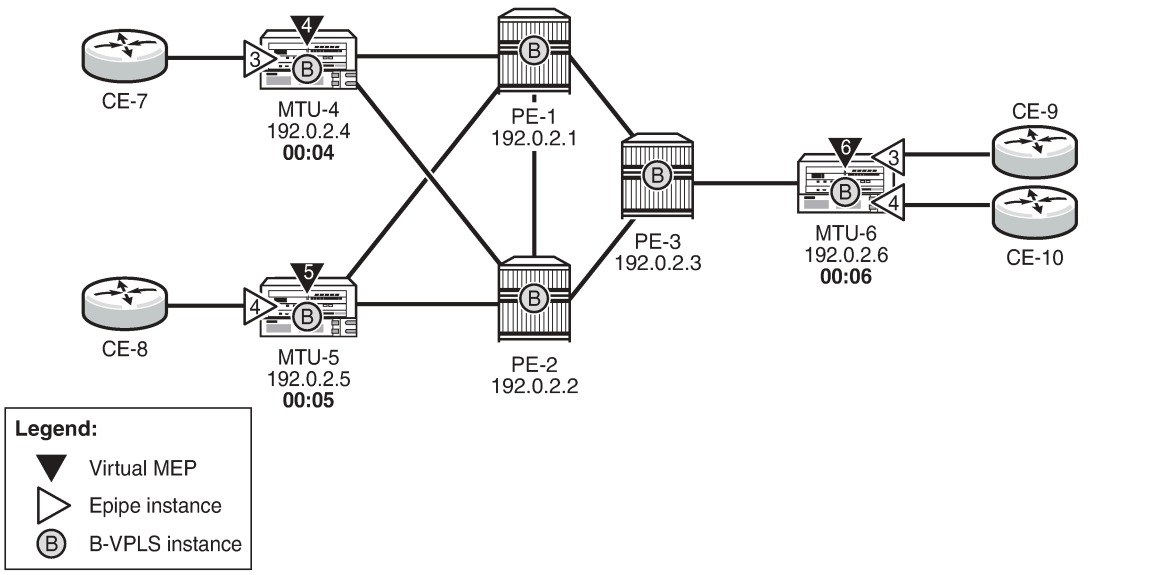

The following configuration example uses MTU-4. First, the general ETH-CFM configuration is made, as follows:

```
# on MTU-4:
configure {
    eth-cfm {
        domain "domain-1" {
            level 3
            name "domain-1"
            md-index 1
            association "assoc-1" {
                icc-based "B-VPLS-000101"
                ma-index 1
                bridge-identifier "B-VPLS 101" {
                }
                remote-mep 5 {
                }
                remote-mep 6 {
                }
            }
```

Then the actual virtual MEP configuration is made:

```
# on MTU-4:
configure {
    service {
        vpls "B-VPLS 101" {
            eth-cfm {
                mep md-admin-name "domain-1" ma-admin-name "assoc-1" mep-id 4 {
                    admin-state enable
                    mac-address 00:04:04:04:04:04
                    ccm true
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
```

The MAC address configured for the MEP4 matches the MAC address configured as the **source-bmac** on MTU-4, which is the **backbone-destination-mac** configured on the Epipe 3 PBB tunnel on MTU-6. The source-BMAC address on MTU-4 is 00:04:04:04:04:04, as follows:

```
# on MTU-4:
configure {
    service {
        pbb {
            source-bmac {
                address 00:04:04:04:04:04
            }
            mac "MTU-4" {
                address 00:04:04:04:04:04
            }
            mac "MTU-5" {
                address 00:05:05:05:05:05
            }
            mac "MTU-6" {
                address 00:06:06:06:06:06
            }
        }
```

In Epipe 3 on MTU-6, the configured backbone destination MAC name is MAC name "MTU-4", which corresponds to MAC address 00:04:04:04:04:04, as follows:

```
# on MTU-6:
configure {
    service {
        pbb {
            source-bmac {
                address 00:06:06:06:06:06
            }
            mac "MTU-4" {
                address 00:04:04:04:04:04
            }
            mac "MTU-5" {
                address 00:05:05:05:05:05
            }
            mac "MTU-6" {
                address 00:06:06:06:06:06
            }
        }
        epipe "Epipe 3" {
            admin-state enable
            description "pbb epipe number 3"
            service-id 3
            customer "1"
            pbb {
                tunnel {
                    backbone-vpls-service-name "B-VPLS 101"
                    isid 3
                    backbone-dest-mac-name "MTU-4"
                }
            }
            sap 1/1/1:9 {
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Once MEP4 has been configured, check that MTU-6 is receiving CC messages from MEP4 with the following command:

```
[]
A:admin@MTU-6# show eth-cfm mep 6 domain 1 association 1 all-remote-mepids

================================================================================
Eth-CFM Remote-Mep Table
================================================================================
R-mepId AD Rx CC RxRdi Port-Tlv If-Tlv Peer Mac Addr     CCM status since
--------------------------------------------------------------------------------
4          True  False Absent   Absent 00:04:04:04:04:04 01/11/2021 16:26:59
5          True  False Absent   Absent 00:05:05:05:05:05 01/11/2021 16:26:59
================================================================================
Entries marked with a 'T' under the 'AD' column have been auto-discovered.
```

As a result of the CC messages coming from MEP4, the MTU-4 MAC is permanently learned in the B-VPLS 101 FDB on node MTU-6 and no flooding takes place. The following output shows that the flooding flag is not set.

```
[]
A:admin@MTU-6# show service id 3 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 3                   Vpn Id            : 0
Service Type      : Epipe
---snip---

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                              Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/1:9                             q-tag        9000    9000    Up   Up
-------------------------------------------------------------------------------
PBB Tunnel Point
-------------------------------------------------------------------------------
B-vpls     Backbone-dest-MAC Isid     AdmMTU OperState Flood Oper-dest-MAC
-------------------------------------------------------------------------------
101        MTU-4             3        2000   Up        No    00:04:04:04:04:04
-------------------------------------------------------------------------------
Last Status Change: 01/11/2021 16:18:43
Last Mgmt Change  : 01/11/2021 16:18:43
===============================================================================
```

Flooding case 4 — Remote node failure

If the node owner of the backbone destination MAC fails or gets isolated, the node where the PBB Epipe is initiated will not detect the failure; that is, if MTU-4 fails, the Epipe 3 remote end will also fail, but MTU-6 will not detect the failure and, as a result of that, MTU-6 will flood the traffic to the network (flooding will occur after MTU-4 MAC is removed from the B-VPLS FDBs, due to either the B-VPLS flushing mechanisms or aging).

In order to avoid/reduce flooding in this case, the following mechanisms are recommended:

- Provision virtual MEPs in the B-VPLS instances terminating PBB Epipes, as already explained. This will guarantee there is no unknown B-MAC unicast being flooded under normal operation.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- CCM timers should be provisioned based on how long the service provider is willing to accept flooding.

```
[ex:configure eth-cfm domain "domain-1" association "assoc-1"]
A:admin@MTU-6# ccm-interval ?

 ccm-interval <keyword>
 <keyword> - (10ms|100ms|1s|10s|60s|600s)
 Default   - 10s

    CCM transmission interval for all MEPs in the association
```

- It is possible to provision **discard-unknown** in the B-VPLS, so that flooded traffic due to the destination MAC being unknown in the B-VPLS is discarded immediately. This can be configured on the PEs and the MTUs. On the MTUs, it is important to configure this in conjunction with the CC messages from the virtual MEPs to ensure that the remote B-MACs are learned in both directions. If, for any reason, the remote B-MACs are not in the MTU B-VPLS, no traffic will be forwarded at all on the PBB-Epipe.

```
configure {
    service {
        vpls "B-VPLS 101" {
            fdb {
                discard-unknown true
            }
```

  As soon as the MTU node recovers, it will start sending CC messages and the backbone MAC address will be learned on the backbone nodes and MTU nodes again.

With the recommended configuration in place, in case MTU-4 fails, the backbone destination MAC configured on the PBB tunnel for Epipe 3 on MTU-6 will be removed from the B-VPLS 101 on all the nodes (either by MAC flush mechanisms on the B-VPLS or by aging). From that point on, traffic originated from CE-9 will be discarded at MTU-6 and won't be flooded further.

As soon as MTU-4 comes back up, MEP4 will start sending CCM and as such the MTU-4 MAC will be learned throughout the B-VPLS 101 domain and in particular in PE-1, PE-3, and MTU-6 (CCM PDUs use a multicast address). From the moment MTU-4 MAC is known on the backbone nodes and MTU-6, the traffic will not be discarded any more, but forwarded to MTU-4.

PBB-Epipe show commands

The following commands can help to check the PBB Epipe configuration and their related parameters.

For the B-VPLS service:

```
[]
A:admin@MTU-4# show service id 101 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 101                Vpn Id            : 0
Service Type      : b-VPLS
MACSec enabled    : no
Name              : B-VPLS 101
Description       : (Not Specified)
Customer Id       : 1                  Creation Origin   : manual
Last Status Change: 01/11/2021 16:10:35
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Last Mgmt Change  : 01/11/2021 16:32:31
Etree Mode        : Disabled
Admin State       : Up                  Oper State        : Up
MTU               : 2000
SAP Count         : 0                   SDP Bind Count    : 2
Snd Flush on Fail : Disabled            Host Conn Verify  : Disabled
SHCV pol IPv4     : None
Propagate MacFlush: Disabled            Per Svc Hashing   : Disabled
Allow IP Intf Bind: Disabled
Fwd-IPv4-Mcast-To*: Disabled            Fwd-IPv6-Mcast-To*: Disabled
Mcast IPv6 scope  : mac-based
Temp Flood Time   : Disabled            Temp Flood        : Inactive
Temp Flood Chg Cnt: 0
SPI load-balance  : Disabled
TEID load-balance : Disabled
Src Tep IP        : N/A
Vxlan ECMP        : Disabled
MPLS ECMP         : Disabled
VSD Domain        : <none>
Oper Backbone Src : 00:04:04:04:04:04
Use SAP B-MAC     : Disabled
i-Vpls Count      : 0
Epipe Count       : 1
Use ESI B-MAC     : Disabled

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                              Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sdp:41:101 S(192.0.2.1)                  Spok         8000    8000    Up   Up
sdp:42:101 S(192.0.2.2)                  Spok         8000    8000    Up   Up
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

For the Epipe service:

```
[]
A:admin@MTU-4# show service id 3 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 3                   Vpn Id            : 0
Service Type      : Epipe
MACSec enabled    : no
Name              : Epipe 3
Description       : pbb epipe number 3
Customer Id       : 1                   Creation Origin   : manual
Last Status Change: 01/11/2021 16:18:26
Last Mgmt Change  : 01/11/2021 16:18:26
Test Service      : No
Admin State       : Up                  Oper State        : Up
MTU               : 1514
Vc Switching      : False
SAP Count         : 1                   SDP Bind Count    : 0
Per Svc Hashing   : Disabled
Vxlan Src Tep Ip  : N/A
Force QTag Fwd    : Disabled
Oper Group        : <none>
-------------------------------------------------------------------------------
Service Access & Destination Points
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
Identifier                              Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/1:7                             q-tag        9000    9000    Up   Up

-------------------------------------------------------------------------------
PBB Tunnel Point
-------------------------------------------------------------------------------
B-vpls     Backbone-dest-MAC Isid     AdmMTU OperState Flood Oper-dest-MAC
-------------------------------------------------------------------------------
101        MTU-6             3        2000   Up        No    00:06:06:06:06:06
-------------------------------------------------------------------------------
Last Status Change: 01/11/2021 16:18:26
Last Mgmt Change  : 01/11/2021 16:18:26
===============================================================================
```

The following command shows all the Epipe instances multiplexed into a particular B-VPLS and its status.

```
[]
A:admin@MTU-4# show service id 101 epipe

===============================================================================
Related Epipe services for b-Vpls service 101
===============================================================================
Epipe SvcId         Oper ISID           Admin               Oper
-------------------------------------------------------------------------------
3                   3                   Up                  Up
-------------------------------------------------------------------------------
Number of Entries : 1
-------------------------------------------------------------------------------
===============================================================================
```

The following command shows the local virtual MEPs configured on MTU-4:

```
[]
A:admin@MTU-4# show eth-cfm cfm-stack-table all-virtuals
===============================================================================
CFM Stack Table Defect Legend:
R = Rdi, M = MacStatus, C = RemoteCCM, E = ErrorCCM, X = XconCCM
A = AisRx, L = CSF LOS Rx, F = CSF AIS/FDI rx, r = CSF RDI rx
G = receiving grace PDU (MCC-ED or VSM) from at least one peer

===============================================================================
CFM Virtual Stack Table
===============================================================================
Service            Lvl Dir Md-index   Ma-index  MepId  Mac-address     Defect G
-------------------------------------------------------------------------------
101                  3 U          1          1    4 00:04:04:04:04:04 ------- -
===============================================================================
```

The following command shows all the information related to the remote MEPs configured in the association, for example, the remote virtual MEPs configured in MTU-5 and MTU-6:

```
[]
A:admin@MTU-4# show eth-cfm mep 4 domain 1 association 1 all-remote-mepids

===============================================================================
Eth-CFM Remote-Mep Table
===============================================================================
R-mepId AD Rx CC RxRdi Port-Tlv If-Tlv Peer Mac Addr     CCM status since
-------------------------------------------------------------------------------
5          True  False Absent   Absent 00:05:05:05:05:05 01/11/2021 16:26:38
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
6          True  False Absent   Absent 00:06:06:06:06:06 01/11/2021 16:26:38
===============================================================================
Entries marked with a 'T' under the 'AD' column have been auto-discovered.
```

The following command shows the detail information and status of the local virtual MEP configured in MTU-4:

```
[]
A:admin@MTU-4# show eth-cfm mep 4 domain 1 association 1
===============================================================================
Eth-Cfm MEP Configuration Information
===============================================================================
Md-index         : 1                        Direction         : Up
Ma-index         : 1                        Admin             : Enabled
MepId            : 4                        CCM-Enable        : Enabled
SvcId            : 101
Description      : (Not Specified)
FngAlarmTime     : 0                        FngResetTime      : 0
FngState         : fngReset                 ControlMep        : False
LowestDefectPri  : macRemErrXcon            HighestDefect     : none
Defect Flags     : None
Mac Address      : 00:04:04:04:04:04        Collect LMM Stats : disabled
LMM FC Stats     : None
LMM FC In Prof   : None
TxAis            : noTransmit               TxGrace           : noTransmit
Facility Fault   : disabled
CcmLtmPriority   : 7                        CcmPaddingSize    : 0 octets
CcmTx            : 88                       CcmSequenceErr    : 0
CcmTxIfStatus    : Absent                   CcmTxPortStatus   : Absent
CcmTxRdi         : False                    CcmTxCcmStatus    : transmit
CcmIgnoreTLVs    : (Not Specified)
Fault Propagation: disabled                 FacilityFault     : n/a
MA-CcmInterval   : 10                       MA-CcmHoldTime    : 0ms
MA-Primary-Vid   : Disabled
Eth-1Dm Threshold: 3(sec)                   MD-Level          : 3
Eth-1Dm Last Dest: 00:00:00:00:00:00
Eth-Dmm Last Dest: 00:00:00:00:00:00
Eth-Ais          : Disabled
Eth-Ais Tx defCCM: allDef
Eth-Tst          : Disabled
Eth-CSF          : Disabled

Eth-Cfm Grace Tx : Enabled                  Eth-Cfm Grace Rx  : Enabled
Eth-Cfm ED Tx    : Disabled                 Eth-Cfm ED Rx     : Enabled
Eth-Cfm ED Rx Max: 0
Eth-Cfm ED Tx Pri: CcmLtmPri (7)

Eth-BNM Receive  : Disabled                 Eth-BNM Rx Pacing : 5

Redundancy:
    MC-LAG State : n/a

CcmLastFailure Frame:
    None

XconCcmFailure Frame:
    None
===============================================================================
```

When there is a failure on a remote Epipe node, as described, the source node keeps sending traffic. The 802.1ag/Y.1731 virtual MEP configured can help to detect and troubleshoot the problem. For instance,

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

when a failure happens in MTU-6 (node goes down or the B-VPLS instance is disabled), the virtual MEP show commands will show the following information:

```
# on MTU-6:
configure {
    service {
        vpls "B-VPLS 101" {
            admin-state disable
        }
```

```
[]A:admin@MTU-4# show eth-cfm mep 4 domain 1 association 1
================================================================================
Eth-Cfm MEP Configuration Information
================================================================================
Md-index         : 1                        Direction         : Up
Ma-index         : 1                        Admin             : Enabled
MepId            : 4                        CCM-Enable        : Enabled
SvcId            : 101
Description      : (Not Specified)
FngAlarmTime     : 0                        FngResetTime      : 0
FngState         : fngDefectReported        ControlMep        : False
LowestDefectPri  : macRemErrXcon            HighestDefect     : defRemoteCCM
Defect Flags     : bDefRDICCM bDefRemoteCCM
Mac Address      : 00:04:04:04:04:04        Collect LMM Stats : disabled
LMM FC Stats     : None
LMM FC In Prof   : None
TxAis            : noTransmit               TxGrace           : noTransmit
Facility Fault   : disabled
CcmLtmPriority   : 7                        CcmPaddingSize    : 0 octets
CcmTx            : 128                      CcmSequenceErr    : 0
CcmTxIfStatus    : Absent                   CcmTxPortStatus   : Absent
CcmTxRdi         : True                     CcmTxCcmStatus    : transmit
CcmIgnoreTLVs    : (Not Specified)
Fault Propagation: disabled                 FacilityFault     : n/a
MA-CcmInterval   : 10                       MA-CcmHoldTime    : 0ms
MA-Primary-Vid   : Disabled
Eth-1Dm Threshold: 3(sec)                   MD-Level          : 3
Eth-1Dm Last Dest: 00:00:00:00:00:00
Eth-Dmm Last Dest: 00:00:00:00:00:00
Eth-Ais          : Disabled
Eth-Ais Tx defCCM: allDef
Eth-Tst          : Disabled
Eth-CSF          : Disabled

Eth-Cfm Grace Tx : Enabled                  Eth-Cfm Grace Rx  : Enabled
Eth-Cfm ED Tx    : Disabled                 Eth-Cfm ED Rx     : Enabled
Eth-Cfm ED Rx Max: 0
Eth-Cfm ED Tx Pri: CcmLtmPri (7)

Eth-BNM Receive  : Disabled                 Eth-BNM Rx Pacing : 5

Redundancy:
    MC-LAG State : n/a

CcmLastFailure Frame:
    None

XconCcmFailure Frame:
    None
================================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The bDefRemoteCCMdefect flag clearly shows that there is a remote MEP in the association which has stopped sending CCMs. In order to find out which node is affected, see the following output:

```
[]
A:admin@MTU-4# show eth-cfm mep 4 domain 1 association 1 all-remote-mepids

===============================================================================
Eth-CFM Remote-Mep Table
===============================================================================
R-mepId AD Rx CC RxRdi Port-Tlv If-Tlv Peer Mac Addr     CCM status since
-------------------------------------------------------------------------------
5          True  True  Absent   Absent 00:05:05:05:05:05 01/11/2021 16:26:38
6          False False Absent   Absent 00:00:00:00:00:00 01/11/2021 16:43:56
===============================================================================
Entries marked with a 'T' under the 'AD' column have been auto-discovered.
```

CCMs are no longer received from virtual MEP 6 (the one defined in MTU-6) since 01/11/2021 16:43:56. This conveys which node has failed and when it failed.

Conclusion

Point-to-Point Ethernet services can use the same operational model followed by PBB VPLS for multipoint services. In other words, Epipes can be linked to the same B-VPLS domain being used by I-VPLS instances and use the existing H-VPLS network infrastructure in the core. The use of PBB Epipes reduces dramatically the number of services and pseudowires in the core and therefore allows the service provider to scale the number of E-Line services in the network.

The example used in this document shows the configuration of the PBB Epipes as well as all the related features which are required for this environment. Show commands have also been suggested so that the operator can verify and troubleshoot the service.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

PBB-EVPN ISID-based CMAC Flush

This chapter provides information about PBB-EVPN ISID-based CMAC Flush.

Topics in this chapter include:

Applicability

This chapter was initially written for SR OS Release 15.0.R4, but the MD-CLI in the current edition is based on SR OS Release 21.2.R2. PBB-EVPN ISID-based CMAC flush is supported on the following objects in an I-VPLS:

- SAPs in a BGP multi-homing site (no Ethernet Segment (ES))-supported in SR OS Release 14.0.R4, and later
- SAPs in ESs or virtual ESs (vESs)-SR OS Release 15.0.R1, and later
- Spoke-SDPs (that may be part of an ES/vES or not)-SR OS Release 15.0.R4, and later.

Chapter EVPN for PBB over MPLS (PBB-EVPN) is prerequisite reading.

Overview

Figure 223: CMAC flush when SAP in BGP multi-homing site fails shows an example topology with PBB-EVPN where a CMAC flush is triggered after a SAP in a BGP multi-homing site fails.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 223: CMAC flush when SAP in BGP multi-homing site fails

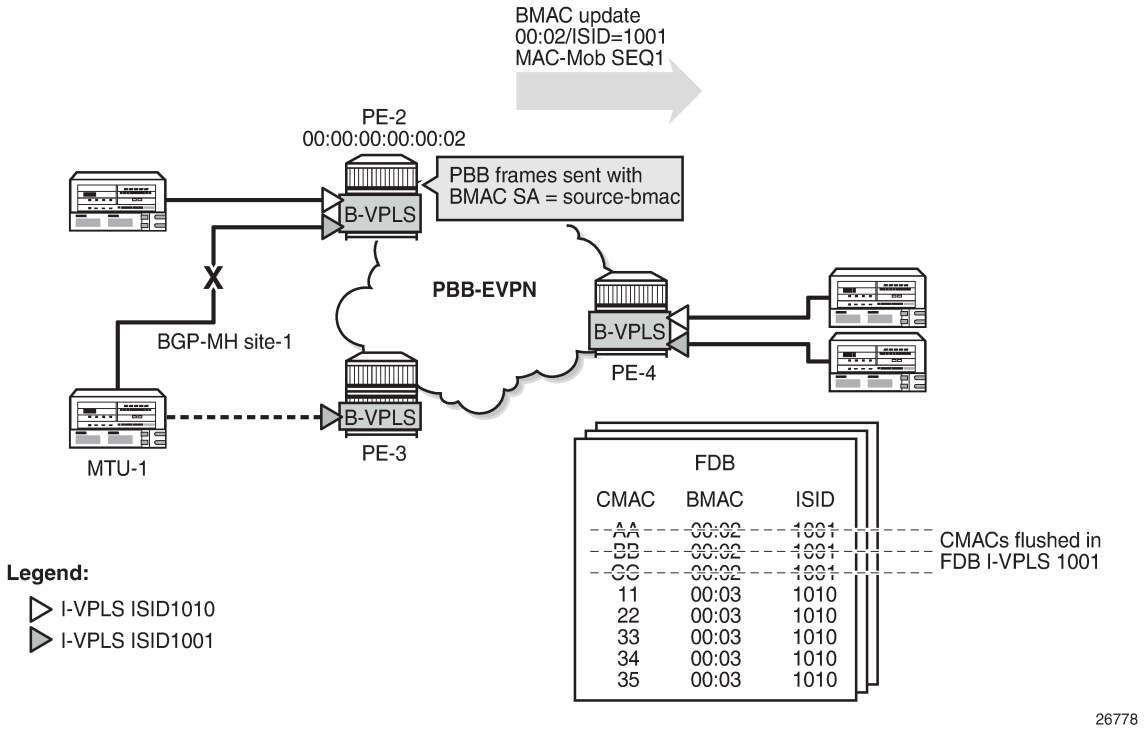

I-VPLS 1001 is configured in PE-2 and PE-3 with **pbb>i-vpls-mac-flush>bgp-evpn>send-to-bvpls true** and connected to MTU-1. In the example, the SAP goes operationally down in I-VPLS 1001 on PE-2. To speed up convergence without flushing CMAC addresses in other I-VPLS services, PE-2 sends a BGP-EVPN BMAC route for ISID 1001 with increased sequence number to trigger a MAC-flush for I-VPLS 1001 on the remote PEs. All CMAC addresses in the FDB for other I-VPLS services, such as I-VPLS 1010 in this example, will be preserved. When PE-4 needs to send traffic to one of the flushed CMAC addresses in I-VPLS 1001, it will flood the frames until the CMAC address is learned again (via PE-3).

When SAPs or SDP-bindings-associated with ESs, vESs, or BGP-MH sites-in an I-VPLS service fail, a BGP-EVPN BMAC route (route type 2) can trigger an ISID-based CMAC flush on the remote PEs. For the CMAC addresses to be flushed from the FDB of the I-VPLS, the existing EVPN BMAC routes will be used with the Ethernet tag equal to the ISID. Figure 224: EVPN BMAC route with ISID indication shows the EVPN BMAC route with ISID indication (BMAC/ISID). A BMAC/ISID update may trigger a selective MAC-flush for a specific I-VPLS, whereas a BMAC/0 update (BMAC/ISID route where ISID=0) may trigger a MAC-flush for all I-VPLS services. This procedure is based on *draft-snr-bess-pbb-evpn-isid-cmacflush*.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 224: EVPN BMAC route with ISID indication

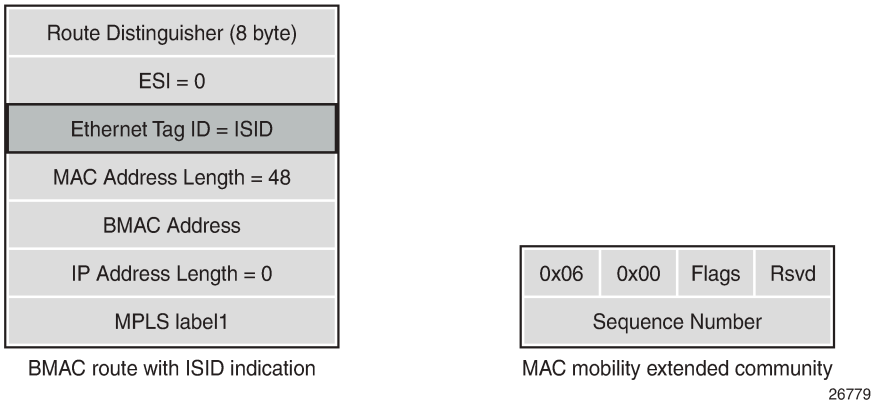

By default, ISID-based CMAC flush is disabled: no I-VPLS will send a B-VPLS EVPN flush message and no B-VPLS will accept any I-VPLS EVPN flush messages. The router only installs CMAC entries corresponding to a zero Ethernet tag and ignores non-zero Ethernet tag MAC routes. However, when the B-VPLS is configured to accept BMAC/ISID routes, non-zero Ethernet tag BMAC routes can be processed for CMAC flush. The CMAC flush trigger will be an EVPN BMAC/ISID route with a sequence number that is higher than before. The receiving PE will then flush all CMACs associated with this BMAC address in the I-VPLS.

The first time that a BMAC/ISID route is received, it is added to the database as a baseline. It does not cause a CMAC flush. Only subsequent BMAC/ISID updates with increased sequence number or withdrawals will cause CMAC flush.

The following command shows that B-VPLS 1000 does not accept any I-VPLS EVPN flush messages. This is the default behavior.

```
[/]
A:admin@PE-2# show service id 1000 bgp-evpn | match "Accept IVPLS Flush"
Accept IVPLS Flush : Disabled
```

At the receiving node, B-VPLS 1000 will accept BMAC/ISID routes when the following command is configured:

```
# on PE-2:
configure {
    service {
        vpls "B-VPLS 1000" {
            bgp-evpn {
                accept-ivpls-evpn-flush true
```

By default, I-VPLS 1001 will not send any B-VPLS EVPN flush messages, as follows:

```
[/]
A:admin@PE-2# show service id 1001 base | match SendBvplsEvpnFlush
SendBvplsEvpnFlush: Disabled
```

The following configuration allows I-VPLS 1001 to send B-VPLS EVPN flush messages when a SAP or SDP-binding fails:

```
# on PE-2:
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
configure {
    service {
        vpls "I-VPLS 1001" {
            pbb {
                i-vpls-mac-flush {
                    bgp-evpn {
                        send-to-bvpls true
```

When enabled, the I-VPLS will send a BMAC/ISID route and subsequent updates with a higher sequence number whenever a SAP fails in the I-VPLS on the node. The default setting for a SAP allows a B-VPLS EVPN flush message to be sent (when enabled in the I-VPLS itself):

```
[/]
A:admin@PE-2# show service id 1001 sap 1/2/1:1001 detail | match SendBvplsEvpnFlush
SendBvplsEvpnFlush : Enabled
```

When no alternative route via another node is available for specific SAPs (single-homed SAPs), no CMAC flush should be triggered. When no B-VPLS EVPN flush messages need to be sent from PE-4 when SAP 1/2/1:1001 goes down, the configuration is as follows:

```
# on PE-4:
configure {
    service {
        vpls "I-VPLS 1001" {
            sap 1/2/1:1001 {
                i-vpls-mac-flush {
                    bgp-evpn {
                        send-to-bvpls false
```

The router only installs the BMACs received in MAC routes that have Ethernet tag zero. When CMAC flush is enabled, MAC routes with Ethernet tag equal to the ISID (always non-zero) are for CMAC flush, but not for installing the conveyed BMACs.

BMAC/ISID routes have the following characteristics:

- BMAC/ISID routes are sent with the static bit flag set as for any other BMAC route. The static bit is ignored at reception because this route is never used to install a BMAC in the FDB.
- BMAC/ISID routes received with non-zero ESI and non-zero Ethernet tag are treated as withdraw by the router at application level. Route Reflectors (RRs) treat such BMAC/ISID routes as valid routes that can be forwarded.
- BMAC/ISID routes are shown as valid in the **show router bgp routes evpn mac** commands, as in the following output, even though they are not used to populate the FDB. This shows that BGP is sending the routes to the application layer for CMAC flush processing. The BMAC/0 route should be sent before the BMAC/ISID routes for the same BMAC. Also, when the B-VPLS goes operationally down, the BMAC/0 should be withdrawn before the BMAC/ISID routes.

```
[/]
A:admin@PE-2# show router bgp routes evpn mac rd 192.0.2.3:1000
===============================================================================
 BGP Router ID:192.0.2.2        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
BGP EVPN MAC Routes
===============================================================================
Flag  Route Dist.          MacAddr           ESI
      Tag                  Mac Mobility      Label1
                           Ip Address
                           NextHop
-------------------------------------------------------------------------------
u*>i  192.0.2.3:1000       00:00:00:00:00:03 ESI-0
      0                    Static            LABEL 524282
                           n/a
                           192.0.2.3

u*>i  192.0.2.3:1000       00:00:00:00:00:03 ESI-0
      1001                 Static            LABEL 524282
                           n/a
                           192.0.2.3

-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

When **pbb>i-vpls-mac-flush>bgp-evpn>send-to-bvpls true** is configured in an I-VPLS that is associated with a B-VPLS, BGP-EVPN BMAC/ISID updates will be sent when certain events take place in the I-VPLS or B-VPLS. Table 12: CMAC flush transmission behavior shows the CMAC flush transmission behavior at the egress PE.

Table 12: CMAC flush transmission behavior

Local Event	pbb>i-vpls-mac-flush>bgp-evpn>send-to-bvpls	sap>i-vpls-mac-flush>bgp-evpn>send-to-bvpls	Action
Reconfigure I-VPLS: enable or disable send-to-bvpls	true or false	N/A	Send update/withdraw source BMAC/ISID with Seq=0
Associate/disassociate I-VPLS to/from B-VPLS	true	N/A	Send update/withdraw source BMAC/ISID with Seq=0
I-VPLS oper-up/oper-down	true	N/A	Send update/withdraw source BMAC/ISID with Seq=0
B-VPLS oper-up/oper-down	true	N/A	Send update/withdraw source BMAC/ISID with Seq=0 Note: All BMACs are also advertised/withdrawn.
B-VPLS bgp-evpn mpls enabled/disabled	true	N/A	Send update/withdraw source BMAC/ISID with Seq=0

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Local Event	pbb>i-vpls-mac-flush>bgp-evpn>send-to-bvpls	sap>i-vpls-mac-flush>bgp-evpn>send-to-bvpls	Action
B-VPLS operational source BMAC change	true	N/A	Send update/withdraw source BMAC/ISID with Seq=0
SAP oper-up	true	N/A	No operation
SAP oper-down	true	true	Send update source BMAC/ ISID Seq=Seq+1
	true	false	No operation

Table 13: CMAC flush reception behavior shows the reception behavior at the ingress PE. For the CMAC flush triggered by a BMAC/ISID update with increased sequence number, the B-VPLS in the receiving PE must be configured with **accept-ivpls-evpn-flush true**. BMAC/0 refers to a BMAC route where the Ethernet Tag is 0.

Table 13: CMAC flush reception behavior

Received route	Action
BMAC/0 withdraw	Flush all CMACs for that BMAC
BMAC/ISID withdraw	Flush all CMACs for that BMAC and ISID
BMAC/0 update + Seq change	Flush all CMACs for that BMAC
BMAC/ISID update + Seq change	Flush all CMACs for that BMAC and ISID
BMAC/0 update + PE (NHop) change	No CMAC-flush
BMAC/ISID update + PE (NHop) change	Flush all CMACs for that BMAC and ISID

BMAC/ISID updates will trigger CMAC flush procedures regardless of the Termination Endpoint (TEP) or Route Distinguisher (RD) with which the update is received. CMAC flush will be processed even if the BMAC-ISID comes from a TEP or RD different from the BMAC/0 route. Even when the sequence number is the same as in the previous BMAC/ISID update, CMAC flush will happen when the TEP is different. When the same BMAC/ISID is received from two PEs, both are accepted and any change in sequence number causes a MAC flush. However, when the same BMAC/ISID route is received from two PEs with the same RD, BGP will select only one, so the router only sees one.

CMAC flush for ES/vES

RFC 7623 (PBB-EVPN) defines the following CMAC Flush notification mechanisms for single-active multi-homing. These notifications do not include the local ISIDs:

- When ES-BMACs are used and the ES goes operationally down, the ES-BMAC will be withdrawn.
- When source-BMACs are used and the ES goes operationally down, a BGP-EVPN BMAC/0 is sent with a higher sequence number.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 225: ISID-independent CMAC flush when ES fails shows the following two scenarios for ISID-independent CMAC flush that are supported in SR OS Release 13.0.R4, and later:

- PBB frames are sent with the source-BMAC. When the ES goes operationally down, a BMAC update is sent with an increased sequence number, triggering a CMAC flush for all CMAC addresses associated with the BMAC address in I-VPLS, regardless of the ISID.
- PBB frames are sent with the ES-BMAC address. When the ES goes operationally down, a BMAC withdraw message is sent, triggering the remote PEs to flush all CMAC addresses associated to the ES-BMAC address, regardless of the ISID.

Figure 225: ISID-independent CMAC flush when ES fails

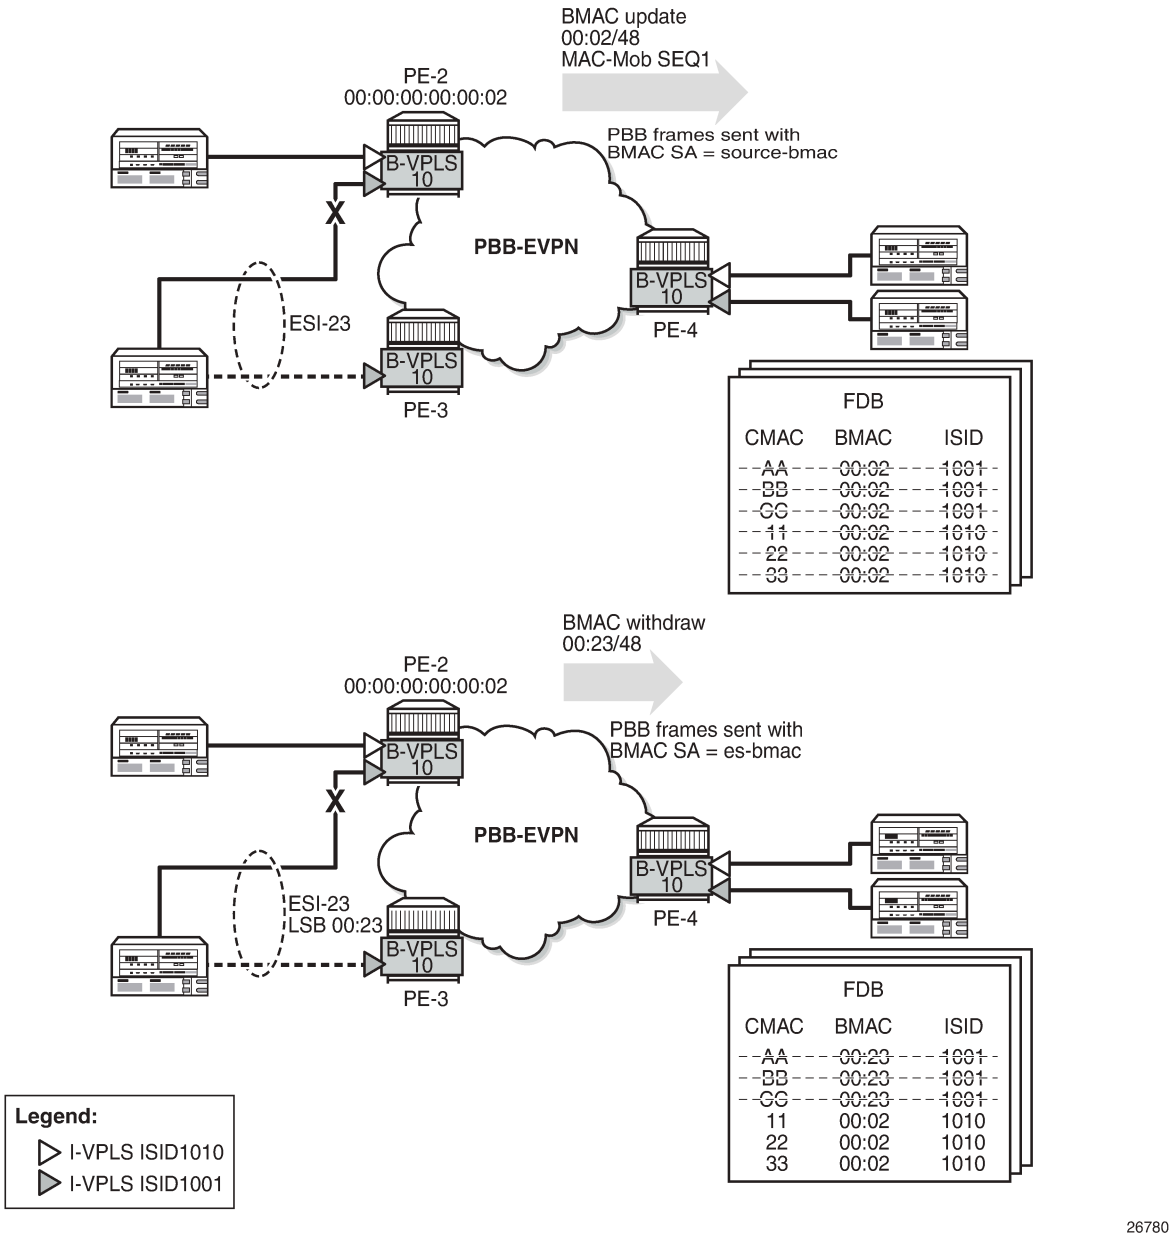

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

In addition to the preceding ISID-independent CMAC flush mechanisms, ISID-based CMAC flush is also supported in I-VPLS services with SAP or spoke-SDPs that are part of an ES or vES. ISID-based CMAC flush is enabled in the I-VPLS with the **pbb>i-vpls-mac-flush>bgp-evpn>send-to-bvpls true** command. An I-VPLS that is configured with **pbb>i-vpls-mac-flush>bgp-evpn>send-to-bvpls true** requires one of the following conditions to be met:

- The SAP or spoke-SDP has **i-vpls-mac-flush>bgp-evpn>send-to-bvpls false** configured.
- The SAP or spoke-SDP has **i-vpls-mac-flush>bgp-evpn>send-to-bvpls true** configured (default) and one of the following conditions is met:
  - The SAP or spoke-SDP is not on an ES.
  - The SAP or spoke-SDP is on an ES or vES with no **src-bmac-lsb** configured.
  - The B-VPLS has **pbb>source-bmac>use-es-bmac-lsb false** configured.

For ES SAPs with **i-vpls-mac-flush>bgp-evpn>send-to-bvpls true** in I-VPLS services that have **pbb>i-vpls-mac-flush>bgp-evpn>send-to-bvpls true** configured, the ISID-based CMAC flush replaces the RFC 7623-based CMAC flush mechanism.

For each ES/vES and B-VPLS, the system will check whether all I-VPLS services in the ES/B-VPLS have ISID-based MAC-flush enabled.

- If all I-VPLSs have **pbb>i-vpls-mac-flush>bgp-evpn>send-to-bvpls true** configured:
  - No BMAC/0 updates with increased sequence number will be triggered when the ES/vES goes operationally down.
  - Only BMAC/ISID updates with increased sequence number will be sent when the I-VPLS attachment circuit goes operationally down.
- If at least one I-VPLS has **pbb>i-vpls-mac-flush>bgp-evpn>send-to-bvpls false** configured:
  - BMAC/0 updates with increased sequence number will be triggered when the ES/vES goes operationally down.
  - Also, BMAC/ISID updates with increased sequence number will be generated for those I-VPLS services that have **pbb>i-vpls-mac-flush>bgp-evpn>send-to-bvpls true** configured.

The number of CMAC addresses that may be flushed at the remote nodes can be reduced by enabling ISID-based MAC-flush for all the I-VPLS services in the ES/vES.

When attempting to set **use-es-bmac-lsb true** in B-VPLS 1000 on PE-4 when the SAP/SDP-binding has default settings (and **pbb>i-vpls-mac-flush>bgp-evpn>send-to-bvpls true** in the I-VPLS), the following error is raised:

```
[ex:/configure service vpls "B-VPLS 1000" pbb source-bmac]
A:admin@PE-4# use-es-bmac-lsb true

*[ex:/configure service vpls "B-VPLS 1000" pbb source-bmac]
A:admin@PE-4# commit
MINOR: MGMT_CORE #4001: configure service vpls "I-VPLS 1024" spoke-sdp 46:1024 - ethernet-
segment ESI-45 using es-bmac and service has send-bvpls-evpn-flush enabled - configure service
 vpls "I-VPLS 1024" pbb i-vpls-mac-flush bgp-evpn send-to-bvpls
MINOR: MGMT_CORE #4001: configure service vpls "I-VPLS 1001" spoke-sdp 46:1001 - ethernet-
segment ESI-45 using es-bmac and service has send-bvpls-evpn-flush enabled - configure service
 vpls "I-VPLS 1001" pbb i-vpls-mac-flush bgp-evpn send-to-bvpls
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

However, when the ES is disabled, the B-VPLS can be configured with **use-es-bmac-lsb true**. When attempting to re-enable the ES afterward, the following error is raised.

```
[ex:/configure service system bgp evpn ethernet-segment "ESI-45"]
A:admin@PE-4# admin-state enable

*[ex:/configure service system bgp evpn ethernet-segment "ESI-45"]
A:admin@PE-4# commit
MINOR: MGMT_CORE #4001: configure service vpls "I-VPLS 1024" spoke-sdp 46:1024 - ethernet-
segment ESI-45 using es-bmac and service has send-bvpls-evpn-flush enabled - configure service
 vpls "I-VPLS 1024" pbb i-vpls-mac-flush bgp-evpn send-to-bvpls
MINOR: MGMT_CORE #4001: configure service vpls "I-VPLS 1001" spoke-sdp 46:1001 - ethernet-
segment ESI-45 using es-bmac and service has send-bvpls-evpn-flush enabled - configure service
 vpls "I-VPLS 1001" pbb i-vpls-mac-flush bgp-evpn send-to-bvpls
```

Configuration

Figure 226: Example topology shows the example topology.

Figure 226: Example topology

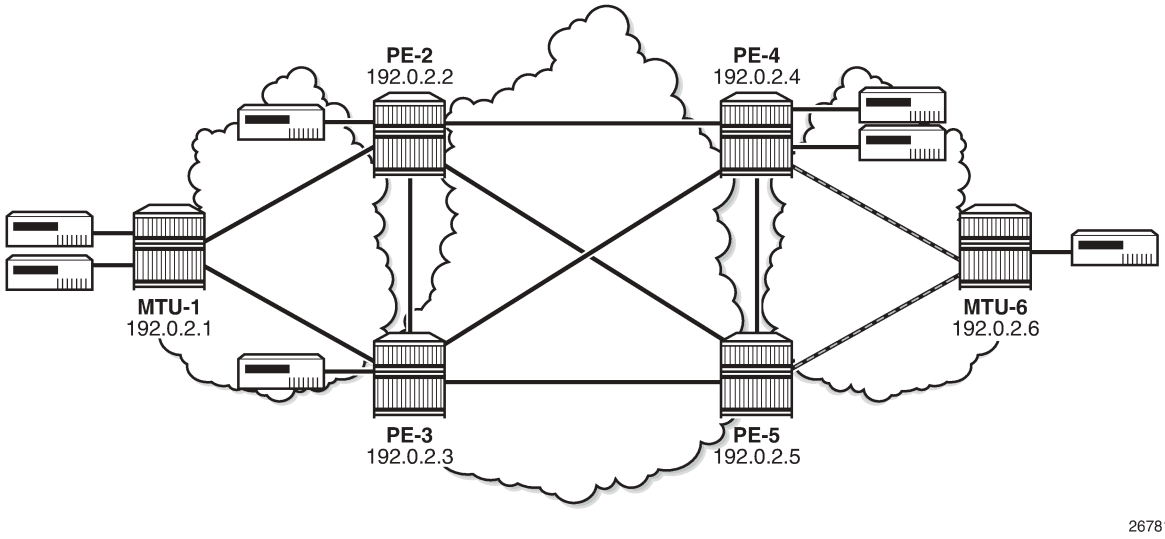

The initial configuration includes the following:

- Cards, MDAs
- Ports: the ports between the MTUs and the PEs are hybrid or access ports with dot1q encapsulation; the ports between the PEs are network ports with null encapsulation
- Router interfaces
- IS-IS on all router interfaces (alternatively, OSPF could be used)
- LDP on all router interfaces

The following use cases are described in this section:

- ISID-based CMAC flush for BGP non-EVPN multi-homing (no ES)
- ISID-based CMAC flush for BGP-EVPN in a single-active ES

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

ISID-based CMAC flush for BGP multi-homing

Figure 227: Example topology with BGP multi-homing shows the example topology with BGP multi-homing site 1 between PE-2 and PE-3. B-VPLS 1000 is configured on all the core nodes (PEs) and I-VPLS 1001 and I-VPLS 1010 are associated with this B-VPLS in the PEs. On MTU-1, regular VPLSs are configured. For more information about BGP non-EVPN multi-homing, see chapter BGP Multi-Homing for VPLS Networks.

Figure 227: Example topology with BGP multi-homing

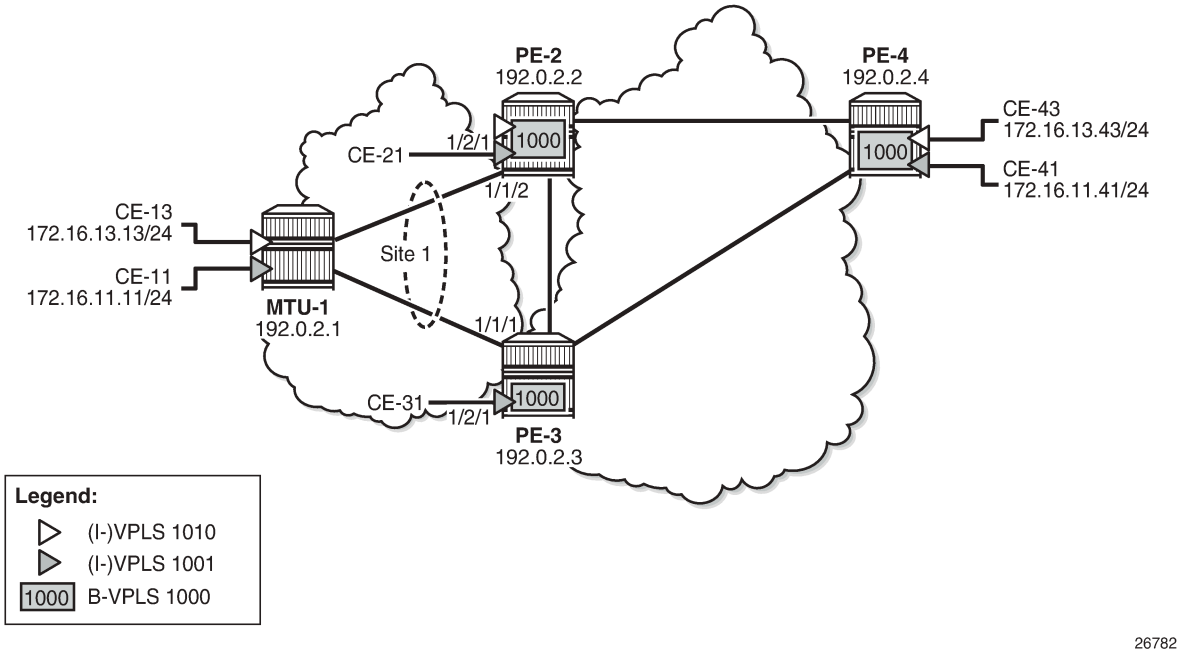

BGP is configured for address family EVPN on all PEs with PE-2 as RR. For BGP multi-homing, address family L2-VPN is enabled between PE-2 and PE-3. The BGP configuration on PE-2 is as follows:

```
# on PE-2:
configure {
    router "Base"
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            split-horizon true
            rapid-update {
                l2-vpn true
                evpn true
            }
            group "internal" {
                peer-as 64500
                family {
                    evpn true
                }
                cluster {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    cluster-id 1.1.1.1
                }
            }
            neighbor "192.0.2.3" {
                group "internal"
                family {
                    l2-vpn true
                    evpn true
                }
            }
            neighbor "192.0.2.4" {
                group "internal"
                family {
                    evpn true
                }
            }
        }
```

The BGP configuration on PE-4 is as follows:

```
# on PE-4:
configure {
    router "Base"
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            split-horizon true
            rapid-update {
                evpn true
            }
            group "internal" {
                peer-as 64500
                family {
                    evpn true
                }
            }
            neighbor "192.0.2.2" {
                group "internal"
            }
        }
```

The configuration of B-VPLS 1000 and I-VPLS 1001 on PE-2 is as follows. ISID-based CMAC flush is disabled by default. BGP multi-homing site "MH-site-1" is configured on PE-2 with SAP 1/1/2:1001 associated with it, whereas SAP 1/2/1:1001 is not associated to the MH site. CE-21 is attached to I-VPLS 1001 with SAP 1/2/1:1001.

```
# on PE-2:
configure {
    service {
        system {
            bgp-auto-rd-range {
                ip-address 192.0.2.2
                community-value {
                    start 1
                    end 999
                }
            }
        }
        vpls "B-VPLS 1000" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            admin-state enable
            service-id 1000
            customer "1"
            service-mtu 2000
            pbb-type b-vpls
            pbb {
                source-bmac {
                    address 00:00:00:00:00:02
                }
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 1000
                mpls 1 {
                    admin-state enable
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
        }
        vpls "I-VPLS 1001" {
            admin-state enable
            service-id 1001
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS 1000" {
                    isid 1001
                }
            }
            bgp 1 {
                route-distinguisher auto-rd
                route-target {
                    export "target:64500:1001"
                    import "target:64500:1001"
                }
            }
            sap 1/1/2:1001 {
            }
            sap 1/2/1:1001 {
            }
            bgp-mh-site "MH-site-1" {
                admin-state enable
                id 1
                sap 1/1/2:1001
            }
        }
        vpls "I-VPLS 1010" {
            admin-state enable
            service-id 1010
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS 1000" {
                    isid 1010
                }
            }
            bgp 1 {
                route-distinguisher auto-rd
                route-target {
                    export "target:64500:1010"
                    import "target:64500:1010"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
            }
            sap 1/1/2:1010 {
            }
        }
```

I-VPLS 1010 is configured without multi-homing. The configuration of VPLS 1001 on PE-3 is similar, but without I-VPLS 1010.

ISID-based CMAC flush is not enabled yet. The PEs exchange BGP-EVPN MAC routes with Ethernet tag zero. PE-3 has received BMAC/0 routes from PE-2 and PE-4, as follows:

```
[/]
A:admin@PE-3# show router bgp routes evpn mac
===============================================================================
 BGP Router ID:192.0.2.3        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN MAC Routes
===============================================================================
Flag  Route Dist.         MacAddr           ESI
      Tag                 Mac Mobility      Label1
                          Ip Address
                          NextHop
-------------------------------------------------------------------------------
u*>i  192.0.2.2:1000      00:00:00:00:00:02 ESI-0
      0                   Static            LABEL 524282
                          n/a
                          192.0.2.2

u*>i  192.0.2.4:1000      00:00:00:00:00:04 ESI-0
      0                   Static            LABEL 524282
                          n/a
                          192.0.2.4

-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

PE-2 and PE-4 have also received BMAC/0 routes from the other PEs.

ISID-based CMAC flush is enabled in I-VPLS 1001 on PE-2 and PE-3. PE-4 has no multi-homing in I-VPLS 1001, so it should not send any CMAC flush. I-VPLS 1010 has no multi-homing in any PE, so ISID-based MAC-flush should not be enabled in I-VPLS 1010.

```
# on PE-2, PE-3:
configure {
    service {
        vpls "I-VPLS 1001" {
            pbb {
                i-vpls-mac-flush {
                    bgp-evpn {
                        send-to-bvpls true
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

PE-2 and PE-3 will send BMAC/1001 updates with sequence number 0 to the other two PEs. As an example, the following EVPN-MAC route for BMAC 00:00:00:00:00:03 with tag 1001 is sent by PE-3:

```
22 2021/04/15 08:07:57.818 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 89
    Flag: 0x90 Type: 14 Len: 44 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.3
        Type: EVPN-MAC Len: 33 RD: 192.0.2.3:1000 ESI: ESI-0, tag: 1001, mac len: 48
                       mac: 00:00:00:00:00:03, IP len: 0, IP: NULL, label1: 8388512
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:1000
        bgp-tunnel-encap:MPLS
        mac-mobility:Seq:0/Static
"
```

PE-4 has received the following BMAC routes from PE-2 and PE-3, with Ethernet tag zero and Ethernet tag 1001. BMAC routes are always static (received with the sticky bit set).

```
[/]
A:admin@PE-4# show router bgp routes evpn mac
===============================================================================
 BGP Router ID:192.0.2.4        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN MAC Routes
===============================================================================
Flag  Route Dist.         MacAddr           ESI
      Tag                 Mac Mobility      Label1
                          Ip Address
                          NextHop
-------------------------------------------------------------------------------
u*>i  192.0.2.2:1000      00:00:00:00:00:02 ESI-0
      0                   Static            LABEL 524282
                          n/a
                          192.0.2.2

u*>i  192.0.2.2:1000      00:00:00:00:00:02 ESI-0
      1001                Static            LABEL 524282
                          n/a
                          192.0.2.2

u*>i  192.0.2.3:1000      00:00:00:00:00:03 ESI-0
      0                   Static            LABEL 524282
                          n/a
                          192.0.2.3

u*>i  192.0.2.3:1000      00:00:00:00:00:03 ESI-0
      1001                Static            LABEL 524282
                          n/a
                          192.0.2.3
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
Routes : 4
===============================================================================
```

When a failure occurs on PE-2, PE-3 and PE-4 should accept the BMAC/ISID with increased sequence number; for a failure on PE-3, PE-2 and PE-4 should accept the BMAC/ISID update. Therefore, the B-VPLS on all PEs should accept the CMAC flush message for ISID 1001, and this is configured as follows:

```
# on PE-2, PE-3, PE-4:
configure {
    service {
        vpls "B-VPLS 1000" {
            bgp-evpn {
                accept-ivpls-evpn-flush true
```

The FDB for VPLS 1001 on PE-4 includes MAC address 00:00:11:11:11:11 with source-identifier 192.0.2.2:524282, so PE-4 will forward traffic toward that MAC address to PE-2.

```
[/]
A:admin@PE-4# show service id 1001 fdb detail

===============================================================================
Forwarding Database, Service 1001
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
1001       00:00:11:11:11:11 b-mpls:                 L/420    04/15/21 08:03:47
                             192.0.2.2:524282
           ldp:65537
1001       00:00:41:41:41:41 sap:1/2/1:1001          L/0      04/15/21 08:11:36
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

A failure is simulated on SAP 1/1/2:1001 in multi-homing site 1 on PE-2 as follows:

```
# on PE-2:
configure {
    service {
        vpls "I-VPLS 1001" {
            sap 1/1/2:1001 {
                admin-state disable
```

SAP 1/1/2:1001 has the default **i-vpls-mac-flush>bgp-evpn>send-to-bvpls true** and I-VPLS 1001 is configured with **pbb>i-vpls-mac-flush>bgp-evpn>send-to-bvpls true**, so PE-2 will send BMAC/ISID updates for BMAC 00:00:00:00:00:02, ISID 1001, and sequence number 1 to its BGP peers. The following BGP update is sent by PE-2 to PE-4:

```
# on PE-2:
64 2021/04/15 08:12:55.058 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.4
"Peer 1: 192.0.2.4: UPDATE
Peer 1: 192.0.2.4 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 89
    Flag: 0x90 Type: 14 Len: 44 Multiprotocol Reachable NLRI:
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
        Type: EVPN-MAC Len: 33 RD: 192.0.2.2:1000 ESI: ESI-0, tag: 1001, mac len: 48
                       mac: 00:00:00:00:00:02, IP len: 0, IP: NULL, label1: 8388512
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:1000
        bgp-tunnel-encap:MPLS
        mac-mobility:Seq:1/Static
"
```

This BMAC/ISID with sequence number 1 triggers a CMAC flush in the FDB for VPLS 1001, so the entry for 00:00:11:11:11:11 will be flushed, along with all other MAC addresses associated with BMAC 00:00:00:00:00:02. The FDB on PE-4 does not contain any entries with source-identifier BMAC 00:00:00:00:00:02, as follows:

```
[/]
A:admin@PE-4# show service id 1001 fdb detail

===============================================================================
Forwarding Database, Service 1001
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
1001       00:00:41:41:41:41 sap:1/2/1:1001          L/150    04/15/21 08:11:36
-------------------------------------------------------------------------------
No. of MAC Entries: 1
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

When the MAC address 00:00:11:11:11:11 is learned via PE-3, the FDB is as follows:

```
[/]
A:admin@PE-4# show service id 1001 fdb detail

===============================================================================
Forwarding Database, Service 1001
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
1001       00:00:11:11:11:11 b-mpls:                 L/0      04/15/21 08:15:16
                             192.0.2.3:524282
           ldp:65538
1001       00:00:41:41:41:41 sap:1/2/1:1001          L/0      04/15/21 08:11:36
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

The CMAC flush is only applied for VPLS 1001, so the FDB for VPLS 1010 on PE-4 will keep entries learned from PE-2, as follows:

```
[/]
A:admin@PE-4# show service id 1010 fdb detail
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Forwarding Database, Service 1010
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
1010       00:00:13:13:13:13 b-mpls:                 L/0      04/15/21 08:03:48
                             192.0.2.2:524282
           ldp:65537
1010       00:00:43:43:43:43 sap:1/2/1:1010          L/0      04/15/21 08:11:36
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

ISID-based CMAC flush in single-active ES

CMAC flush only makes sense for single-active multi-homing. Also, CMAC flush only works for single-active multi-homing; not for all-active multi-homing, because ES-BMAC is required in all-active multi-homing. Figure 228: Example topology with single-active ES shows the example topology with a single-active ES "ESI-45" configured in PE-4 and PE-5.

Figure 228: Example topology with single-active ES

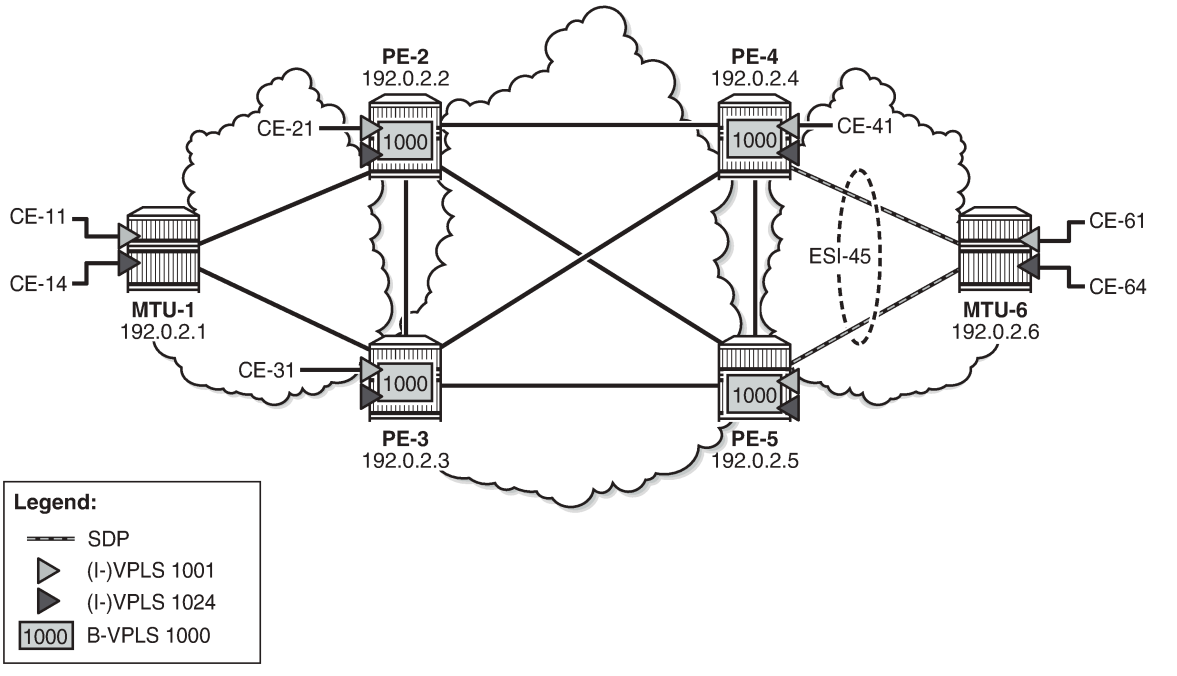

The multi-homing configuration has been removed from PE-2 and PE-3, so no CMAC flush should be sent by PE-2 or PE-3. VPLS 1001 is configured as follows on PE-2 and PE-3:

```
# on PE-2, PE-3:
configure {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    service {
        vpls "I-VPLS 1001" {
            admin-state enable
            service-id 1001
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS 1000" {
                    isid 1001
                }
            }
            bgp 1 {
                route-distinguisher auto-rd
                route-target {
                    export "target:64500:1001"
                    import "target:64500:1001"
                }
            }
            sap 1/2/1:1001 {
            }
            sap lag-1:1001 {
            }
        }
```

SDPs are configured between PE-4 and MTU-6, and between PE-5 and MTU-6. These SDPs are associated with the single-active ES "ESI-45".

The configuration of B-VPLS 1000 on PE-4 is as follows. The B-VPLS configuration on the other PEs is similar, but with a different source BMAC.

```
# on PE-4:
configure {
    service {
        vpls "B-VPLS 1000" {
            admin-state enable
            service-id 1000
            customer "1"
            service-mtu 2000
            pbb-type b-vpls
            pbb {
                source-bmac {
                    address 00:00:00:00:00:04
                }
            }
            bgp 1 {
            }
            bgp-evpn {
                accept-ivpls-evpn-flush true
                evi 1000
                mpls 1 {
                    admin-state enable
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
        }
```

The service configuration on PE-4 includes an SDP toward PE-6 and a single-active multi-homing ES, as follows:

```
# on PE-4:
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "ESI-45" {
                        admin-state enable
                        esi 01:00:00:00:00:45:00:00:00:01
                        multi-homing-mode single-active
                        df-election {
                            es-activation-timer 3
                        }
                        association {
                            sdp 46 {
                            }
                        }
                        pbb {
                            source-bmac-lsb 45-04
                        }
                    }
                }
            }
        }
        sdp 46 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.6
            }
        }
```

The configuration on PE-5 is similar. The configuration of B-VPLS 1000 is similar to the one for PE-2, with only a different BMAC. The configuration of I-VPLS 1001 on PE-4 is as follows:

```
# on PE-4:
configure {
    service {
        vpls "I-VPLS 1001" {
            admin-state enable
            service-id 1001
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS 1000" {
                    isid 1001
                }
                i-vpls-mac-flush {
                    bgp-evpn {
                        send-to-bvpls true
                    }
                }
            }
            bgp 1 {
                route-distinguisher auto-rd
                route-target {
                    export "target:64500:1001"
                    import "target:64500:1001"
                }
            }
            spoke-sdp 46:1001 {
            }
            sap 1/2/1:1001 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
        }
```

ISID-based MAC-flush is enabled in B-VPLS 1000 and I-VPLS 1001 on all PEs.

I-VPLS 1024 is also associated with B-VPLS 1000 and contains one object (SAP or spoke-SDP) in each PE. The configuration of I-VPLS 1024 is identical on PE-2 and PE-3, as follows:

```
# on PE-2, PE-3:
configure {
    service {
        vpls "I-VPLS 1024" {
            admin-state enable
            service-id 1024
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS 1000" {
                    isid 1024
                }
            }
            sap lag-1:1024 {
            }
        }
```

The configuration of I-VPLS 1024 on PE-4 has **pbb>i-vpls-mac-flush>bgp-evpn>send-to-bvpls true** configured and contains a spoke-SDP instead of a SAP, as follows. The configuration on PE-5 is similar, but with a different SDP.

```
# on PE-4:
configure {
    service {
        vpls "I-VPLS 1024" {
            admin-state enable
            service-id 1024
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS 1000" {
                    isid 1024
                }
                i-vpls-mac-flush {
                    bgp-evpn {
                        send-to-bvpls true
                    }
                }
            }
            spoke-sdp 46:1024 {
            }
        }
```

ISID-based MAC-flush is enabled on PE-4 and PE-5 for both I-VPLS 1001 and I-VPLS 1024, and BMAC/ISID updates are sent for ISID 1001 and ISID 1024, as follows:

```
[/]
A:admin@PE-3# show router bgp routes evpn mac rd 192.0.2.4:1000
===============================================================================
 BGP Router ID:192.0.2.3        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN MAC Routes
===============================================================================
Flag  Route Dist.          MacAddr           ESI
      Tag                  Mac Mobility      Label1
                           Ip Address
                           NextHop
-------------------------------------------------------------------------------
u*>i  192.0.2.4:1000       00:00:00:00:00:04 ESI-0
      0                    Static            LABEL 524282
                           n/a
                           192.0.2.4

u*>i  192.0.2.4:1000       00:00:00:00:00:04 ESI-0
      1001                 Static            LABEL 524282
                           n/a
                           192.0.2.4

u*>i  192.0.2.4:1000       00:00:00:00:00:04 ESI-0
      1024                 Static            LABEL 524282
                           n/a
                           192.0.2.4

-------------------------------------------------------------------------------
Routes : 3
===============================================================================
```

PE-5 is the DF for VPLS 1001 in the single-active ES "ESI-45", but not for VPLS 1024, as follows:

```
[/]
A:admin@PE-5# show service id 1001 ethernet-segment
No sap entries

===============================================================================
SDP Ethernet-Segment Information
===============================================================================
SDP                   Eth-Seg                         Status
-------------------------------------------------------------------------------
56:1001               ESI-45                          DF
===============================================================================
No vxlan instance entries
```

```
[/]
A:admin@PE-5# show service id 1024 ethernet-segment
No sap entries

===============================================================================
SDP Ethernet-Segment Information
===============================================================================
SDP                   Eth-Seg                         Status
-------------------------------------------------------------------------------
56:1024               ESI-45                          NDF
===============================================================================
No vxlan instance entries
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The following FDB for VPLS 1001 on PE-5 shows that traffic toward CMAC 00:00:11:11:11:11 (CE-11) in VPLS 1001 will be forwarded to PE-3:

```
[/]
A:admin@PE-5# show service id 1001 fdb detail

===============================================================================
Forwarding Database, Service 1001
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
1001       00:00:11:11:11:11 b-mpls:                 L/0      04/15/21 08:19:47
                             192.0.2.3:524282
           ldp:65539
1001       00:00:41:41:41:41 b-mpls:                 L/0      04/15/21 08:19:47
                             192.0.2.4:524282
           ldp:65537
1001       00:00:61:61:61:61 sdp:56:1001             L/0      04/15/21 08:19:42
-------------------------------------------------------------------------------
No. of MAC Entries: 3
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

The following FDB for VPLS 1024 on PE-4 shows that traffic toward CMAC 00:00:14:14:14:14 (CE-14) will be forwarded to PE-2:

```
[/]
A:admin@PE-4# show service id 1024 fdb detail

===============================================================================
Forwarding Database, Service 1024
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
1024       00:00:14:14:14:14 b-mpls:                 L/0      04/15/21 08:19:48
                             192.0.2.2:524282
           ldp:65537
1024       00:00:64:64:64:64 sdp:46:1024             L/0      04/15/21 08:19:48
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

The following FDB for VPLS 1001 on PE-3 shows that traffic toward CMAC 00:00:61:61:61:61 (CE-61) will be forwarded to PE-5:

```
[/]
A:admin@PE-3# show service id 1001 fdb detail

===============================================================================
Forwarding Database, Service 1001
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
1001       00:00:11:11:11:11 sap:lag-1:1001          L/0      04/15/21 08:19:47
1001       00:00:41:41:41:41 b-mpls:                 L/0      04/15/21 08:19:47
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                            192.0.2.4:524282
          ldp:65538
1001      00:00:61:61:61:61 b-mpls:                L/0      04/15/21 08:19:42
                            192.0.2.5:524282
          ldp:65539
-------------------------------------------------------------------------------
No. of MAC Entries: 3
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

The following FDB for VPLS 1024 on PE-2 shows that traffic toward CMAC 00:00:64:64:64:64 (CE-64) will be forwarded to PE-4:

```
[/]
A:admin@PE-2# show service id 1024 fdb detail

===============================================================================
Forwarding Database, Service 1024
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
1024       00:00:14:14:14:14 sap:lag-1:1024          L/0      04/15/21 08:19:48
1024       00:00:64:64:64:64 b-mpls:                 L/0      04/15/21 08:19:48
                            192.0.2.4:524282
          ldp:65538
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

PE-5 is the DF for VPLS 1001 in "ESI-45". A failure is simulated by disabling the SDP toward PE-5 on MTU-6, as follows:

```
# on MTU-6:
configure {
    service {
        sdp 65 {
            admin-state disable
```

PE-5 sends the following BMAC/ISID with increased sequence number for ISID 1001 to the RR PE-2:

```
50 2021/04/15 08:24:35.567 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 89
    Flag: 0x90 Type: 14 Len: 44 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.5
        Type: EVPN-MAC Len: 33 RD: 192.0.2.5:1000 ESI: ESI-0, tag: 1001, mac len: 48
                       mac: 00:00:00:00:00:05, IP len: 0, IP: NULL, label1: 8388496
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:1000
        bgp-tunnel-encap:MPLS
        mac-mobility:Seq:1/Static
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
"
```

When PE-3 receives this BMAC/ISID, all MAC routes with next-hop PE-5 are flushed and the FDB will contain the following MAC entries:

```
[/]
A:admin@PE-3# show service id 1001 fdb detail

===============================================================================
Forwarding Database, Service 1001
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
1001       00:00:11:11:11:11 sap:lag-1:1001          L/0      04/15/21 08:19:47
1001       00:00:41:41:41:41 b-mpls:                 L/0      04/15/21 08:19:47
                             192.0.2.4:524282
           ldp:65538
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

If MAC address 00:00:61:61:61:61 is learned again, the next hop will be PE-4 instead of PE-5.

The configuration is restored as follows:

```
# on MTU-6:
configure {
    service {
        sdp 65 {
            admin-state enable
```

No CMAC/ISID update will be sent when the last SAP/SDP-binding in a service goes operationally down. VPLS 1024 only has one SAP/SDP-binding in DF PE-4: spoke-SDP 46:1024. A failure of the spoke-SDP is simulated as follows:

```
# on MTU-6:
configure {
    service {
        sdp 64 {
            admin-state disable
```

When the last SAP/SDP-binding is down, the service will be operationally down, as follows:

```
[/]
A:admin@PE-4# show service id 1024 base | match "Oper State"
Admin State       : Up                 Oper State        : Down
```

PE-4 sends the following withdrawal message instead of a CMAC/ISID:

```
56 2021/04/15 08:26:10.691 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 61
    Flag: 0x90 Type: 15 Len: 57 Multiprotocol Unreachable NLRI:
        Address Family EVPN
        Type: EVPN-INCL-MCAST Len: 17 RD: 192.0.2.4:1000, tag: 1024,
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                              orig_addr len: 32, orig_addr: 192.0.2.4
        Type: EVPN-MAC Len: 33 RD: 192.0.2.4:1000 ESI: ESI-0, tag: 1024, mac len: 48
                       mac: 00:00:00:00:00:04, IP len: 0, IP: NULL, label1: 0
"
```

The configuration is restored as follows:

```
# on MTU-6:
configure {
    service {
        sdp 64 {
            admin-state enable
```

ISID-based and regular CMAC flush in ES

When ISID-based CMAC flush is not enabled in all I-VPLS services using the ES, a failure in the ES will trigger BMAC/0 updates and BMAC/ISID updates with increased sequence number. An additional I-VPLS is configured on the nodes with **pbb>i-vpls-mac-flush>bgp-evpn>send-to-bvpls false** (default). The configuration of I-VPLS 1021 on PE-5 is as follows:

```
# on PE-5:
configure {
    service {
        vpls "I-VPLS 1021" {
            admin-state enable
            service-id 1021
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS 1000" {
                    isid 1021
                }
            }
            spoke-sdp 56:1021 {
            }
            sap 1/2/1:1021 {
            }
        }
```

The configuration on PE-4 is similar; PE-2 and PE-3 have SAP lag-1:1021 instead of the spoke-SDP.

On MTU-6, SDP 65 is disabled, which will cause an ES failure on PE-5:

```
# on MTU-6:
configure {
    service {
        sdp 65 {
            admin-state disable
```

The following BMAC updates are sent by PE-5:

- BMAC/0 with increased sequence number, which will trigger a CMAC flush for all entries received from PE-5 for all I-VPLS services (ISID-independent)
- BMAC/ISID with increased sequence number, which will trigger a CMAC flush for all entries received from PE-5 for VPLS 1001

```
73 2021/04/15 08:32:57.204 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 89
    Flag: 0x90 Type: 14 Len: 44 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.5
        Type: EVPN-MAC Len: 33 RD: 192.0.2.5:1000 ESI: ESI-0, tag: 0, mac len: 48
                       mac: 00:00:00:00:00:05, IP len: 0, IP: NULL, label1: 8388496
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:1000
        bgp-tunnel-encap:MPLS
        mac-mobility:Seq:1/Static
"

74 2021/04/15 08:32:57.204 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 89
    Flag: 0x90 Type: 14 Len: 44 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.5
        Type: EVPN-MAC Len: 33 RD: 192.0.2.5:1000 ESI: ESI-0, tag: 1001, mac len: 48
                       mac: 00:00:00:00:00:05, IP len: 0, IP: NULL, label1: 8388496
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:1000
        bgp-tunnel-encap:MPLS
        mac-mobility:Seq:3/Static
"
```

Conclusion

ISID-based MAC-flush speeds up convergence after a SAP or spoke-SDP failure, triggering a selective CMAC flush on the receiving nodes, which flushes all CMAC entries associated with that ISID and BMAC. The feature can be enabled per I-VPLS and disabled for those SAPs or spoke-SDPs for which no alternative route is available, or for those SAPs that are contained in an all-active Ethernet Segment. The BMAC/ISID update always contains the source-BMAC, not the ES-BMAC. CMAC flush based on ES-BMAC is not performed per ISID.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

PBB-EVPN ISID-based Route Targets

This chapter provides information about PBB-EVPN ISID-based Route Targets.

Topics in this chapter include:

Applicability

This chapter was initially written based on SR OS Release 15.0.R4, but the MD-CLI in the current edition corresponds to SR OS Release 21.5.R1. PBB-EVPN ISID-based route targets are supported in SR OS Release 15.0.R1, and later.

Overview

The following BGP-EVPN routes are used in PBB-EVPN according to RFC 7623:

- B-MAC routes—based on BGP-EVPN route type 2—are sent with the B-VPLS Route Target (RT), so they are sent to all the PEs where the B-VPLS is defined.
- Ethernet Segment (ES) routes—route type 4—are used for multi-homing. ES routes are sent with an RT auto-derived from the ES Identifier (ESI). If the RT-constraint is enabled, the routes are sent to only those PEs that are part of the ES.
- Inclusive Multicast Ethernet Tag (IMET) routes—route type 3—are used for the setup of per-ISID flooding domains and can be sent with a B-VPLS RT or with an ISID-based RT.
  - IMET routes are, by default, sent with a B-VPLS RT (referred to as IMET/0 routes), so they are imported by all the PEs where the B-VPLS is defined, as per RFC 7623, and supported in SR OS Release 13.0.R4, and later.
  - IMET routes with an ISID-based RT (referred to as IMET/ISID routes) are imported by only the PEs where the ISID is defined. RFC 7623 recommends these routes for deployments where the ISIDs are sparsely distributed in the network. This is supported in SR OS Release 15.0.R1, and later. The service ISID is encoded in the Ethernet tag field.

Figure 229: PBB-EVPN B-VPLS-based RT shows how the B-MAC and IMET routes with a B-VPLS RT sent by PE-1 are advertised to all other PEs (via the Route Reflector (RR)), regardless of the ISID.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 229: PBB-EVPN B-VPLS-based RT

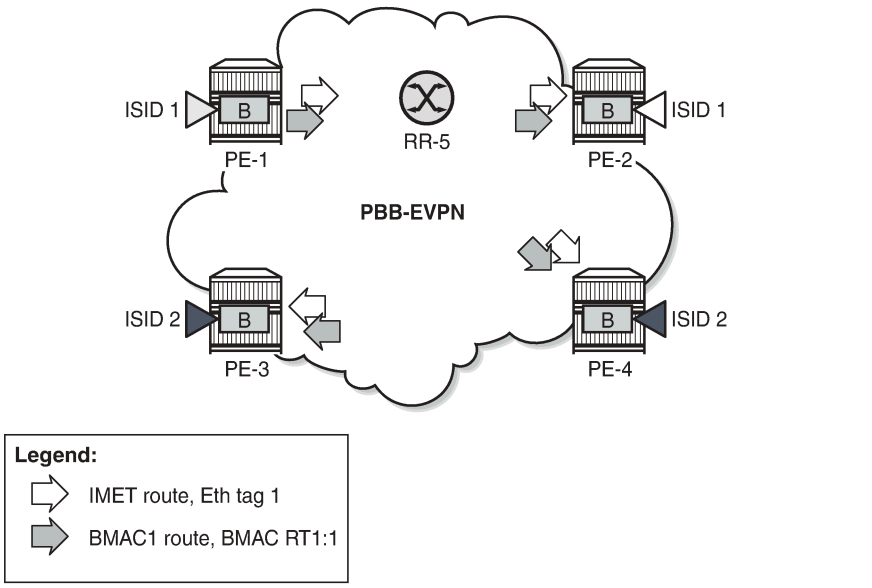

Figure 230: PBB-EVPN ISID-based RT shows how the B-MAC routes are sent to all PEs within the B-VPLS, whereas the IMET routes sent by PE-1 are selectively reflected by the RR (due to RT-constraints) and only sent to PE-2, which is the only PE with the same ISID.

Figure 230: PBB-EVPN ISID-based RT

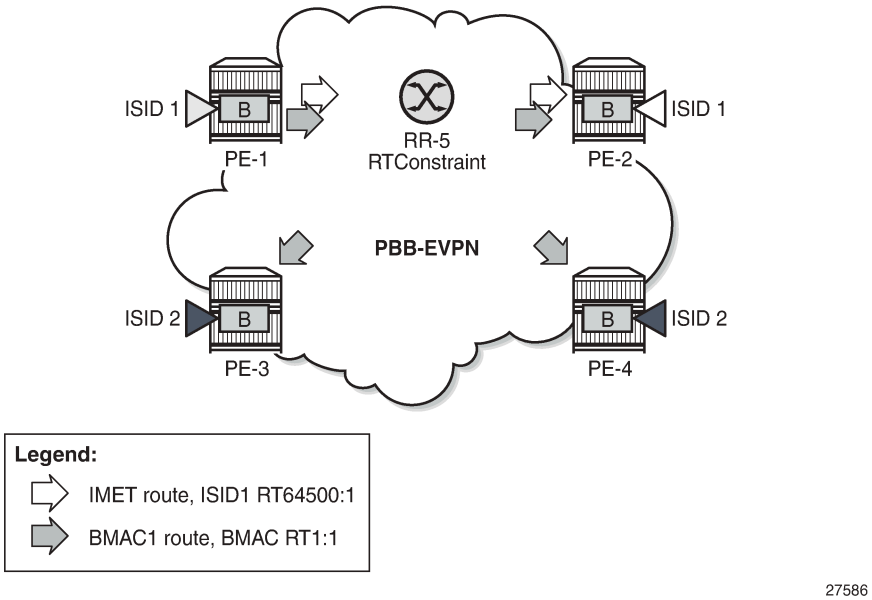

IMET routes with ISID-based RTs (IMET/ISID) can significantly reduce the number of IMET/ISID routes distributed by the RRs. The RT for the IMET/ISID route can be auto-derived from the corresponding Ethernet tag (ISID).

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

In addition to RFC 7623, the ISID-derived RTs can be used for BMAC/ISID routes if ISID-based CMAC flush is enabled, as per *draft-snr-bess-pbb-evpn-isid-cmacflush*. The service ISID is encoded in the Ethernet tag field.

PBB-EVPN ISID-based RT format

Figure 231: PBB-EVPN ISID-based RT format shows the ISID-based RT format:

Figure 231: PBB-EVPN ISID-based RT format

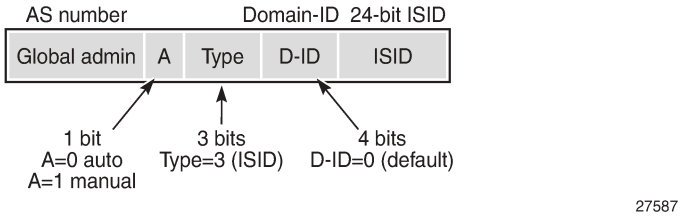

For an auto-derived ISID-based RT, the values are as follows:

- The Autonomous System (AS) number is obtained from the **config>router>autonomous-system** command:
  - Value = 2-byte AS number
  - For AS numbers with more than 2 bytes, the low-order 16-bit value is used.
- A = 0 for auto-derivation
- Type = 011 = 3 for ISID-based RT
- Domain ID = 0000 (default)
- ISID value

The auto-derived RT will be AS:00110000+ISID = AS:0x30+ISID Hex.

The type and sub-type of the BGP extended community is 0x00 and 0x02.

Enabling ISID-based RT

The following command is used to enable ISID-based RT for specific ISID ranges for IMET/ISID and BMAC/ISID routes.

```
*[ex:/configure service vpls "B-VPLS 100" bgp-evpn]
A:admin@PE-1# isid-route-target ?

 isid-route-target

 range                 + Enter the range list instance
```

The ISID range is configured as follows:

```
[ex:/configure service vpls "B-VPLS 100" bgp-evpn isid-route-target]
A:admin@PE-1# range ?

 [start] <number>
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
 <number>  - <1..16777215>

    Starting value of the isid-range entry
```

```
[ex:/configure service vpls "B-VPLS 100" bgp-evpn isid-route-target range 1]
A:admin@PE-1# end ?

 end <number>
 <number>  - <1..16777215>

 'end' is: mandatory

    Ending value of the isid-range entry
```

The RT to be used for the I-VPLS can be auto-derived (default) or explicitly configured. The following configures an ISID range from 20 to 29 with auto-derived RT (default: type auto), whereas ISID 30 has a manually configured RT of 64500:30.

```
# on PE-1:
configure {
    service {
        vpls "B-VPLS 100" {
            bgp-evpn {
                isid-route-target {
                    range 20 {
                        end 29
                        # type auto   # default
                    }
                    range 30 {
                        end 30
                        type configured
                        route-target "target:64500:30"
                    }
                }
```

If **isid-route-target** is enabled, the IMET/ISID and BMAC/ISID route processing is modified in the export and import directions:

- "Exported IMET/ISID and BMAC/ISID routes:
  - IMET/ISID routes are sent with an ISID-based RT for the local I-VPLS ISIDs and static ISIDs, unless the ISID is contained in an ISID policy for which **advertise-local false** is configured.
  - When **isid-route-target** and **ivpls-mac-flush>bgp-evpn>send-to-bvpls** are both enabled for an I-VPLS, the BMAC/ISID route will also be sent with the ISID-based RT instead of the B-VPLS-based RT.
  - The **isid-route-target** command has impact only on IMET/ISID and BMAC/ISID, not on IMET/0, BMAC/0, or ES routes.
  - When a new ISID-based RT is added for an I-VPLS, a BGP update is sent for the existing IMET/ISID and BMAC/ISID routes. The new RT will be added when the routes are advertised.
- Imported IMET/ISID and BMAC/ISID routes:
  - When **isid-route-target** is enabled for an I-VPLS, BGP will start importing IMET/ISID routes and—if **ivpls-mac-flush>bgp-evpn>send-to-bvpls** is enabled—BMAC/ISID routes with ISID-based RTs.
  - ISID-based RTs are added for import operations when the I-VPLS is associated with the B-VPLS (regardless of the operational state of the I-VPLS) and/or when the static ISID has been added.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

  - Ensure that the ISID-based RTs are configured consistently in the network. The system does not keep a mapping of RTs and ISIDs for imported routes.
  - The system will not check the format of the received auto-derived RTs. Routes will be imported when the RT is on the list of RTs for the B-VPLS.
- When **isid-route-target** is configured for an I-VPLS, VSI import/export policies are blocked in the B-VPLS, whereas BGP import/export policies are allowed and matching on the export ISID-based RT is supported.

Some other considerations:

- ISID ranges cannot overlap within a B-VPLS, but they can overlap across different B-VPLSs.
- The explicitly configured RT is meant to be used in two cases:
  - ISID aggregation - when multiple ISIDs are using the same ISID RT
  - Interoperability - in case the peer sends an RT in a different format

ISID-based RTs and RT-constraint

The use of the RT-constraint feature (BGP family route-target) maximizes the benefits of using different RTs per ISID; therefore, service providers are expected to enable both ISID-based RTs and RT-constraint. RT-constraint is enabled by adding the BGP address family route-target in the general BGP settings, per group, or per neighbor, as follows:

```
configure {
    router "Base" {
        bgp {
            family {
                route-target true
                ---snip---

            group "internal" {
                family {
                    route-target true
                    ---snip---

            neighbor 192.0.2.4 {
                family {
                    route-target true
                    ---snip---
```

The system will advertise the RT-constraint route when the I-VPLS is associated with the B-VPLS, regardless of the operational state of the I-VPLS. However, the IMET/ISID and the BMAC/ISID routes are sent based on the I-VPLS operational state.

Configuration

Figure 232: Example topology shows the example topology with three PEs and an RR.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 232: Example topology

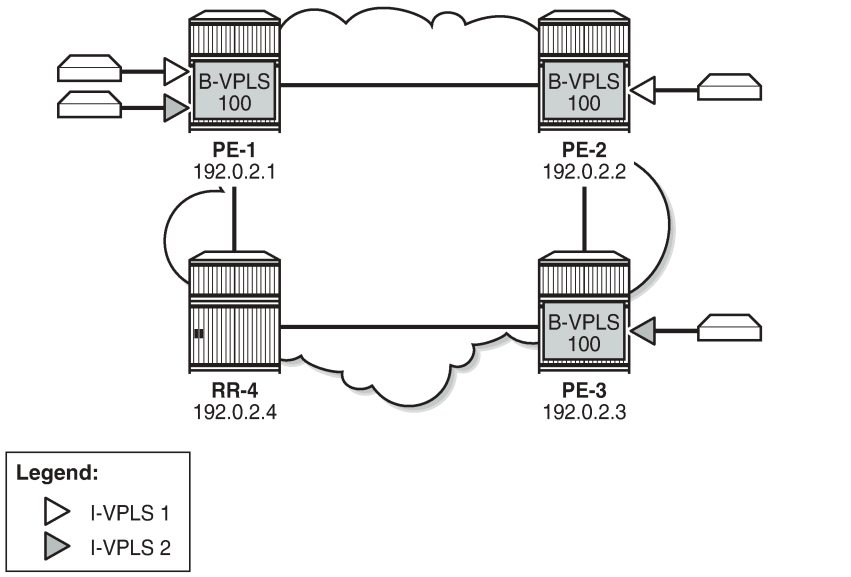

Initial configuration

The initial configuration on the nodes includes the following:

- Cards, MDAs, ports
- Router interfaces
- IS-IS enabled on all router interfaces (alternatively, OSPF could be used)
- SR-ISIS enabled on the PEs (but disabled on the RR)

BGP is configured on all PEs for address family EVPN, as follows.

```
# on PE-1, PE-2, PE-3:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            rapid-withdrawal true
            split-horizon true
            family {
                ipv4 false
                evpn true
            }
            rapid-update {
                evpn true
            }
            group "internal" {
                peer-as 64500
            }
            neighbor "192.0.2.4" {
                group "internal"
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
```

On RR-4, BGP is configured as follows:

```
# on RR-4:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            rapid-withdrawal true
            split-horizon true
            family {
                ipv4 false
                evpn true
            }
            rapid-update {
                evpn true
            }
            group "internal" {
                peer-as 64500
                cluster {
                    cluster-id 1.1.1.1
                }
            }
            neighbor "192.0.2.1" {
                group "internal"
            }
            neighbor "192.0.2.2" {
                group "internal"
            }
            neighbor "192.0.2.3" {
                group "internal"
            }
        }
```

For the RT-constraint feature, the route-target address family can be configured in combination with the EVPN address family; see section ISID-based RTs and RT-constraint.

The initial service configuration on PE-1 without ISID-based RTs is as follows:

```
# on PE-1:
configure {
    service {
        system {
            bgp-auto-rd-range {
                ip-address 192.0.2.1
                community-value {
                    start 10
                    end 99
                }
            }
        }
        vpls "B-VPLS 100" {
            admin-state enable
            service-id 100
            customer "1"
            service-mtu 2000
            pbb-type b-vpls
            pbb {
                source-bmac {
                    address 00:00:00:00:00:01
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            bgp 1 {
            }
            bgp-evpn {
                evi 100
                mpls 1 {
                    admin-state enable
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
        }
        vpls "I-VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS 100" {
                    isid 1
                }
            }
            bgp 1 {
                route-distinguisher auto-rd
                route-target {
                    export "target:64500:1"
                    import "target:64500:1"
                }
            }
            sap 1/2/1:1 {
            }
        }
        vpls "I-VPLS 2" {
            admin-state enable
            service-id 2
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS 100" {
                    isid 2
                }
            }
            bgp 1 {
                route-distinguisher auto-rd
                route-target {
                    export "target:64500:2"
                    import "target:64500:2"
                }
            }
            sap 1/2/1:2 {
            }
        }
```

The service configuration on PE-2 is similar, but only I-VPLS 1 is configured. On PE-3, only I-VPLS 2 is configured.

PE-1 sends the following default BGP-EVPN IMET/0 update to the RR:

```
# on PE-1:
2 2021/05/28 08:55:18.406 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.4
"Peer 1: 192.0.2.4: UPDATE
Peer 1: 192.0.2.4 - Send BGP UPDATE:
    Withdrawn Length = 0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    Total Path Attr Length = 77
    Flag: 0x90 Type: 14 Len: 28 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.1
        Type: EVPN-INCL-MCAST Len: 17 RD: 192.0.2.1:100, tag: 0, orig_addr len: 32,
                              orig_addr: 192.0.2.1
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:100
        bgp-tunnel-encap:MPLS
    Flag: 0xc0 Type: 22 Len: 9 PMSI:
        Tunnel-type Ingress Replication (6)
        Flags: (0x0)[Type: None BM: 0 U: 0 Leaf: not required]
        MPLS Label 8388560
        Tunnel-Endpoint 192.0.2.1
"
```

The following BGP-EVPN IMET routes are received on PE-1. Toward each other PE, there is a route with Ethernet tag 0; toward PE-2, there is a route with Ethernet tag 1 for ISID 1; toward PE-3, there is a route with Ethernet tag 2.

```
[/]
A:admin@PE-1# show router bgp routes evpn incl-mcast
===============================================================================
 BGP Router ID:192.0.2.1        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Inclusive-Mcast Routes
===============================================================================
Flag  Route Dist.         OrigAddr
      Tag                 NextHop
-------------------------------------------------------------------------------
u*>i  192.0.2.2:100       192.0.2.2
      0                   192.0.2.2

u*>i  192.0.2.2:100       192.0.2.2
      1                   192.0.2.2

u*>i  192.0.2.3:100       192.0.2.3
      0                   192.0.2.3

u*>i  192.0.2.3:100       192.0.2.3
      2                   192.0.2.3

-------------------------------------------------------------------------------
Routes : 4
===============================================================================
```

All these routes have a B-VPLS-based RT equal to 64500:100, as follows:

```
[/]
A:admin@PE-1# show router bgp routes evpn incl-mcast detail | match Community
Community      : target:64500:100 bgp-tunnel-encap:MPLS
Community      : target:64500:100 bgp-tunnel-encap:MPLS
Community      : target:64500:100 bgp-tunnel-encap:MPLS
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Community      : target:64500:100 bgp-tunnel-encap:MPLS
Community      : target:64500:100 bgp-tunnel-encap:MPLS
Community      : target:64500:100 bgp-tunnel-encap:MPLS
Community      : target:64500:100 bgp-tunnel-encap:MPLS
Community      : target:64500:100 bgp-tunnel-encap:MPLS
```

In the preceding output, each of the four inclusive multicast routes occurs twice: the first time with the original attributes, the second time with the modified attributes, but in this example, the attribute did not change.

For the EVPN MAC routes, the output is similar. ISID-based CMAC flush is not enabled yet, so there are only BMAC/0 routes, no BMAC/ISID routes, as follows:

```
[/]
A:admin@PE-1# show router bgp routes evpn mac
===============================================================================
 BGP Router ID:192.0.2.1        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN MAC Routes
===============================================================================
Flag  Route Dist.         MacAddr           ESI
      Tag                 Mac Mobility      Label1
                          Ip Address
                          NextHop
-------------------------------------------------------------------------------
u*>i  192.0.2.2:100       00:00:00:00:00:02 ESI-0
      0                   Static            LABEL 524285
                          n/a
                          192.0.2.2

u*>i  192.0.2.3:100       00:00:00:00:00:03 ESI-0
      0                   Static            LABEL 524285
                          n/a
                          192.0.2.3

-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

Both EVPN MAC routes have the same B-VPLS-based RT with value 64500:100, as follows:

```
[/]
A:admin@PE-1# show router bgp routes evpn mac detail | match Community
Community      : target:64500:100 bgp-tunnel-encap:MPLS
Community      : target:64500:100 bgp-tunnel-encap:MPLS
Community      : target:64500:100 bgp-tunnel-encap:MPLS
Community      : target:64500:100 bgp-tunnel-encap:MPLS
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

ISID-based RTs

On the PEs, B-VPLS 100 is configured with ISID-based RTs, but initially without ISID-based CMAC flush, as follows:

```
# on PE-1, PE-2:
configure {
    service {
        vpls "B-VPLS 100" {
            bgp-evpn {
                isid-route-target {
                    range 1 {
                        end 2
                    }
                    range 10 {
                        end 11
                        type configured
                        route-target "target:64500:10"
                    }
                }
            }
```

B-VPLS 100 has two ISID-ranges configured:

- For ISIDs 1 and 2, the RT is auto-derived. The hexadecimal value for ISID 1 is 0x30000001, which corresponds to decimal value 805306369. The hexadecimal value for ISID 2 is 0x30000002 (decimal value 805306370). For ISID 1, the RT is 64500: 805306369; for ISID 2, the RT is 64500: 805306370.
- For ISIDs 10 and 11, the RT is manually configured as 64500:10.

The configuration is identical on PE-2. On PE-3, only ISID range 2 is configured, as follows:

```
# on PE-3:
configure {
    service {
        vpls "B-VPLS 100" {
            bgp-evpn {
                isid-route-target {
                    range 2 {
                        end 2
                    }
                }
            }
```

On PE-1, the same four BGP-EVPN IMET routes are shown, as follows:

```
[/]
A:admin@PE-1# show router bgp routes evpn incl-mcast
===============================================================================
 BGP Router ID:192.0.2.1        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Inclusive-Mcast Routes
===============================================================================
Flag  Route Dist.         OrigAddr
      Tag                 NextHop
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
u*>i  192.0.2.2:100       192.0.2.2
      0                   192.0.2.2

u*>i  192.0.2.2:100       192.0.2.2
      1                   192.0.2.2

u*>i  192.0.2.3:100       192.0.2.3
      0                   192.0.2.3

u*>i  192.0.2.3:100       192.0.2.3
      2                   192.0.2.3

-------------------------------------------------------------------------------
Routes : 4
```

The IMET route with Ethernet tag 1 now has RT 64500:805306369 (ISID 1) and the IMET route with Ethernet tag 2 has RT 64500:805306370 (ISID 2), as follows:

```
[/]
A:admin@PE-1# show router bgp routes evpn incl-mcast detail | match Community
Community      : target:64500:100 bgp-tunnel-encap:MPLS
Community      : target:64500:100 bgp-tunnel-encap:MPLS
Community      : target:64500:805306369 bgp-tunnel-encap:MPLS
Community      : target:64500:805306369 bgp-tunnel-encap:MPLS
Community      : target:64500:100 bgp-tunnel-encap:MPLS
Community      : target:64500:100 bgp-tunnel-encap:MPLS
Community      : target:64500:805306370 bgp-tunnel-encap:MPLS
Community      : target:64500:805306370 bgp-tunnel-encap:MPLS
```

Again, each route has two identical entries in the preceding command: one with the original attributes and another with the modified attributes.

The following BGP-EVPN IMET/ISID route is sent by PE-1 for ISID 1. The Ethernet tag is 1 and the RT is 64500:805306369.

```
# on PE-1:
11 2021/05/28 08:59:47.220 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.4
"Peer 1: 192.0.2.4: UPDATE
Peer 1: 192.0.2.4 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 77
    Flag: 0x90 Type: 14 Len: 28 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.1
        Type: EVPN-INCL-MCAST Len: 17 RD: 192.0.2.1:100, tag: 1, orig_addr len: 32,
                              orig_addr: 192.0.2.1
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:805306369
        bgp-tunnel-encap:MPLS
    Flag: 0xc0 Type: 22 Len: 9 PMSI:
        Tunnel-type Ingress Replication (6)
        Flags: (0x0)[Type: None BM: 0 U: 0 Leaf: not required]
        MPLS Label 8388560
        Tunnel-Endpoint 192.0.2.1
"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The following BGP-EVPN IMET/ISID route is sent by PE-1 for ISID 2. The Ethernet tag is 2 and the RT is 64500:805306370.

```
# on PE-1:
12 2021/05/28 08:59:47.220 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.4
"Peer 1: 192.0.2.4: UPDATE
Peer 1: 192.0.2.4 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 77
    Flag: 0x90 Type: 14 Len: 28 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.1
        Type: EVPN-INCL-MCAST Len: 17 RD: 192.0.2.1:100, tag: 2, orig_addr len: 32,
                              orig_addr: 192.0.2.1
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:805306370
        bgp-tunnel-encap:MPLS
    Flag: 0xc0 Type: 22 Len: 9 PMSI:
        Tunnel-type Ingress Replication (6)
        Flags: (0x0)[Type: None BM: 0 U: 0 Leaf: not required]
        MPLS Label 8388560
        Tunnel-Endpoint 192.0.2.1
"
```

When a SAP (or SDP binding) is added with static ISID 11, RT 64500:10 will be added. The service configuration on PE-1 is modified as follows:

```
# on PE-1:
configure {
    service {
        vpls "B-VPLS 100" {
            bgp-evpn {
                isid-route-target {
                    range 1 {
                        end 2
                    }
                    range 10 {
                        end 11
                        type configured
                        route-target "target:64500:10"
                    }
                }
            }
            sap 1/1/1:100 {
                static-isid {
                    range 1 {
                        start 11
                        end 11
                    }
                }
            }
            isid-policy {
                entry 10 {
                    range {
                        start 11
                        end 11
                    }
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The configuration is similar on PE-2. Only on PE-1 and PE-2, SAPs are configured, with static ISID 11. The following IMET/ISID route with RT 64500:10 is sent by PE-1:

```
# on PE-1:
13 2021/05/28 08:59:47.251 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.4
"Peer 1: 192.0.2.4: UPDATE
Peer 1: 192.0.2.4 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 77
    Flag: 0x90 Type: 14 Len: 28 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.1
        Type: EVPN-INCL-MCAST Len: 17 RD: 192.0.2.1:100, tag: 11, orig_addr len: 32,
                              orig_addr: 192.0.2.1
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:10
        bgp-tunnel-encap:MPLS
    Flag: 0xc0 Type: 22 Len: 9 PMSI:
        Tunnel-type Ingress Replication (6)
        Flags: (0x0)[Type: None BM: 0 U: 0 Leaf: not required]
        MPLS Label 8388560
        Tunnel-Endpoint 192.0.2.1
"
```

This RT 64500:10 is not auto-derived, but configured manually for ISID range 10 to 11.

ISID-based CMAC flush

ISID-based CMAC flush is described in chapter PBB-EVPN ISID-based CMAC Flush and requires the following configuration on PE-1:

```
# on PE-1:
configure {
    service {
        vpls "I-VPLS 1" {
            pbb {
                i-vpls-mac-flush {
                    bgp-evpn {
                        send-to-bvpls true
                    }
                }
            }
        }
        vpls "I-VPLS 2" {
            pbb {
                i-vpls-mac-flush {
                    bgp-evpn {
                        send-to-bvpls true
                    }
                }
            }
        }
        vpls "B-VPLS 100" {
            bgp-evpn {
                accept-ivpls-evpn-flush true
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The configuration on PE-2 and PE-3 is similar, but only needs to be applied for I-VPLS 1 on PE-2 (I-VPLS 2 is not configured on PE-2) and for I-VPLS 2 on PE-3. The configuration for B-VPLS 100 is the same on all PEs.

When ISID-based CMAC flush is enabled on the PEs, additional BGP-EVPN MAC routes are sent by PE-1 for ISIDs 1 and 2:

```
# on PE-1:
27 2021/05/28 09:02:38.769 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.4
"Peer 1: 192.0.2.4: UPDATE
Peer 1: 192.0.2.4 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 89
    Flag: 0x90 Type: 14 Len: 44 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.1
        Type: EVPN-MAC Len: 33 RD: 192.0.2.1:100 ESI: ESI-0, tag: 2, mac len: 48
                       mac: 00:00:00:00:00:01, IP len: 0, IP: NULL, label1: 8388560
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:805306370
        bgp-tunnel-encap:MPLS
        mac-mobility:Seq:0/Static
"

25 2021/05/28 09:02:38.769 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.4
"Peer 1: 192.0.2.4: UPDATE
Peer 1: 192.0.2.4 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 89
    Flag: 0x90 Type: 14 Len: 44 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.1
        Type: EVPN-MAC Len: 33 RD: 192.0.2.1:100 ESI: ESI-0, tag: 1, mac len: 48
                       mac: 00:00:00:00:00:01, IP len: 0, IP: NULL, label1: 8388560
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:805306369
        bgp-tunnel-encap:MPLS
        mac-mobility:Seq:0/Static
"
```

The BGP-EVPN MAC routes for ISIDs 1 and 2 use the same auto-derived RT values as the IMET/ISID routes. The following four BGP-EVPN MAC routes are received in PE-1:

```
[/]
A:admin@PE-1# show router bgp routes evpn mac
===============================================================================
 BGP Router ID:192.0.2.1        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN MAC Routes
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Flag  Route Dist.          MacAddr           ESI
      Tag                  Mac Mobility      Label1
                           Ip Address
                           NextHop
-------------------------------------------------------------------------------
u*>i  192.0.2.2:100        00:00:00:00:00:02 ESI-0
      0                    Static            LABEL 524285
                           n/a
                           192.0.2.2

u*>i  192.0.2.2:100        00:00:00:00:00:02 ESI-0
      1                    Static            LABEL 524285
                           n/a
                           192.0.2.2

u*>i  192.0.2.3:100        00:00:00:00:00:03 ESI-0
      0                    Static            LABEL 524285
                           n/a
                           192.0.2.3

u*>i  192.0.2.3:100        00:00:00:00:00:03 ESI-0
      2                    Static            LABEL 524285
                           n/a
                           192.0.2.3

-------------------------------------------------------------------------------
Routes : 4
===============================================================================
```

The BMAC/0 routes have an RT based on the B-VPLS, whereas the BMAC/ISID routes have an RT derived from the ISID, as follows:

```
[/]
A:admin@PE-1# show router bgp routes evpn mac detail | match Community
Community      : target:64500:100 bgp-tunnel-encap:MPLS
Community      : target:64500:100 bgp-tunnel-encap:MPLS
Community      : target:64500:805306369 bgp-tunnel-encap:MPLS
Community      : target:64500:805306369 bgp-tunnel-encap:MPLS
Community      : target:64500:100 bgp-tunnel-encap:MPLS
Community      : target:64500:100 bgp-tunnel-encap:MPLS
Community      : target:64500:805306370 bgp-tunnel-encap:MPLS
Community      : target:64500:805306370 bgp-tunnel-encap:MPLS
```

ISID-based RTs and RT-constraint

To show that RT BGP updates are sent when the I-VPLS is associated with the B-VPLS, the I-VPLSs are initially disassociated from B-VPLS 100 on PE-1, as follows:

```
# on PE-1:
configure {
    service {
        vpls "I-VPLS 1" {
            pbb {
                delete backbone-vpls "B-VPLS 100"
            }
        }
        vpls "I-VPLS 2" {
            pbb {
                delete backbone-vpls "B-VPLS 100"
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
```

The BGP configuration is modified on all nodes to include address families route-target and EVPN, as follows:

```
# on PE-1, PE-2, PE-3, RR-4:
configure {
    router "Base" {
        bgp {
            family {
                route-target true
                evpn true
            }
```

The following RT-constraint route is sent by PE-1 after I-VPLS 1 is associated with B-VPLS 100. The RT is auto-derived from the ISID 1:

```
# on PE-1:
configure {
    service {
        vpls "I-VPLS 1" {
            pbb {
                backbone-vpls "B-VPLS 100" {
                    isid 1
                }
            }
        }
        vpls "I-VPLS 2"
            pbb {
                backbone-vpls "B-VPLS 100" {
                    isid 2
                }
            }
        }
```

```
# on PE-1:
73 2021/05/28 09:09:34.587 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.4
"Peer 1: 192.0.2.4: UPDATE
Peer 1: 192.0.2.4 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 47
    Flag: 0x90 Type: 14 Len: 22 Multiprotocol Reachable NLRI:
        Address Family RTC_V4
        NextHop len 4 NextHop 192.0.2.1
        [RT-Const-V4] origin-as 64500, Target target:64500:805306369
    Flag: 0x40 Type: 1 Len: 1 Origin: 2
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x80 Type: 4 Len: 4 MED: 0
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
"
```

When the I-VPLS goes operationally down, the IMET/ISID and BMAC/ISID routes are withdrawn, but not the RT-constraint route.

```
# on PE-1:
configure {
    service {
        vpls "I-VPLS 1" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            admin-state disable
```

```
# on PE-1:
83 2021/05/28 09:10:33.458 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.4
"Peer 1: 192.0.2.4: UPDATE
Peer 1: 192.0.2.4 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 61
    Flag: 0x90 Type: 15 Len: 57 Multiprotocol Unreachable NLRI:
        Address Family EVPN
        Type: EVPN-INCL-MCAST Len: 17 RD: 192.0.2.1:100, tag: 1, orig_addr len: 32,
                      orig_addr: 192.0.2.1
        Type: EVPN-MAC Len: 33 RD: 192.0.2.1:100 ESI: ESI-0, tag: 1, mac len: 48
                      mac: 00:00:00:00:00:01, IP len: 0, IP: NULL, label1: 0
"
```

The RT-constraint route is withdrawn when the I-VPLS is disassociated from B-VPLS 100, as follows:

```
# on PE-1:
configure {
    service {
        vpls "I-VPLS 1" {
            pbb {
                delete backbone-vpls "B-VPLS 100"
```

```
# on PE-1:
84 2021/05/28 09:11:28.205 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.4
"Peer 1: 192.0.2.4: UPDATE
Peer 1: 192.0.2.4 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 20
    Flag: 0x90 Type: 15 Len: 16 Multiprotocol Unreachable NLRI:
        Address Family RTC_V4
        [RT-Const-V4] origin-as 64500, Target target:64500:805306369
"
```

Conclusion

PBB-EVPN ISID-based RTs, in combination with RT-constraint, reduce the number of advertised IMET routes to only those nodes where the ISID is configured. The ISID-based RT can be auto-derived from the ISID or configured manually. When ISID-based CMAC flush is also enabled, the BMAC/ISID routes will contain the same auto-derived RT.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

PBB-VPLS

This chapter provides information about Provider Backbone Bridging (PBB) in a Multi-Protocol Label Switching (MPLS) based network.

Topics in this chapter include:

Applicability

This chapter is applicable to SR OS and was initially written for SR OS Release 7.0.R6. The MD-CLI in the current edition is based on SR OS Release 20.10.R2.

Note:
Although it can be used in an MPLS-based PBB network as described in this document, the MAC notification feature for dual-homed access is normally used in native PBB networks.

Overview

RFC 7041, *Extensions to the Virtual Private LAN Service (VPLS) Provider Edge (PE) Model for Provider Backbone Bridging*, describes the PBB-VPLS model supported by SR OS. This model expands the VPLS PE model to support PBB as defined by the IEEE 802.1ah.

PBB-VPLS combines the best of the PBB and VPLS technologies to deliver the most scalable multi-point Layer 2 VPN in the market. PBB-VPLS inherits all the benefits derived from MPLS (for example, sub-50ms Fast Reroute (FRR) protection, Traffic Engineering (TE), no need for Multiple Spanning Tree Protocol (MSTP) in the backbone) while greatly increasing the scalability of the network by providing MAC hiding, service multiplexing, and pseudowire aggregation.

The SR OS PBB-VPLS implementation also includes support for:

- Multiple MAC Registration Protocol (MMRP), application within IEEE 802.1ak for flood containment in the backbone instances, as specified in Section 6 of RFC 7041.
- Extensions to LDP signaling for PBB-VPLS, according to *draft-balus-l2vpn-pbb-ldp-ext-00*. These extensions avoid network black-hole issues, as described in the Section 3 of the mentioned draft.

This chapter describes how to configure and troubleshoot a PBB-VPLS network.

Knowledge of the VPLS and H-VPLS (RFC 4762, *Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling*) architecture and functionality is assumed throughout this chapter. The most relevant concepts are briefly described in this chapter. For further information, see the relevant Nokia documentation.

Figure 233: Example topology including B-VPLS, I-VPLSs, and protocol stacks shows the example topology that are used throughout the rest of the chapter.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 233: Example topology including B-VPLS, I-VPLSs, and protocol stacks

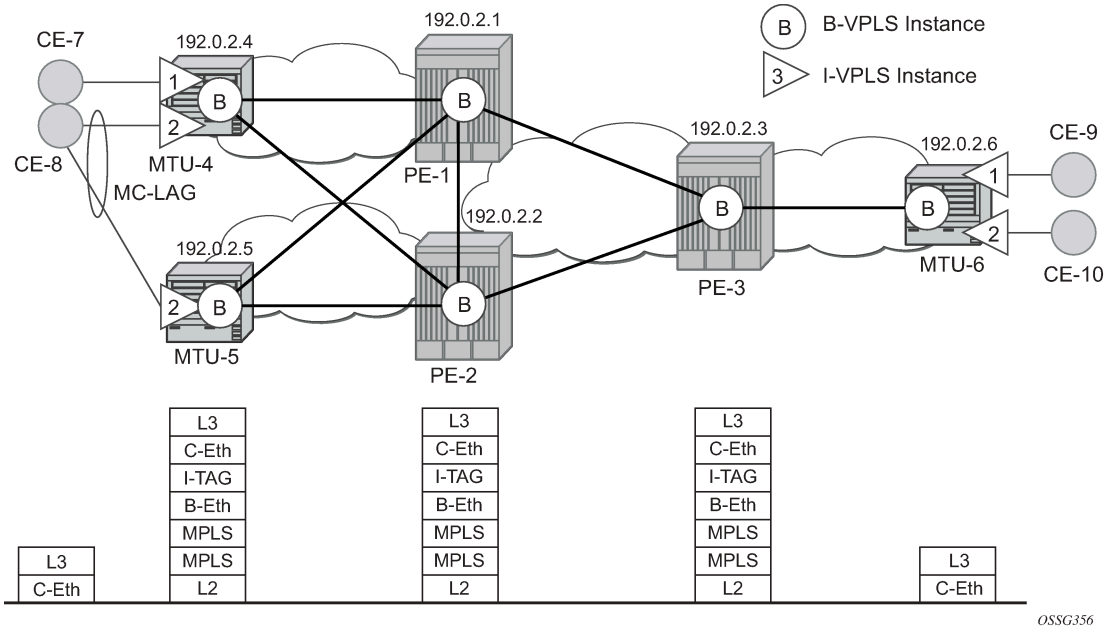

The topology consists of three core nodes (PE-1, PE-2, and PE-3) and three Multi-Tenant Unit (MTU) nodes connected to the core. A backbone VPLS instance (B-VPLS 100) will be defined in all the six nodes, whereas a few customer I-VPLS instances will be defined on the three MTU nodes.

Those I-VPLS instances will be multiplexed into the common B-VPLS, using the ISID field within the I-TAG as the demultiplexer field at the egress MTU to differentiate each specific customer.

The B-VPLS domain constitutes an H-VPLS network itself, with spoke-SDPs from the MTUs to the core PE layer. Active/standby spoke-SDPs can be used from the MTUs to the PEs (for example, in the MTU-4 and MTU-5 cases) or single non-redundant spoke-SDPs (for example, MTU-6). CE-8 is dual-connected to the service provider network through MC-LAG.

The protocol stack being used along the path between the CEs is shown in Figure 233: Example topology including B-VPLS, I-VPLSs, and protocol stacks.

Configuration

This section describes all the relevant PBB-VPLS configuration tasks for the setup shown in Figure 233: Example topology including B-VPLS, I-VPLSs, and protocol stacks. The appropriate associated IP/MPLS configuration is out of the scope of this example. In this particular example, the following protocols will be configured beforehand:

- ISIS-TE as IGP with all the interfaces being Level-2 (OSPF-TE could have been used instead).
- RSVP-TE as the MPLS protocol to signal the transport tunnels (LDP could have been used instead).
- LSPs between core PEs will be fast reroute protected (facility bypass tunnels) whereas LSP tunnels between MTUs and PEs will not be protected.
- The protection between MTU-4, MTU-5 and PE-1, PE-2 will be based on the active/standby pseudowire protection configured in the B-VPLS.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- BGP is configured for auto-discovery (Layer 2-VPN family), because FEC 129 will be used for the pseudowires between PEs in the core.

When the IP/MPLS infrastructure is up and running, the service configuration tasks described in the following sections can be implemented.

PBB-VPLS M:1 service configuration

This section describes the process to configure PBB-VPLS services in a M:1 fashion, M being the number of customer I-VPLS services multiplexed into the same B-VPLS instance (instance 100). An alternative configuration is 1:1, where each customer I-VPLS has its own B-VPLS. MTU-4 and PE-1 will be picked to show the relevant CLI configuration commands. The bold digits separated by colons **00:xx** are abbreviations for the backbone MAC addresses.

Figure 234: Example topology with port numbers and IP addresses

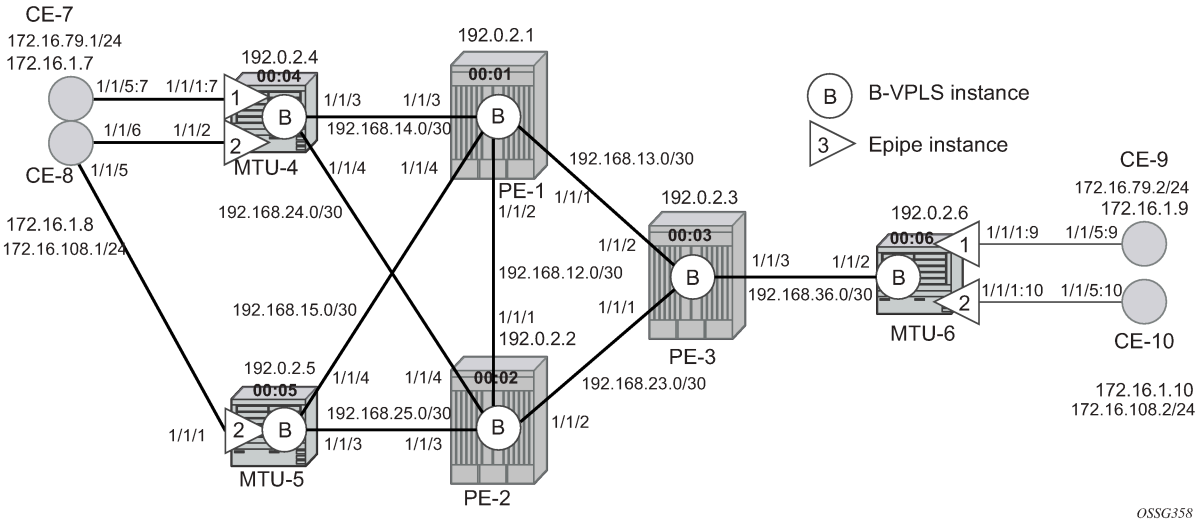

B-VPLS configuration

The first step is to configure the B-VPLS instance that will carry the PBB traffic. The following shows the B-VPLS configuration on MTU-4 and PE-1. The configuration on MTU-5 and MTU-6 resembles the configuration on MTU-4; the configuration on PE-2 and PE-3 resembles the configuration on PE-1.

The configuration for B-VPLS 100 on MTU-4 is as follows:

```
# on MTU-4:
configure {
    service {
        vpls "B-VPLS 100" {
            admin-state enable
            service-id 100
            customer "1"
            service-mtu 2000
            pbb-type b-vpls
            pbb {
                source-bmac {
                    address 00:04:04:04:04:04
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
            }
            endpoint "core" {
                suppress-standby-signaling false
            }
            spoke-sdp 41:100 {
                endpoint {
                    name "core"
                    precedence primary
                }
                stp {
                    admin-state disable
                }
            }
            spoke-sdp 42:100 {
                endpoint {
                    name "core"
                }
                stp {
                    admin-state disable
                }
            }
        }
```

On PE-1, B-VPLS 100 is configured as follows:

```
# on PE-1:
configure {
    service {
        pw-template "PW1" {
            pw-template-id 1
            provisioned-sdp use
            split-horizon-group {
                name "CORE"
            }
        }
        vpls "B-VPLS 100" {
            admin-state enable
            service-id 100
            customer "1"
            service-mtu 2000
            pbb-type b-vpls
            pbb {
                source-bmac {
                    address 00:01:01:01:01:01
                }
            }
            bgp 1 {
                route-target {
                    export "target:65000:100"
                    import "target:65000:100"
                }
                pw-template-binding "PW1" {
                }
            }
            bgp-ad {
                admin-state enable
                vpls-id "65000:100"
            }
            spoke-sdp 14:100 {
            }
            spoke-sdp 15:100 {
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
```

The B-VPLS is a regular VPLS instance in terms of configuration, with the following exceptions:

- The B-VPLS service MTU must be at least 18 bytes greater than the I-VPLS MTU of the multiplexed instances. In this example, the I-VPLS instances will have the default service MTU (1500 bytes); therefore, any MTU equal to or greater than 1518 bytes must be configured. In this particular example, a MTU of 2000 bytes is configured in the B-VPLS instance throughout the network.
- The source B-MAC is the MAC address that will be sourced when the PBB traffic is originated from that node. A source B-MAC per B-VPLS instance can be configured (if there are more than one B-VPLS) or a common source B-MAC that will be shared by all the B-VPLS instances in the node. If no specific source B-MAC is provisioned, the system MAC address is used as the source B-MAC. When using the access multi-homing feature for native PBB, the source B-MAC must be a configured one and never the chassis MAC address. The way to configure a common B-MAC for all the B-VPLS instances on MTU-4 is as follows:

```
# on MTU-4:
configure {
    service {
        pbb {
            source-bmac {
                address 00:04:04:04:04:04
            }
```

The following considerations will be taken into account when configuring the B-VPLS:

- B-VPLS SAPs:
  - Ethernet null, dot1q, and qinq encapsulations are supported
  - Default SAP (:*) types are blocked in the CLI for the B-VPLS SAP
- B-VPLS SDPs:
  - For MPLS, both mesh and spoke-SDPs with split-horizon groups are supported.
  - Similar to regular pseudowires, the outgoing PBB frame on an SDP (for example, B-pseudowire) contains a BVID qtag only if the pseudowire type is Ethernet VLAN. If the pseudowire type is **Ethernet**, the BVID q-tag is stripped before the frame goes out.
  - BGP-AD is supported in the B-VPLS; therefore, spoke-SDPs in the B-VPLS can be signaled using FEC 128 or FEC 129. In this example, BGP-AD and FEC 129 are used. A split-horizon group (SHG) has been configured to emulate the behavior of mesh-SDPs in the core.
- If a local I-VPLS instance is associated with the B-VPLS, local frames originated/terminated on local I-VPLS(s) are PBB encapsulated/de-encapsulated using the PBB Ethertype provisioned under the related port or SDP component.

By default, the PBB Ethertype is 0x88e7 (which is the standard one defined in 802.1ah for the I-TAG) but this PBB Ethertype can be changed if required due to interoperability reasons. This is the way to change it at port and/or SDP level:

```
[ex:configure port 1/1/3 ethernet]
A:admin@MTU-4# pbb-etype ?

 pbb-etype <number>
 <number> - <0x600..0xffff>
 Default  - 35047
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    Ethertype for PBB encapsulation on the Ethernet port
```

```
[ex:configure service sdp 41]
A:admin@MTU-4# pbb-etype ?

 pbb-etype <number>
 <number> - <0x600..0xffff>
 Default  - 0x88E7

    Ethertype used in frames sent out on this SDP when VC type is 'vlan' for
    Provider Backbone Bridging frames as 0xXXYY with range 0x0600-0xFFFF.
```

The following commands are useful to check the actual PBB Ethertype:

```
[]
A:admin@MTU-4# show port 1/1/3 | match PBB
PBB Ethertype      : 0x88e7
```

```
[]
A:admin@MTU-4# show service sdp 41 detail | match PBB
Bw BookingFactor     : 100                   PBB Etype          : 0x88e7
```

I-VPLS configuration

When the common B-VPLS is configured, the next step is to provision the customer I-VPLS instances. The following shows the relevant configuration on MTU-4 for the two I-VPLS instances represented in Figure 234: Example topology with port numbers and IP addresses. The I-VPLS instances are configured on the MTU devices, whereas the core PEs are customer-unaware nodes.

```
# on MTU-4:
configure {
    service {
        vpls "I-VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS 100" {
                    isid 1
                }
            }
            sap 1/1/1:7 {
            }
        }
        vpls "I-VPLS 2" {
            admin-state enable
            service-id 2
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS 100" {
                    isid 2
                }
            }
            sap lag-1 {
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
```

The I-VPLS instance has to be linked to its corresponding transport B-VPLS instance. That link is specified by the **backbone-vpls <b-vpls> isid <isid>** command. The ISID is mandatory when configuring the backbone VPLS.

The following considerations will be taken into account when configuring the I-VPLS:

- I-VPLS SAPs:
  - SAPs can be defined on ports with any Ethernet encapsulation type (null, dot1q, and qinq)
  - The I-VPLS SAPs can coexist on the same port with SAPs for other business services, for example, VLL and VPLS SAPs.
- I-VPLS SDPs:
  - GRE and MPLS SDPs are supported.
  - No mesh-SDPs are supported, only spoke-SDP. Mesh-SDPs can be emulated by using SHGs.

Existing SAP processing rules still apply for the I-VPLS case; the SAP encapsulation definition on Ethernet ingress ports defines which VLAN tags are used to determine the service that the packet belongs to:

- Null encapsulation defined on ingress — Any VLAN tags are ignored and the packet goes to a default service for the SAP.
- Dot1q encapsulation defined on ingress — only first VLAN tag is considered.
- QinQ encapsulation defined on ingress — both VLAN tags are considered; wildcard for the inner VLAN tag is supported.
- For dot1q/qinq encapsulations, traffic encapsulated with VLAN tags for which there is no definition is discarded.
- Any VLAN tag used for service selection on the I-SAP is stripped before the PBB encapsulation is added. Appropriate VLAN tags are added at the remote PBB PE when sending the packet out on the egress SAP.

MMRP for flooding optimization

When the M:1 model is used (as in this example), any I-VPLS broadcast, unknown unicast, or multicast (BUM) frame is flooded throughout the B-VPLS domain regardless of the nodes where the originating I-VPLS is defined. In other words, in our example in Figure 233: Example topology including B-VPLS, I-VPLSs, and protocol stacks, any BUM frame coming from CE-7 would be flooded in the B domain and would reach PE-2 and MTU-5, even though that traffic only needs to go to PE-3 and MTU-6. To build customer-based flooding trees and optimize the flooding, Multiple MAC Registration Protocol (MMRP) must be configured on the B-VPLS.

MMRP can be enabled with its default settings just by executing the following command on all nodes:

```
# on all nodes:
configure {
    service {
        vpls "B-VPLS 100" {
            mrp {
                admin-state enable
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

There are specific B-VPLS MRP settings that can be modified. These are the default values:

```
[ex:configure service vpls "B-VPLS 100" mrp]
A:admin@MTU-4# info detail
    admin-state enable
    mmrp {
        admin-state enable
        end-station-only false
     ## flood-time
        attribute-table {
            high-wmark 95
            low-wmark 90
            size 2048
        }
    }
```

These attributes can be changed to control the number of MMRP attributes per B-VPLS and optimize the convergence time in case of failures in the B-VPLS:

- Controlling the number of attributes per B-VPLS

  The MMRP exchanges create one entry per attribute (group B-MAC) in the B-VPLS where MMRP protocol is running. PBB uses a group B-MAC address—built using a specific OUI (00:1e:83) with the multicast bit set, and the ISID value for the last 24 bits—as a destination MAC address for flooding any BUM frame into the B-domain.

  When the first registration is received for an attribute, an MFIB entry is created for it. The **attribute-table size** allows the user to control the number of MMRP attributes (group B-MACs) created on a per B-VPLS basis, between 1 and 2048. Based on the configured size, high and low watermarks can be set (in percentage) so that alarms can be triggered upon exceeding the watermarks. This ensures that no B-VPLS will take up all the resources from the total pool. The maximum number of attributes per B-VPLS is 2048 and 4000 can be configured globally on the system.

- Optimizing the convergence time

  Assuming that MMRP is used in a certain B-VPLS, under failure conditions, the time it takes for the B-VPLS forwarding to resume may depend on the data plane and control plane convergence plus the time it takes for MMRP exchanges to stabilize the flooding trees on a per ISID basis. In order to minimize the convergence time, the PBB SR OS implementation offers the selection of a mode where B-VPLS forwarding reverts for a short time to flooding so that MMRP has enough time to converge. This mode can be selected through configuration using the **flood-time <value>** command where value represents the amount of time in seconds (between 3 and 600) that flooding will be enabled. If this behavior is selected, the forwarding plane starts with B-VPLS flooding for a configurable time period, then it reverts back to the MFIB entries installed by MMRP. The following B-VPLS events initiate the switch from per I-VPLS (MMRP) MFIB entries to B-VPLS flooding:

  – Reception or local triggering of a Spanning Tree Topology Change Notification (TCN)

  – B-SAP failure

  – Failure of a B-SDP binding

  – Pseudowire activation in a primary/standby H-VPLS resiliency solution

  – SF/CPM switchover due to STP reconvergence

The IEEE 802.1ak standard, which defines MRP, requires the implementation of different state machines with associated timers that can be tuned. A full MRP participant maintains the following state machines:

- Registrar state machine

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- Applicant state machine
- LeaveAll state machine
- PeriodicTransmission state machine

The two first state machines are maintained for each attribute in which the participant is interested, whereas the two latter are global to all the attributes.

The job of the registrar function is to record declarations of the attribute made by other participants on the LAN. A registrar does not send any protocol messages, because the applicant looks after the interests of all would-be participants.

The job of the applicant is twofold: first, to ensure that this participant's declaration is correctly registered by other participants' registrars, and next, to prompt other participants to register again after one withdraws a declaration.

The associated timers can be tuned on a per SAP/SDP basis:

```
[ex:configure service vpls "B-VPLS 100" spoke-sdp 41:100]
A:admin@MTU-4# mrp ?

 mrp

 apply-groups          - Apply a configuration group at this level
 apply-groups-exclude  - Exclude a configuration group at this level
 join-time             - Set the maximum rate for attribute join messages to be sent
                         on the SDP.
 leave-all-time        - Set the frequency where all attribute declarations on the
                         SDP are refreshed.
 leave-time            - Set the time an attribute is held in leave state before
                         registration is removed.
 periodic-time         - Set the freqeuency of retransmission of attribute
                         declarations.
 periodic-timer        - Enable/Disable retransmission of attribute declarations.
 policy                - Specify they MRP policy to control which Group BMAC
                         attributes will advertise on the egress SDP Bind.
```

```
[ex:configure service vpls "B-VPLS 100" spoke-sdp 41:100 mrp]
A:admin@MTU-4# info detail
 ## apply-groups
 ## apply-groups-exclude
    join-time 2
    leave-time 30
    leave-all-time 100
    periodic-time 10
    periodic-timer false
 ## policy
```

A brief description of the MRP SAP/SDP attributes follows:

- **Join-time** — This command controls the interval between transmit opportunities that are applied to the applicant state machine. An instance of this join period timer is required on a per-port, per-MRP participant basis. For more information, see IEEE 802.1ak-2007 section 10.7.4.1.
- **Leave-time** — This command controls the period of time that the registrar state machine will wait in the leave state before transitioning to the MT state when it is removed. An instance of the timer is required for each state machine that is in the leave state. The leave period timer is set to the value leave-time when it is started. A registration is normally in "in" state where there is an MFIB entry and traffic being forwarded. When a "leave all" is performed (periodically around every 10-15 seconds per SAP/SDP binding – see leave-all-time below), a node sends a message to its peer indicating a leave

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

all is occurring and puts all of its registrations in leave state. The peer refreshes its registrations based on the leave all PDU it receives and sends a PDU back to the originating node with the state of all its declarations. See IEEE 802.1ak-2007 section 10.7.4.2.

- **Leave-all-time** — This command controls the frequency with which the leaveall state machine generates leaveall PDUs. The timer is required on a per-port, per-MRP participant basis. The leaveall period timer is set to a random value, T, in the range leavealltime<T<1.5*leave-all-time when it is started. See IEEE 802.1ak-2007, section 10.7.4.3.
- **Periodic-time** — This command controls the frequency the periodic transmission state machine generates periodic events if the periodic transmission timer is enabled. The timer is required on a per-port basis. The periodic transmission timer is set to one second when it is started.
- **Periodic-timer** — This command enables or disables the periodic transmission timer.

The following command shows the MRP configuration and statistics on a per SAP/SDP basis within the B-VPLS:

```
[]
A:admin@MTU-4# show service id 100 all | match MRP post-lines 10
Sdp Id 41:100 MRP Information
-------------------------------------------------------------------------------
Join Time          : 0.2 secs                Leave Time        : 3.0 secs
Leave All Time     : 10.0 secs               Periodic Time     : 1.0 secs
Periodic Enabled   : false
Mrp Policy         : N/A
Rx Pdus            : 234                     Tx Pdus           : 252
Dropped Pdus       : 0
Rx New Event       : 0                       Rx Join-In Event  : 246
Rx In Event        : 0                       Rx Join Empty Evt : 217
Rx Empty Event     : 0                       Rx Leave Event    : 0
SDP MMRP Information
-------------------------------------------------------------------------------
MAC Address       Registered        Declared
-------------------------------------------------------------------------------
01:1e:83:00:00:01 Yes               Yes

01:1e:83:00:00:02 Yes               Yes

-------------------------------------------------------------------------------
Number of MACs=2 Registered=2 Declared=2
-------------------------------------------------------------------------------
Sdp Id 42:100 MRP Information
-------------------------------------------------------------------------------
Join Time          : 0.2 secs                Leave Time        : 3.0 secs
Leave All Time     : 10.0 secs               Periodic Time     : 1.0 secs
Periodic Enabled   : false
Mrp Policy         : N/A
Rx Pdus            : 0                       Tx Pdus           : 0
Dropped Pdus       : 0
Rx New Event       : 0                       Rx Join-In Event  : 0
Rx In Event        : 0                       Rx Join Empty Evt : 0
Rx Empty Event     : 0                       Rx Leave Event    : 0
SDP MMRP Information
-------------------------------------------------------------------------------
MAC Address       Registered        Declared
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Number of MACs=0 Registered=0 Declared=0
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Number of SDPs : 2
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
* indicates that the corresponding row element may have been truncated.
Service MRP Information
===============================================================================
Admin State          : enabled
-------------------------------------------------------------------------------
MMRP
-------------------------------------------------------------------------------
Admin Status         : enabled              Oper Status      : up
Register Attr Cnt    : 2                    Declared Attr Cnt: 2
End-station-only     : disabled
Max Attributes       : 2048                 Attribute Count  : 2
Hi Watermark         : 95%                  Low Watermark    : 90%
Failed Registers     : 0                    Flood Time       : Off
-------------------------------------------------------------------------------
MVRP
-------------------------------------------------------------------------------
MRP SAP Table
===============================================================================
SAP                            Join      Leave     Leave All Periodic
                               Time(sec) Time(sec) Time(sec) Time(sec)
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
MRP SDP-BIND Table
===============================================================================
SDP-BIND                       Join      Leave     Leave All Periodic
                               Time(sec) Time(sec) Time(sec) Time(sec)
-------------------------------------------------------------------------------
41:100                         0.2       3.0       10.0      1.0
42:100                         0.2       3.0       10.0      1.0
===============================================================================
===============================================================================

-------------------------------------------------------------------------------
```

The following command is useful to check the MRP configuration and status.

```
[]
A:admin@MTU-4# show service id 100 mrp

===============================================================================
Service MRP Information
===============================================================================
Admin State          : enabled
-------------------------------------------------------------------------------
MMRP
-------------------------------------------------------------------------------
Admin Status         : enabled              Oper Status      : up
Register Attr Cnt    : 2                    Declared Attr Cnt: 2
End-station-only     : disabled
Max Attributes       : 2048                 Attribute Count  : 2
Hi Watermark         : 95%                  Low Watermark    : 90%
Failed Registers     : 0                    Flood Time       : Off
-------------------------------------------------------------------------------
MVRP
-------------------------------------------------------------------------------
Admin Status         : disabled             Oper Status      : down
Max Attr             : 4095                 Failed Register  : 0
Register Attr Count  : 0                    Declared Attr    : 0
Hi Watermark         : 95%                  Low Watermark    : 90%
Hold Time            : disabled             Attr Count       : 0
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
MRP SAP Table
===============================================================================
SAP                           Join      Leave     Leave All Periodic
                              Time(sec) Time(sec) Time(sec) Time(sec)
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
MRP SDP-BIND Table
===============================================================================
SDP-BIND                      Join      Leave     Leave All Periodic
                              Time(sec) Time(sec) Time(sec) Time(sec)
-------------------------------------------------------------------------------
41:100                        0.2       3.0       10.0      1.0
42:100                        0.2       3.0       10.0      1.0
===============================================================================
===============================================================================
```

In the example throughout the chapter, as soon as MMRP is enabled, an optimized flooding tree will be built for ISID 1, because the I-VPLS 1 is only defined in MTU-4 and MTU-6, but not in MTU-5. A good way to track the flooding tree for a particular ISID is the following command:

```
[]
A:admin@MTU-4# show service id 100 mmrp mac
-------------------------------------------------------------------------------
SAP/SDP                               MAC Address       Registered  Declared
-------------------------------------------------------------------------------
sdp:41:100                            01:1e:83:00:00:01 Yes         Yes
sdp:41:100                            01:1e:83:00:00:02 Yes         Yes
-------------------------------------------------------------------------------
Number of Entries=2 SAPs=0 SDPs=2
-------------------------------------------------------------------------------
```

```
[]
A:admin@MTU-5# show service id 100 mmrp mac
-------------------------------------------------------------------------------
SAP/SDP                               MAC Address       Registered  Declared
-------------------------------------------------------------------------------
sdp:52:100                            01:1e:83:00:00:01 Yes         No
sdp:52:100                            01:1e:83:00:00:02 Yes         No
-------------------------------------------------------------------------------
Number of Entries=2 SAPs=0 SDPs=2
-------------------------------------------------------------------------------
```

The group B-MAC ending in **01** corresponds to the I-VPLS 1 whereas the one ending in **02** to the I-VPLS 2. MMRP PDUs for the two attributes are sent throughout the loop-tree topology (not over STP blocked ports or standby spoke-SDPs and observing the split-horizon rules). The two attributes are registered on every B-VPLS virtual port; however, the tree is only built on those ports where the attribute is also declared, and not only registered. For instance, the spoke-SDP 52:100 in MTU-5 will not be part of the ISID 1 or ISID 2 flooding trees. Neither attribute is declared because I-VPLS 1 does not exist on MTU-5 and I-VPLS 2 is operationally down on MTU-5 (MC-LAG SAP is in standby state, so the I-VPLS is down).

As soon as a group B-MAC attribute is registered on a particular port, an MFIB entry is added for that B-MAC on that port, regardless of the declaration state for that attribute on the port. For instance, neither

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

B-MAC is declared on MTU-5, however, the two MFIB entries are created as soon as the attributes are registered:

```
[]
A:admin@MTU-5# show service id 100 mfib

===============================================================================
Multicast FIB, Service 100
===============================================================================
Source Address  Group Address         Port Id                      Svc Id   Fwd
                                                                          Blk
-------------------------------------------------------------------------------
*               01:1e:83:00:00:01     b-sdp:52:100                 Local    Fwd
*               01:1e:83:00:00:02     b-sdp:52:100                 Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 2
===============================================================================
```

MAC flush: avoiding black-holes

Both the I-VPLS and B-VPLS components inherit the MAC flush capabilities of a regular VPLS clearing the related C-MAC and respectively B-MAC FIBs. All types of MAC flush—all-but-mine and all-from-me—are supported together with the related CLI. In addition to these features, some extensions have been added so that MAC flush can be triggered on the B-VPLS based on some events happening on the I-VPLS. The following diagram shows a potential scenario where black-holes can occur if the correct configuration is not added.

Figure 235: Black-hole

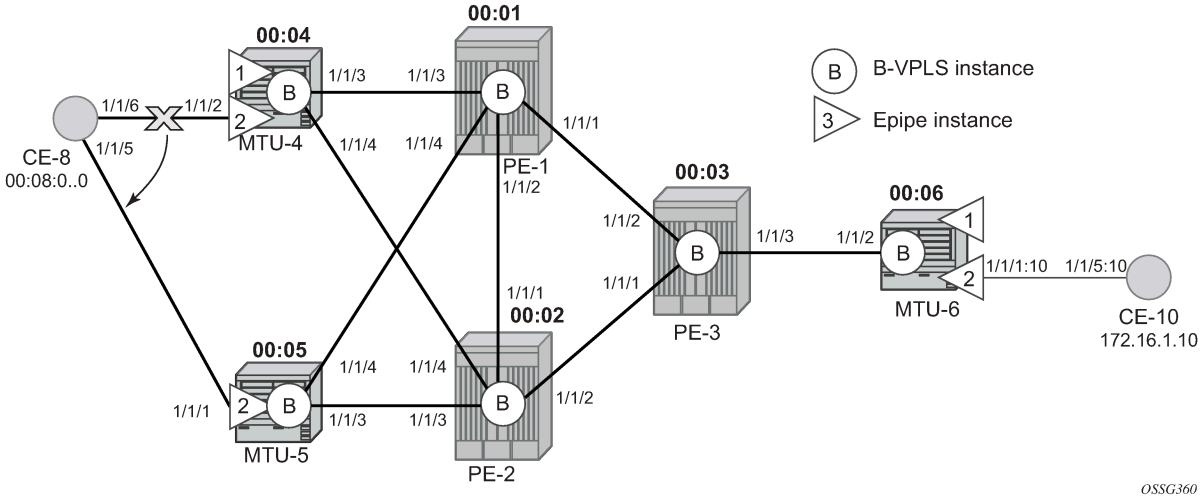

Under normal conditions, the I-VPLS 2 FIB on MTU-6 shows that CE-8 MAC address is learned through B-MAC 00:04 of MTU-4:

```
[]
A:admin@MTU-6# show service id 2 fdb pbb

===============================================================================
Forwarding Database, i-Vpls Service 2
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
MAC               Source-Identifier     B-Svc     b-Vpls MAC        Type/Age
 Transport:Tnl-Id
-------------------------------------------------------------------------------
00:08:00:00:00:00 b-sdp:63:100          100       00:04:04:04:04:04 L/60
00:10:00:00:00:00 sap:1/1/1:10          100       N/A               L/60
===============================================================================
```

When a failure happens in the CE-8 MC-LAG active link, the link to MTU-5 takes over. However, the FIB on MTU-6 still points at the BMAC of MTU-4 and that will still be the B-MAC used in the PBB encapsulation. Therefore, a black-hole occurs until either bidirectional traffic is sent or the FIB aging timer expires.

The configuration in the I-VPLS can be modified to trigger a MAC flush in the B-VPLS with the following command:

```
[ex:configure service vpls "I-VPLS 2" pbb i-vpls-mac-flush tldp]
A:admin@MTU-4# send-to-bvpls ?

 send-to-bvpls

 all-but-mine          - Generate LDP MAC withdraw message to b-VPLS
 all-from-me           - Generate LDP MAC withdraw all from me message to b-VPLS
```

The following command is executed on all MTUs to solve the black-hole:

```
# on MTU-4, MTU-5, MTU-6:
configure {
    service {
        vpls "I-VPLS 2" {
            pbb {
                i-vpls-mac-flush {
                    tldp {
                        send-to-bvpls {
                            all-from-me true
                        }
                    }
                }
            }
```

By configuring **send-to-bvpls all-from-me true** on I-VPLS 2, a failure on the MC-LAG active link on I-VPLS 2 will trigger an LDP MAC **flush-all-from-me** into the B-VPLS that will flush the FIB in MTU-6 for I-VPLS 2, avoiding the black-hole. A MC-LAG failure is emulated by disabling the LAG on MTU-4, as follows:

```
# on MTU-4:
configure {
    lag 1 {
        admin-state disable
    }
```

MTU-4 sends the following LDP MAC flush for all MAC addresses learned from MTU-4:

```
1 2021/01/12 17:02:25.211 UTC MINOR: DEBUG #2001 Base LDP
"LDP: LDP
Send Address Withdraw packet (msgId 263) to 192.0.2.1:0
Protocol version = 1
MAC Flush (All MACs learned from me)
Service FEC PWE3: ENET(5)/100 Group ID = 0 cBit = 0
Number of PBB-BMACs = 1
BMAC 1 = 00:04:04:04:04:04
Number of PBB-ISIDs = 1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
ISID 1 = 2
Number of Path Vectors : 1
Path Vector(  1) = 192.0.2.4
"
```

On MTU-6:

```
1 2021/01/12 17:02:25.227 UTC MINOR: DEBUG #2001 Base LDP
"LDP: LDP
Recv Address Withdraw packet (msgId 206) from 192.0.2.3:0
Protocol version = 1
MAC Flush (All MACs learned from me)
Service FEC PWE3: ENET(5)/100 Group ID = 0 cBit = 0
Number of PBB-BMACs = 1
BMAC 1 = 00:04:04:04:04:04
Number of PBB-ISIDs = 1
ISID 1 = 2
Number of Path Vectors : 3
Path Vector(  1) = 192.0.2.4
Path Vector(  2) = 192.0.2.1
Path Vector(  3) = 192.0.2.3
```

Immediately after receiving the MAC flush, the CE-8 MAC is flushed. The CE-8 MAC is learned again, but this time linked to the B-MAC 00:05, which is the B-MAC of MTU-5:

```
[]
A:admin@MTU-6# show service id 2 fdb pbb

===============================================================================
Forwarding Database, i-Vpls Service 2
===============================================================================
MAC               Source-Identifier     B-Svc     b-Vpls MAC        Type/Age
 Transport:Tnl-Id
-------------------------------------------------------------------------------
00:08:00:00:00:00 b-sdp:63:100          100       00:05:05:05:05:05 L/0
00:10:00:00:00:00 sap:1/1/1:10          100       N/A               L/120
===============================================================================
```

The following I-VPLS events are propagated into the B-VPLS depending on the all-but-mine or all-from-me keywords used in the configuration:

If the all-but-mine keyword is configured (positive flush), the following events in the I-VPLS trigger a MAC flush into the B-VPLS:

1. TCN event in one or more of the related I-VPLS/M-VPLS.
2. Pseudowire/SDP binding activation with active/standby pseudowire (standby to active or down to up).
3. Reception of an LDP MAC withdraw flush-all-but-mine in the related I-VPLS.

If the all-from-me keyword is configured (negative flush) the following events in the I-VPLS trigger a MAC flush into the B-VPLS:

1. MC-LAG active link failure (in our example).
2. Failure of a local SAP – requires **mac-flush>tldp>send-on-failure true** to be enabled in I-VPLS.
3. Failure of a local pseudowire/SDP binding – requires **mac-flush>tldp>send-on-failure true** to be enabled in I-VPLS.
4. Reception of an LDP MAC withdraws flush-all-from-me in the related I-VPLS.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

In addition to this and regardless of what type, MAC flush has been optimized to avoid flushing in the core PEs, flushing only the C-MACs mapped to a specific B-MAC (belonging to a specific ISID FIB) and the ability to indicate to core PEs which messages should always be forwarded endpoint-to-endpoint toward all PBB PEs regardless of the propagate-mac-flush setting in B-VPLS. All of this is implemented without the need of any additional CLI commands and it is part of **draft-balus-l2vpn-pbb-ldp-ext-00**.

Another extension supported to avoid black-holes within this mix of I- and B-VPLS environments is the **block-on-mesh-failure** feature in PBB. When the VPLS mesh exists only in I-VPLS or in B-VPLS, and the **block-on-mesh-failure** feature is enabled, the regular VPLS behavior will apply (when all the mesh-SDPs go down an LDP notification with pseudowire status bits = 0x01—Pseudo Wire Not Forwarding—is sent over the spoke-SDPs). When the active/standby pseudowire resiliency is implemented in I-VPLS such that the PBB PE performs the role of a PE-rs, the B-VPLS core replaces the pseudowire (SDP binding) mesh. The block-on-mesh notification (LDP notification indicating pseudowire not forwarding) will be sent to the MTUs only when the related B-VPLS is operationally down. The B-VPLS core is operationally down only when all of its SAPs and SDPs are down.

The final feature that can be enabled in an I-VPLS with CLI is the **send-on-bvpls-failure** feature, as follows:

```
# on MTU-4, MTU-5, MTU-6:
configure {
    service {
        vpls "I-VPLS 2" {
            pbb {
                i-vpls-mac-flush {
                    tldp {
                        send-on-bvpls-failure true
                        }
                    }
```

This feature is required to avoid black-holes when there is a full mesh of pseudowires in the I-VPLS domain and the B-VPLS instance can go operationally down. Figure 236: Send flush on B-VPLS failure example shows a typical scenario where this feature is needed (normally when PBB-VPLS and multi-chassis end point are combined together).

Figure 236: Send flush on B-VPLS failure example

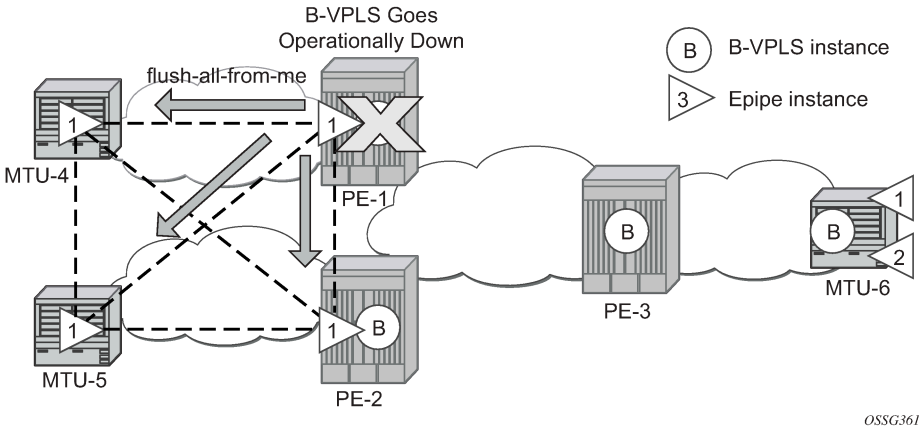

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Access dual-homing and MAC notification

Although this section is focused on PBB in a MPLS based network, the Nokia PBB implementation also allows the operator to use a native Ethernet infrastructure in the PBB core. Native Ethernet tunneling can be emulated using Ethernet SAPs to interconnect the related B-VPLS instances. In those cases, there is no LDP signaling available; therefore, there is no MAC flush sent when the active link in a multi-homed access device fails.

The SR OS supports a mechanism to avoid potential black-holes in native Ethernet PBB networks. In addition to the source B-MAC associated with each B-VPLS, an additional B-MAC is associated with each MC-LAG supporting Multi-homed I-VPLS SAPs. The nodes that are in a multi-homed MC-LAG configuration share a common B-MAC on the related MC-LAG interfaces. When the MAC notification is enabled, an Ethernet CFM notification message is sent from the node holding the active link. That message will be flooded in the B-VPLS domain using the MC-LAG SAP B-MAC as the source MAC address. The remote nodes will learn the customer MAC addresses behind the MC-LAG and will link them to this new SAP B-MAC. MC-LAG will keep track of the active link for each particular LAG associated with a SAP B-MAC. Should MC-LAG detect any new active link in a node, a new CFM notification message will be flooded from the new active node.

The following restrictions and considerations must be taken into account:

- Only MC-LAG is supported as dual-home mechanism.
- This mechanism is supported for native PBB and/or MPLS-based PBB-VPLS. Although it is mostly beneficial when native PBB is used in the core, it can also help to optimize the re-learning process in a MPLS-based core in case of MC-LAG failures, in addition to the existing LDP MAC flush procedures.

The example of this configuration shows the setup being used in this configuration example. MAC-notification will be configured in MTU-4 and MTU-5 for the dual-homed CE-8.

The first step is to configure the SAP B-MAC that will be used for the MAC notification messages. The **source-bmac-lsb** (source backbone MAC least significant bits) command has been added to the MC-LAG branch so that the operator can decide the two last octets to be used in the SAP B-MAC. Those two last octets can be derived from the LACP key (if the **use-lacp-key** statement is used) or can be specifically defined.

```
[ex:configure redundancy multi-chassis peer 192.0.2.5 mc-lag]
A:admin@MTU-4# lag 1 ?

 lag

 Immutable fields      - lacp-key, system-id, system-priority
 apply-groups          - Apply a configuration group at this level
 apply-groups-exclude  - Exclude a configuration group at this level
 lacp-key              - Key based on the remote MC-LAG
 remote-lag            - Lag ID of the remote MC-LAG
 source-bmac-lsb       - MAC address value to apply to all ingress traffic
 system-id             - ID based on the remote MC-LAG
 system-priority       - Priority based on the remote MC-LAG
```

There must be a different SAP B-MAC per MC-LAG. The use of the LACP key as a default for two least significant octets makes the operations simpler. In this example, the last two octets of the SAP BMAC will come from the lacp-key. The configuration on MTU-4 is as follows:

```
# on MTU-4:
configure {
    redundancy {
        multi-chassis {
            peer 192.0.2.5 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                admin-state enable
                mc-lag {
                    admin-state enable
                    lag 1 {
                        lacp-key 15
                        system-id 00:00:00:00:00:01
                        system-priority 65535
                        source-bmac-lsb use-lacp-key
                    }
                }
            }
        }
```

Therefore, the SAP B-MAC will be formed in the following way:

[SAP BMAC = 4 first bytes of the source BMAC + 2 bytes from source-bmac-lsb]

MAC notification in B-VPLS 100 is enabled on all MTUs, as follows:

```
# on MTU-4, MTU-5, MTU-6:
configure {
    service {
        vpls "B-VPLS 100"
            pbb {
                mac-notification {
                    admin-state enable
                }
```

The **mac-notification** command activates the described mechanism and has the following parameters:

```
[ex:configure service vpls "B-VPLS 100" pbb]
A:admin@MTU-4# mac-notification ?

 mac-notification

 admin-state          - Administrative state of MAC notification
 count                - MAC notification messages count
 interval             - Interval for MAC notification messages
 renotify             - Re-notify interval for MAC-notification messages
```

Where:

- interval <value> controls how often the subsequent MAC notification messages are sent. Default = 100 ms. Required values: 100 ms – 10 sec, in increments of 100 ms.
- count <value> controls how often the MAC notification messages are sent. Default: 3. Range: 1–10.

The "count" and "interval" parameters can also be configured at the service context. The settings configured at the B-VPLS service context take precedence though.

```
[ex:configure service pbb]
A:admin@MTU-4# mac-notification ?

 mac-notification

 apply-groups          - Apply a configuration group at this level
 apply-groups-exclude  - Exclude a configuration group at this level
 count                 - MAC notification messages count
 interval              - Interval for MAC-notification messages
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Finally, the B-VPLS is instructed to use the SAP B-MAC. The **use-mclag-bmac-lsb** statement enables the use of the source B-MAC allocated to the multi-homed SAPs (assigned to the MC-LAG) in the related I-VPLS service (could be Epipe service as well). The command will fail if the value of the source B-MAC assigned to the B-VPLS is the hardware (chassis) B-MAC. In other words, the source B-MAC must be a configured one. The **use-mclag-bmac-lsb** statement is by default false.

```
# on MTU-4:
configure {
    service {
        vpls "B-VPLS 100"
            pbb {
                source-bmac {
                    address 00:aa:aa:aa:aa:04
                    use-mclag-bmac-lsb true
                }
```

```
# on MTU-5:
configure {
    service {
        vpls "B-VPLS 100"
            pbb {
                source-bmac {
                    address 00:aa:aa:aa:aa:05
                    use-mclag-bmac-lsb true
                }
```

```
[]
A:admin@MTU-6# show service id 2 fdb pbb

===============================================================================
Forwarding Database, i-Vpls Service 2
===============================================================================
MAC               Source-Identifier     B-Svc     b-Vpls MAC        Type/Age
 Transport:Tnl-Id
-------------------------------------------------------------------------------
00:08:00:00:00:00 b-sdp:63:100          100       00:aa:aa:aa:00:0f L/0
00:10:00:00:00:00 sap:1/1/1:10          100       N/A               L/0
===============================================================================
```

As soon as MAC notification is enabled, an Ethernet CFM notification message is sent from MTU-4, which is the node where the active MC-LAG link resides. The CFM message will have the source MAC address "00:aa:aa:aa:00:0f" (4 first bytes of the configured source BMAC + 2 bytes from the configured source-bmac-lsb, which is 15 in hex) and will be flooded throughout the B-VPLS domain. Should the link between CE-8 and MTU-4 fail, the MC-LAG protocol will activate the redundant link and MTU-5 will immediately issue a CFM message with the shared sourced SAP B-MAC that will be flooded in the B-VPLS domain.

PBB and IGMP snooping

IGMP snooping can be enabled on I-VPLS SAPs and SDPs (it cannot be enabled on B-VPLS). SR OS can keep track of IGMP joins received over individual B-SDPs or B-SAPs, and it starts flooding the multicast group (and only the multicast group) to all B-components (using the group B-MAC for I-SID) as soon as the first IGMP join for that multicast group is received in one of the B-SAP/SDP components.

The first IGMP join message received over the local B-VPLS will add all the B-VPLS SAP/SDP components into the related multicast table associated with the I-VPLS context. When the querier is connected to a remote I-VPLS instance, over the B-VPLS infrastructure, its location is identified by the B-

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

VPLS SDP/SAP on which the query was received and also by the source B-MAC address used in the PBB header for the query message, the B-MAC associated with the B-VPLS instance on the remote PBB PE.

The following configuration on MTU-4 enables IGMP snooping in I-VPLS 1 and adds some static groups on a SAP. The location of the querier is configured by adding the B-MAC where the querier is connected to (in this example, MTU-6) and adding the two B-VPLS spoke-SDPs as mrouter ports (B-VPLS mrouter ports are added in the I-VPLS backbone-vpls context).

The **mac** command translates MAC address into strings so that the names can be used instead of typing the entire MAC address every time we need to.

```
# on MTU-4:
configure {
    service {
        pbb {
            source-bmac {
                address 00:04:04:04:04:04
            }
            mac "MTU-4" {
                address 00:04:04:04:04:04
            }
            mac "MTU-5" {
                address 00:05:05:05:05:05
            }
            mac "MTU-6" {
                address 00:06:06:06:06:06
            }
        }
        vpls "I-VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS 100" {
                    isid 1
                    igmp-snooping {
                        mrouter-destination "MTU-6" { }
                    }
                    spoke-sdp 41:100 {
                        igmp-snooping
                            mrouter-port true
                        }
                    }
                    spoke-sdp 42:100 {
                        igmp-snooping
                            mrouter-port true
                        }
                    }
                }
            }
            igmp-snooping
                admin-state enable
            }
            sap 1/1/1:7 {
                igmp-snooping {
                    static {
                        group 228.0.0.1 {
                            starg
                        }
                        group 228.0.0.2 {
                            starg
                        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        group 239.0.0.1 {
                            source 172.16.99.99 { }
                        }
                    }
                }
            }
```

As in regular VPLS instances, mrouter ports are added to all the multicast groups:

```
[]
A:admin@MTU-4# show service id 1 mfib

===============================================================================
Multicast FIB, Service 1
===============================================================================
Source Address  Group Address         Port Id                      Svc Id   Fwd
                                                                          Blk
-------------------------------------------------------------------------------
*               *                     b-sdp:41:100                 100      Fwd
                                      b-sdp:42:100                 100      Fwd
*               228.0.0.1             sap:1/1/1:7                  Local    Fwd
                                      b-sdp:41:100                 100      Fwd
                                      b-sdp:42:100                 100      Fwd
*               228.0.0.2             sap:1/1/1:7                  Local    Fwd
                                      b-sdp:41:100                 100      Fwd
                                      b-sdp:42:100                 100      Fwd
172.16.99.99    239.0.0.1             sap:1/1/1:7                  Local    Fwd
                                      b-sdp:41:100                 100      Fwd
                                      b-sdp:42:100                 100      Fwd
-------------------------------------------------------------------------------
Number of entries: 4
===============================================================================
```

When the **show service id x mfib** command is issued in an I-VPLS as in the preceding output, the IGMP (S,G) and (*,G) entries for the I and B components are shown if IGMP snooping is enabled. However, when the same command is launched in a B-VPLS as in the following output, the group B-MAC entries are shown.

```
[]
A:admin@MTU-4# show service id 100 mfib

===============================================================================
Multicast FIB, Service 100
===============================================================================
Source Address  Group Address         Port Id                      Svc Id   Fwd
                                                                          Blk
-------------------------------------------------------------------------------
*               01:1e:83:00:00:01     b-sdp:41:100                 Local    Fwd
*               01:1e:83:00:00:02     b-sdp:41:100                 Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 2
===============================================================================
```

MMRP policies and ISID-based filtering for PBB inter-domain expansion

As described in the MMRP for flooding optimization section, MMRP is used in the backbone VPLS instances to build per I-VPLS flooding trees. Each I-VPLS has an associated group B-MAC in the B-

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

VPLS, which is derived from the ISID, and is advertised by MMRP throughout the whole B-VPLS context, regardless of whether a specific I-VPLS is present in one or all the B-VPLS PEs.

In an inter-domain environment, the same B-VPLS can be defined in different domains and therefore MMRP will advertise all the group B-MACs in every domain. The group B-MACs are consuming resources in all the PEs no matter if a particular ISID—and therefore its group B-MAC—is required in one of the domains or not. When MMRP is enabled in a particular PE, data plane and control plane resources are consumed and they must be taken into consideration when designing PBB-VPLS networks:

- Control plane – MRRP processing takes CPU cycles and the number of attributes that can be advertised is not unlimited
- Data plane – each group B-MAC registration takes one MFIB entry (the MFIB is shared between MMRP and IGMP/PIM snooping)

SR OS routers support MMRP policies and ISID-based filters so that control plane and data plane resources can be saved when I-VPLS instances are not defined in all the domains.

Figure 237: Inter-domain B-VPLS and MMRP policies/ISID-based filters example illustrates an example of usage for MMRP policies and ISID-based filters that will be configured in this section. "Domain 1" and "domain 2" will have a range of local ISIDs each and a range of "inter-domain" ISIDs:

- Domain 1 local ISIDs: from 1 to 100
- Domain 2 local ISIDs: from 101 to 200
- Inter-domain ISIDs: from 1000 to 2000

By applying the MMRP policies indicated in Figure 237: Inter-domain B-VPLS and MMRP policies/ISID-based filters example, domain 1 attributes will be prevented from being declared and registered in domain 2 and the other way around, domain 2 attributes from being declared and registered in domain 1. The egress MAC filters will drop any traffic sourced from a local ISID preventing it to be transmitted to the remote domain.

Figure 237: Inter-domain B-VPLS and MMRP policies/ISID-based filters example

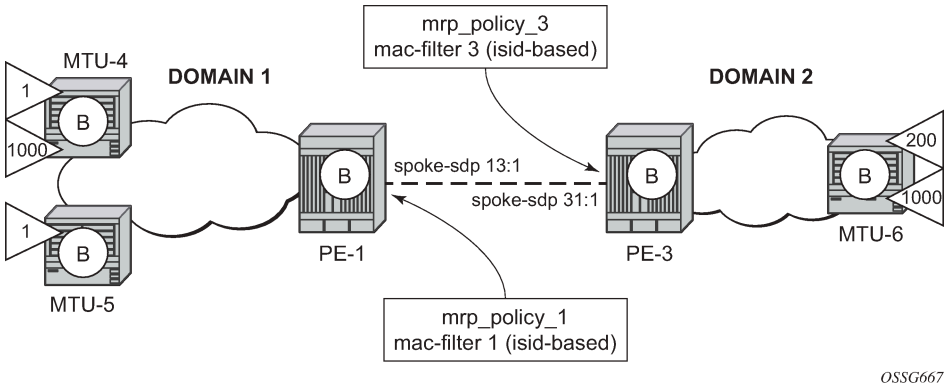

MMRP policies

The following shows the MMRP policy configuration on node PE-1. This policy will block any registration/declaration except those for ISIDs 1000-2000. Packets will be compared against the configured matching ISIDs as long as the PBB Etype matches the one configured on the port or SDP.

```
# on PE-1:
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
configure {
    service {
        mrp {
            policy "mrp_policy_1" {
                description "allow-inter-domain-isids"
                default-action block
                entry 10 {
                    action allow
                    match {
                        isid 1000 {
                            higher-value 2000
                        }
                    }
                }
            }
```

After the MMRP policy is configured, it must be applied on the corresponding SAP or SDP-binding. An MRP policy can be applied to a B-VPLS SAP, B-VPLS spoke-SDP or B-VPLS mesh-SDP:

```
# on PE-1:
configure {
    service {
        vpls "B-VPLS 100" {
            spoke-sdp 14:100 {
                mrp {
                    policy "mrp_policy_1"
                }
            }
            spoke-sdp 15:100 {
                mrp {
                    policy "mrp_policy_1"
                }
            }
```

In the same way, mrp_policy_3 will be configured in PE-3.

Some additional considerations about the MMRP policies:

- Different entries within the same MRP policy can have overlapping ISID ranges. The entries will be evaluated in the order of their IDs and the first match will cause the implementation to execute the associated action for that entry and then to exit the MRP policy.
- If no ISID is specified in the match condition then:
  - If the action is "end-station", no entry is added and the action is block.
  - If the action is different from "end-station", every ISID is considered for that action.
- The MRP policy specifies either a forward or a drop action for the group B-MAC attributes associated with the ISIDs specified in the match criteria.

```
[ex:configure service mrp policy "mrp_policy_1" entry 10]
A:admin@PE-1# action ?

 action <keyword>
 <keyword>  - (block|allow|end-station)

    Specify the action to take for packets that match this mrp-policy entry
```

- There is an additional action called end-station. This action specifies that an end-station emulation is present on the SAP/SDP-binding where the policy has been applied. The matching ISIDs will not get declared/registered in the SAP/SDP-binding (just like the block action). However, those attributes will

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

get mapped as static MMRP entries on the SAP/SDP-binding, which implicitly get instantiated in the data plane as MFIB entries associated with that SAP/SDP-binding for the related group B-MAC. When the action is "end-station", the default action must be block:

```
*[ex:configure service mrp policy "mrp_policy_3"]
A:admin@PE-3# default-action allow

*[ex:configure service mrp policy "mrp_policy_3"]
A:admin@PE-3# commit
MINOR: MGMT_CORE #4001: configure service mrp policy "mrp_policy_3" entry 10 action -
Mrp-policy default-action must be block when end-station action exists - configure
 service mrp policy "mrp_policy_3" default-action
```

- The end-station action can be used in the inter-domain gateways when, for instance, we do not want MMRP control plane exchanges between domains. The following output shows how to define the static MMRP entries 1000-2000 in PE-3 without receiving any declaration for any of those attributes or having any of those locally configured.

```
# on PE-3:
configure {
    service {
        mrp {
             policy "mrp_policy_3" {
                default-action block
                entry 10 {
                    action end-station
                    match {
                        isid 1000 {
                            higher-value 2000
                        }
                    }
                }
            }
```

```
[]
A:admin@PE-3# show service id 100 mfib

===============================================================================
Multicast FIB, Service 100
===============================================================================
Source Address  Group Address         Port Id                     Svc Id   Fwd
                                                                          Blk
-------------------------------------------------------------------------------
*               01:1e:83:00:03:e8     b-sdp:36:100                Local    Fwd
*               01:1e:83:00:03:e9     b-sdp:31:4294967294         Local    Fwd
                                      b-sdp:36:100                Local    Fwd
---snip---
*               01:1e:83:00:07:ce     b-sdp:36:100                Local    Fwd
*               01:1e:83:00:07:cf     b-sdp:36:100                Local    Fwd
*               01:1e:83:00:07:d0     b-sdp:36:100                Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 1001
===============================================================================
```

- The MRP policy can be applied to multiple B-VPLS services as long as the scope of the policy is template (the scope can also be exclusive).
- Any changes made to the existing policy will be applied immediately to all services where this policy is applied. For this reason, when many changes are required on a MRP policy, Nokia recommends copying the policy to a work-in-progress policy. That work-in-progress policy can be modified until

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

complete and then written over the original MRP policy. You can use the **configure service mrp copy** command to work with the policies in this manner.

```
[ex:configure service mrp]
A:admin@PE-1# copy policy "mrp_policy_3" to policy ?

 [policy-name] <string>
 <string>  - <1..32 characters>

    Specify the policy name associated with the MRP
```

The **rename** command can help to change the entries sequence order.

```
[ex:configure service mrp policy "mrp_policy_3"]
A:admin@PE-3# rename entry 10 to ?

 [entry-id] <number>
 <number>  - <1..65535>

    Sepcify an id for the MRP policy entry
```

An MRP policy cannot be deleted until it is removed from all the SAPs/SDP-bindings where it is applied.

ISID-based filters

The MMRP policies help to control the exchange of group B-MAC attributes across domains. Based on the registration state of a specific group B-MAC on a SAP/SDP-binding, the BUM traffic for a particular I-VPLS will be allowed or dropped. However, to avoid that any local ISID packet is flooded to the remote B-VPLS domain, all the packets tagged with the local ISIDs at the gateway PEs need to be filtered at the data plane. ISID-based filters will prevent the local ISIDs from sending any packet with unicast B-MAC to the remote domain. This is particularly useful for PBB-Epipe services across domains, where all the frames use unicast B-MACs and MMRP policies cannot help because they only act on group B-MAC packets.

The following CLI output shows how to configure an ISID-based filter that drops all the traffic sourced from the local ISIDs on PE-1 (the default action is drop and it does not show up in the configuration).

```
# on PE-1:
configure {
    filter {
        mac-filter "MAC 1" {
            description "drop_local_isids"
            type isid
            filter-id 1
            entry 10 {
                log 101
                match {
                    frame-type 802dot3
                    isid {
                        range {
                            start 1000
                            end 2000
                        }
                    }
                }
                action {
                    accept
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Once the filter is configured, it must be applied on a B-VPLS SAP or SDP-binding and always at egress.

```
# on PE-1:
configure {
    service {
        vpls "B-VPLS 100" {
            spoke-sdp 14:100 {
                egress {
                    filter {
                        mac "MAC 1"
                    }
                }
            }
            spoke-sdp 15:100 {
                egress {
                    filter {
                        mac "MAC 1"
                    }
                }
            }
```

Some additional comments about ISID-based filters:

- The **type isid** statement must be added when ISIDs are defined in the match command, otherwise the system will show an error, as follows:

```
*[ex:configure filter mac-filter "MAC 2"]
A:admin@PE-1# commit
MINOR: MGMT_CORE #4001: configure filter mac-filter "MAC 2" entry 10 match isid value
 - The match criteria entered are not compatible with the Mac filter type - Allowed
only with mac-filter type ISID - configure filter mac-filter "MAC 2" type
```

- When the operator sets the "type isid", the filter cannot be applied at ingress. Only egress ISID-based filters are allowed:

```
[ex:configure service vpls "B-VPLS 100" spoke-sdp 14:100 ingress filter]
A:admin@PE-1# mac "MAC 1"

*[ex:configure service vpls "B-VPLS 100" spoke-sdp 14:100 ingress filter]
A:admin@PE-1# commit
MINOR: SVCMGR #2050: configure service vpls "B-VPLS 100" spoke-sdp 14:100 ingress
 filter mac - Can not apply filter of type 'isid' on ingress - configure filter
 mac-filter "MAC 1" type
```

- Like any filter or MMRP policy, the filter can be applied to multiple B-VPLS services as long as the scope of the policy is "template" (the scope can also be "exclusive").
- The following command shows the filter configuration and packets that have matched the filter (field "Egr. Matches"):

```
[]
A:admin@PE-1# show filter mac 1

===============================================================================
Mac Filter
===============================================================================
Filter Id           : 1                             Applied        : Yes
Scope               : Template                      Def. Action    : Drop
Entries             : 1                             Type           : isid
Description         : drop_local_isids
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Filter Name         : MAC 1
-------------------------------------------------------------------------------
Filter Match Criteria : Mac
-------------------------------------------------------------------------------
Entry               : 10                              FrameType      : Ethernet
Description         : (Not Specified)
Log Id              : 101
ISID                : 1000..2000
Primary Action      : Forward
Ing. Matches        : 0 pkts
Egr. Matches        : 5 pkts (580 bytes)

===============================================================================
```

- Like any other filter, the matching packets can be logged. An example follows (the Ethertype is 0x88e7, which is the default standard Ethertype for PBB):

```
[]
A:admin@PE-1# show filter log 101

===============================================================================
Filter Log
===============================================================================
Admin state : Enabled
Description : Default filter log
Destination : Memory
Wrap        : Enabled
-------------------------------------------------------------------------------
Maximum entries configured : 1000
Number of entries logged   : 5
-------------------------------------------------------------------------------
2021/01/12 17:13:40  Mac Filter: 1:10  Desc:
Interface: int-PE-1-MTU-4  Direction: Egress  Action: Forward
VID match: 0
Src MAC: 00-06-06-06-06-06  Dst MAC: 00-aa-aa-aa-00-0f  EtherType: 88e7
Hex: 00 00 03 e9 00 08 00 00 00 00 00 10 00 00 00 00
     08 00 45 00 00 54 27 97 00 00 40 01 22 ee ac 10
     6c 02 ac 10 6c 01 00 00 f1 ff 00 fb 80 01 5f fd*

2021/01/12 17:13:41  Mac Filter: 1:10  Desc:
Interface: int-PE-1-MTU-4  Direction: Egress  Action: Forward
VID match: 0
Src MAC: 00-06-06-06-06-06  Dst MAC: 00-aa-aa-aa-00-0f  EtherType: 88e7
Hex: 00 00 03 e9 00 08 00 00 00 00 00 10 00 00 00 00
     08 00 45 00 00 54 27 99 00 00 40 01 22 ec ac 10
     6c 02 ac 10 6c 01 00 00 41 05 00 fb 80 02 5f fd*

---snip---
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

B-VPLS and I-VPLS show and debug commands

For the following output, the MRP policies and ISID-based MAC filters have been removed from the spoke-SDPs on PE-1 and PE-3. The following commands can help to check the B-VPLS and I-VPLS configuration and their related parameters. The first is for the B-VPLS on MTU-4:

```
[]
A:admin@MTU-4# show service id 100 base
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Service Basic Information
===============================================================================
Service Id        : 100                Vpn Id            : 0
Service Type      : b-VPLS
MACSec enabled    : no
Name              : B-VPLS 100
Description       : (Not Specified)
Customer Id       : 1                  Creation Origin   : manual
Last Status Change: 01/12/2021 16:08:29
Last Mgmt Change  : 01/12/2021 17:03:38
Etree Mode        : Disabled
Admin State       : Up                 Oper State        : Up
MTU               : 2000
SAP Count         : 0                  SDP Bind Count    : 2
Snd Flush on Fail : Disabled           Host Conn Verify  : Disabled
SHCV pol IPv4     : None
Propagate MacFlush: Disabled           Per Svc Hashing   : Disabled
Allow IP Intf Bind: Disabled
Fwd-IPv4-Mcast-To*: Disabled           Fwd-IPv6-Mcast-To*: Disabled
Mcast IPv6 scope  : mac-based
Temp Flood Time   : Disabled           Temp Flood        : Inactive
Temp Flood Chg Cnt: 0
SPI load-balance  : Disabled
TEID load-balance : Disabled
Src Tep IP        : N/A
Vxlan ECMP        : Disabled
MPLS ECMP         : Disabled
VSD Domain        : <none>
Oper Backbone Src : 00:aa:aa:aa:aa:04
Use SAP B-MAC     : Enabled
i-Vpls Count      : 2
Epipe Count       : 0
Use ESI B-MAC     : Disabled

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                              Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sdp:41:100 S(192.0.2.1)                 Spok         8000    8000    Up   Up
sdp:42:100 S(192.0.2.2)                 Spok         8000    8000    Up   Up
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

For the I-VPLS on MTU-4:

```
[]
A:admin@MTU-4# show service id 1 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 1                  Vpn Id            : 0
Service Type      : i-VPLS
MACSec enabled    : no
Name              : I-VPLS 1
Description       : (Not Specified)
Customer Id       : 1                  Creation Origin   : manual
Last Status Change: 01/12/2021 16:17:52
Last Mgmt Change  : 01/12/2021 16:17:52
Etree Mode        : Disabled
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Admin State       : Up                  Oper State        : Up
MTU               : 1514
SAP Count         : 1                   SDP Bind Count    : 0
Snd Flush on Fail : Disabled            Host Conn Verify  : Disabled
SHCV pol IPv4     : None
Propagate MacFlush: Disabled            Per Svc Hashing   : Disabled
Allow IP Intf Bind: Disabled
Fwd-IPv4-Mcast-To*: Disabled            Fwd-IPv6-Mcast-To*: Disabled
Mcast IPv6 scope  : mac-based
Temp Flood Time   : Disabled            Temp Flood        : Inactive
Temp Flood Chg Cnt: 0
SPI load-balance  : Disabled
TEID load-balance : Disabled
Src Tep IP        : N/A
Vxlan ECMP        : Disabled
MPLS ECMP         : Disabled
VSD Domain        : <none>
b-Vpls Id         : 100                 Oper ISID         : 1
b-Vpls Status     : Up
Snd Flush in bVpls: None
Flsh On bVpls Fail: Disabled            Prop Flsh fr bVpls: Disabled
Force QTag Fwd    : Disabled
SendBvplsEvpnFlush: Disabled

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                               Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/1:7                              q-tag        1518    1518    Up   Up
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

The following command shows all the I-VPLS instances multiplexed into a particular B-VPLS.

```
[]
A:admin@MTU-4# show service id 100 i-vpls

===============================================================================
Related i-Vpls services for b-Vpls service 100
===============================================================================
i-Vpls SvcId        Oper ISID           Admin               Oper
-------------------------------------------------------------------------------
1                   1                   Up                  Up
2                   2                   Up                  Up
-------------------------------------------------------------------------------
Number of Entries : 2
-------------------------------------------------------------------------------
===============================================================================
```

Some useful commands to check the I and B VPLS FIBs correlating C-MACs and B-MACs:

```
[]
A:admin@MTU-4# show service id 1 fdb pbb

===============================================================================
Forwarding Database, i-Vpls Service 1
===============================================================================
MAC               Source-Identifier     B-Svc     b-Vpls MAC        Type/Age
 Transport:Tnl-Id
-------------------------------------------------------------------------------
00:07:00:00:00:00 sap:1/1/1:7           100       N/A               L/0
00:09:00:00:00:00 b-sdp:41:100          100       00:06:06:06:06:06 L/0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
```

```
[]
A:admin@MTU-4# show service id 100 fdb pbb

===============================================================================
Forwarding Database, b-Vpls Service 100
===============================================================================
MAC               Source-Identifier     iVplsMACs  Epipes     Type/Age
 Transport:Tnl-Id
-------------------------------------------------------------------------------
00:06:06:06:06:06 sdp:41:100            2          0          L/0
02:0f:ff:00:00:00 sdp:41:100            0          0          L/0
===============================================================================
```

If MAC names are used in the configuration, the following commands can show the translations:

```
[]
A:admin@MTU-4# show service pbb mac-name

===============================================================================
MAC Name Table
===============================================================================
MAC-Name                          MAC-Address
-------------------------------------------------------------------------------
MTU-4                             00:04:04:04:04:04
MTU-5                             00:05:05:05:05:05
MTU-6                             00:06:06:06:06:06
===============================================================================
```

```
[]
A:admin@MTU-4# show service pbb mac-name mac-name-1 "MTU-6" detail

===============================================================================
Services Using MAC name='MTU-6' addr='00:06:06:06:06:06'
===============================================================================
Svc-Id                            ISID
-------------------------------------------------------------------------------
1                                 N/A
-------------------------------------------------------------------------------
Number of services: 1
===============================================================================
```

The following command shows the base MAC notification parameters as well as the source B-MAC configured at the service PBB level. Those values are overridden by any potential MAC notification or source B-MAC values configured under the B-VPLS service context.

```
[]
A:admin@MTU-4# show service pbb base

===============================================================================
PBB MAC Information
===============================================================================
MAC-Notif Count                  : 3
MAC-Notif Interval               : 1
Source BMAC                      : 00:04:04:04:04:04
Leaf Source BMAC                 : Default
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

If MAC notification is used in a particular B-VPLS, the configured least significant bits for the SAP B-MAC on a particular MC-LAG can be shown by using the detailed view of the **show lag** command:

```
[]
A:admin@MTU-4# show lag 1 detail

===============================================================================
LAG Details
===============================================================================
Description        : N/A
-------------------------------------------------------------------------------
Details
-------------------------------------------------------------------------------
Lag-id              : 1                     Mode                : access
Adm                 : up                    Opr                 : up

---snip---

MC Peer Address     : 192.0.2.5             MC Peer Lag-id      : 1
MC System Id        : 00:00:00:00:00:01     MC System Priority  : 65535
MC Admin Key        : 15                    MC Active/Standby   : active
MC Lacp ID in use   : true                  MC extended timeout : false
MC Selection Logic  : local master decided
MC Config Mismatch  : no mismatch
Source BMAC LSB     : use-lacp-key          Oper Src BMAC LSB   : 00:0f

---snip---
===============================================================================
```

The following debug commands (in classic CLI) allow the operator to check the LDP label mapping, label withdrawal, messages and also the MAC-flush messages for regular VPLS, for I-VPLS and B-VPLS including the PBB extensions and TLVs.

```
debug
    router "Base"
        ldp
            peer 192.0.2.1
                event
                exit
                packet
                    init detail
                    label detail
                exit
            exit
            peer 192.0.2.2
                event
                exit
                packet
                    init detail
                    label detail
                exit
            exit
        exit
    exit
exit
```

The following debug commands (in classic CLI) can help the operator to troubleshoot MMRP.

```
A:MTU-4# debug service id 100 mrp ?
  - mrp
  - no mrp
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
     all-events      - Enable/disable MRP debugging for all events
[no] applicant-sm    - Enable/disable MRP debugging for applicant state machine
                       changes
[no] leave-all-sm    - Enable/disable MRP debugging for leave all state machine
                       changes
[no] mmrp-mac        - Enable/disable MRP debugging for a particular MAC address
[no] mrpdu           - Enable/disable MRP debugging for Rx/Tx MRP PDUs
[no] mvrp-vlan       - Enable/disable debugging for a particular vlan
[no] periodic-sm     - Enable/disable MRP debugging for periodic state  machine
                       changes
[no] registrant-sm   - Enable/disable MRP debugging for registrant state machine
                       changes
[no] sap             - Enable/disable MRP debugging for a particular SAP
[no] sdp             - Enable/disable MRP debugging for a particular SDP
```

Conclusion

PBB-VPLS allows the service providers to scale VPLS services by multiplexing customer I-VPLS instances into one or more B-VPLS instances. This multiplexing dramatically reduces the number of services, pseudowires, and MAC addresses in the core and therefore allows the service provider to scale Layer 2 multi-point networks and provide services across international backbones.

The example used in this chapter shows the configuration of the customer and backbone VPLS instances as well as all the related features which are required for this environment. Show and debug commands have also been suggested so that the operator can verify and troubleshoot the service.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

PIM Snooping for IPv4 in EVPN-MPLS Services

This chapter provides information about PIM snooping for IPv4 in EVPN-MPLS services.

Topics in this chapter include:

Applicability

This chapter was initially written based on SR OS Release 15.0.R5, but the CLI in the current edition corresponds to SR OS Release 23.7.R1. PIM snooping for IPv4 is supported in EVPN-MPLS services in SR OS Release 15.0.R1, and later. PIM snooping in single-active multi-homing mode without ESI label is supported in SR OS Release 15.0.R1, and later, whereas PIM snooping in single-active multi-homing mode with ESI label is supported in SR OS Release 15.0.R4, and later. PIM snooping in all-active multi-homing mode is supported in SR OS Release 15.0.R4, and later. Data-driven PIM state synchronization is supported in SR OS Release 15.0.R4, and later.

Overview

Figure 238: Multicast in VPLS without PIM Snooping shows the example topology with four CEs that have IGMP and PIM enabled (L3) and four PEs configured with VPLS 1 (L2). Source-specific multicast is used in this example. The following description applies to all VPLSs, with or without EVPN.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 238: Multicast in VPLS without PIM Snooping

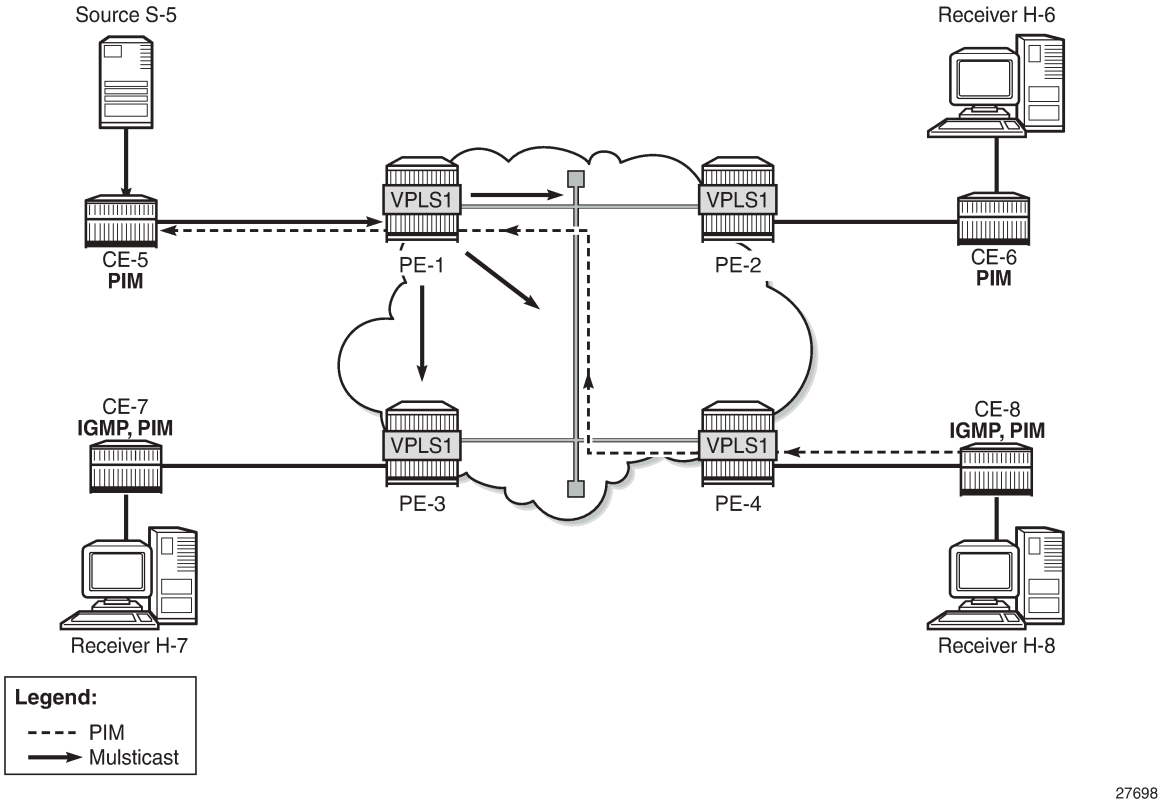

The VPLS emulates a LAN interconnecting sites with L3-capable devices that use PIM to join or leave multicast groups. When receiver H-8 sends an IGMP report message to join a multicast group, CE-8 sends a PIM join message to CE-5. The PEs forward the PIM message without learning any PIM-related information, such as which CE sent the PIM join and for which multicast group.

The source S-5 is sending the multicast stream to CE-5. When CE-5 receives a PIM join message for this multicast group from CE-8, it forwards the multicast stream to CE-8. By default, all PEs flood the multicast stream on all their connections in the VPLS domain, regardless of whether a PIM join was received from that connection. L2 flooding is not aware of the PIM join/prune messages from the L3 edge routers, resulting in an inefficient use of network resources. To avoid this L2 flooding, PIM snooping can be enabled in the VPLS by the following command:

```
configure {
    service {
        vpls "VPLS 1" {
            pim-snooping ?

 pim-snooping

 apply-groups          - Apply a configuration group at this level
 apply-groups-exclude  - Exclude a configuration group at this level
 group-policy          - Group policy name
 hold-time             - Duration that allows the PIM-snooping switch to snoop all the PIM
 states in the VPLS
 ipv4                  + Enter the ipv4 context
 ipv6                  + Enter the ipv6 context
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
configure {
    service {
        vpls "VPLS 1" {
            pim-snooping {
            ipv4 ?

 ipv4

 admin-state           - Administrative state of snooping for multicast traffic
 apply-groups          - Apply a configuration group at this level
 apply-groups-exclude  - Exclude a configuration group at this level
```

The default mode is proxy, but PIM snooping can also use snooping mode, depending on the information in the received PIM hello messages. In snooping mode, the PE does not modify the PIM messages; in proxy mode, the PE terminates incoming PIM messages and generates its own PIM messages.

PIM snooping is used for router multicast registration, whereas IGMP snooping is used for host/client multicast registration. IGMP snooping in EVPN-MPLS services is described in chapter P2MP mLDP Tunnels for BUM Traffic in EVPN-MPLS Services. Optionally, PIM snooping and IGMP snooping can be enabled simultaneously.

With PIM enabled, the CEs send PIM hello messages to the well-known multicast address for PIM, 224.0.0.13. PIM hello messages are used to form PIM neighbors and can be used to form the Forwarding Database (FDB). With PIM snooping enabled in the VPLS in the PEs, the PEs snoop PIM messages. The PEs only forward multicast traffic downstream when required, as determined from the received PIM messages. This provides a more efficient use of network resources.

PIM snooping states in a PE are maintained per VPLS instance. When PIM snooping is enabled, IP multicast traffic to a multicast group that is not learned via snooping is dropped by default, unless it is received from a directly connected source.

PIM Snooping in Snooping Mode

Figure 239: Multicast in VPLS with PIM Snooping in Snooping Mode shows that the multicast stream is not flooded in PE-1 when PIM snooping is enabled and operating in snooping mode.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 239: Multicast in VPLS with PIM Snooping in Snooping Mode

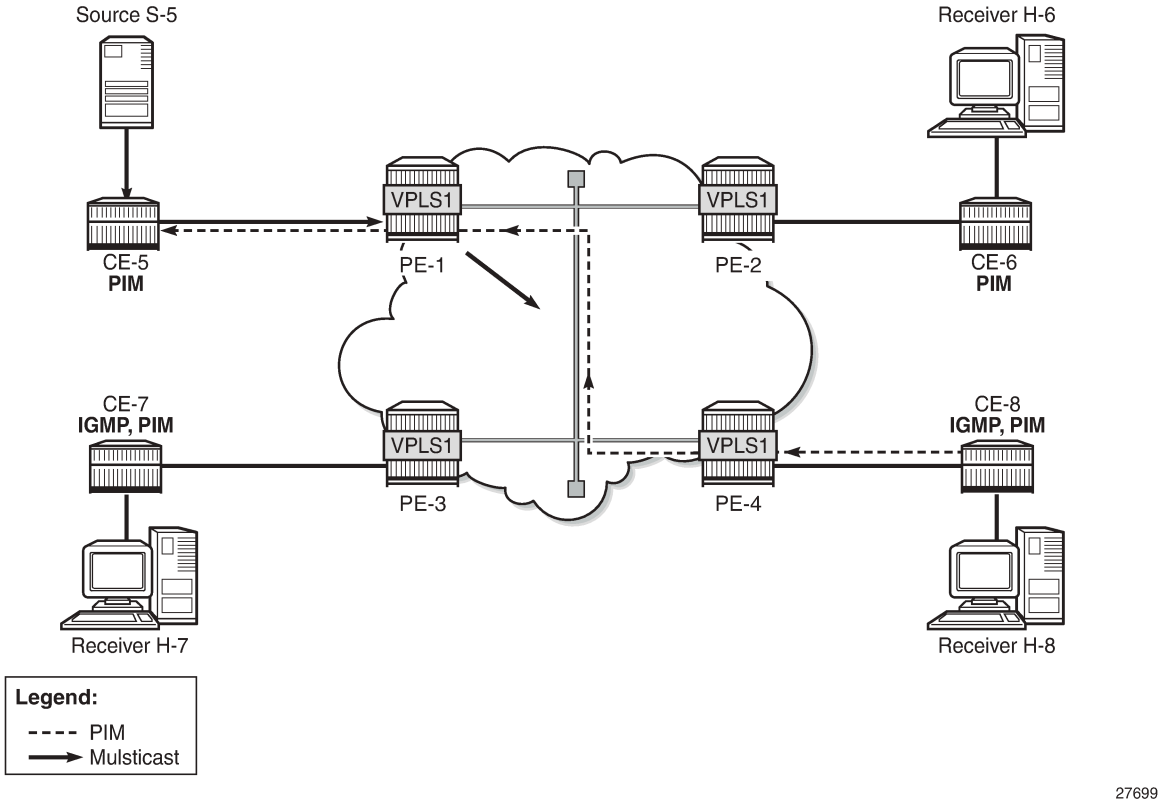

When H-8 sends an IGMP report message to join the multicast stream from source S-5 to CE-8, CE-8 sends a PIM join message to CE-5. PE-4 snoops the PIM join message and builds the FDB. PE-4 forwards the PIM join message to PE-1 by matching the upstream neighbor address in the join with the neighbor database. PE-1 snoops the PIM join message, builds its Multicast Forwarding Information Base (MFIB), and performs a similar lookup in its FDB. PE-1 forwards the PIM join to CE-5. The Source Path Tree (SPT) between receiver CE-8 and sender CE-5 is now built and CE-5 forwards multicast data frames to CE-8. PE-1 does not flood multicast frames, but forwards them to CE-8 only, based on the MFIB.

Figure 240: Multicast in VPLS with PIM Snooping in Snoop Mode – Multiple CEs shows how the number of PIM messages in the control plane increases when multiple client CEs are connected to PE-4.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 240: Multicast in VPLS with PIM Snooping in Snoop Mode – Multiple CEs

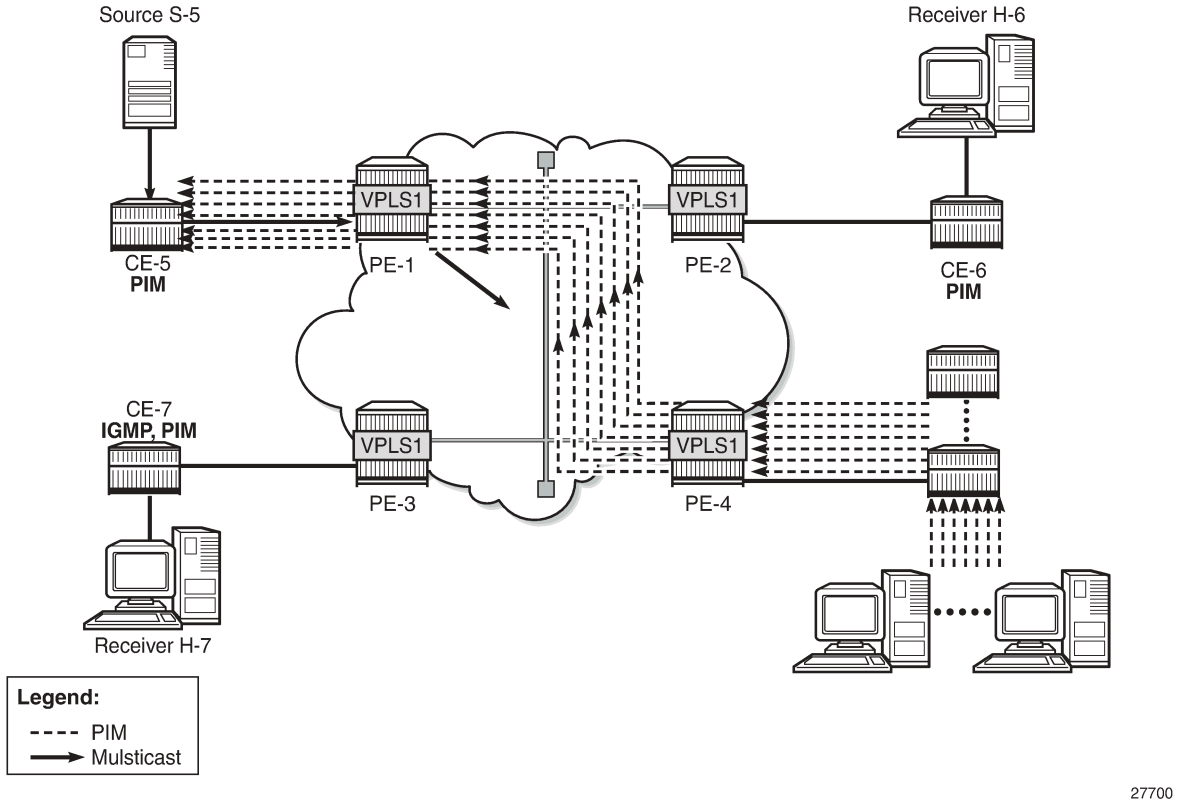

PIM Snooping in Proxy Mode

When H-8 sends an IGMP report message to join a multicast stream, CE-8 again sends a PIM join message to CE-5. PE-4 terminates the incoming PIM join message and generates its own PIM join message using CE-5 as the source address, learned from the PIM hello messages. PE-4 builds its MFIB and sends a new PIM join message to S-5. PE-1 terminates the incoming PIM join message and builds its MFIB. PE-1 generates its own PIM join message using CE-5 as the source address. PE-1 forwards the PIM join to CE-5. The SPT between CE-8 and CE-5 is now built and the multicast stream flows from source S-5 to receiver H-8. No multicast traffic is sent to CE-6 and CE-7, because they do not have receivers attached that joined the multicast stream.

The default mode for PIM snooping is proxy mode.

Figure 241: Multicast in VPLS with PIM Snooping in Proxy Mode - Multiple CEs shows that the number of PIM messages in the control plane does not increase when multiple client CEs are connected to PE-4, compared to snooping mode.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 241: Multicast in VPLS with PIM Snooping in Proxy Mode - Multiple CEs

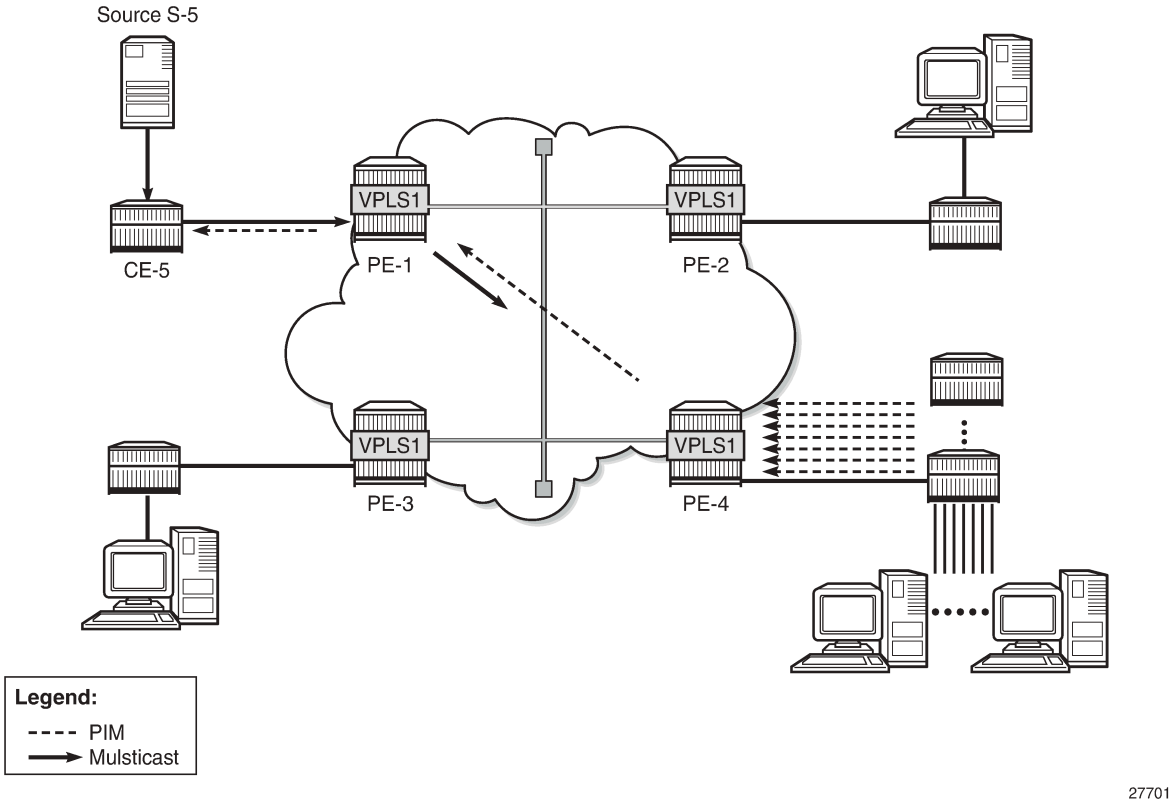

PIM snooping in proxy mode can be configured with a delay to avoid existing traffic interruption. PIM snooping in proxy mode does not program the MFIB until a hold timer has expired. This hold time is useful in the following cases:

- PIM snooping being enabled on the VPLS
- PIM snooping states being manually cleared by an operator

When the hold timer is started, but not expired yet, multicast traffic is flooded in the VPLS as if PIM snooping was not enabled. VPLS flooding ensures flow delivery during the hold time.

PIM Snooping in VPLS with EVPN-MPLS

PIM snooping in an EVPN-MPLS service supports the following:

- Regular PIM snooping on SAPs/SDP-bindings
  - PIM messages received on EVPN-MPLS endpoints are forwarded to SAPs/SDP-bindings.
  - IP multicast traffic received on an EVPN-MPLS binding is forwarded to SAPs/SDP-bindings from which a PIM join was received, or to ports configured as mrouter ports.
- The EVPN-MPLS endpoints are treated as a single PIM interface:
  - IP multicast traffic and PIM messages received on an EVPN-MPLS endpoint are not forwarded to other EVPN-MPLS endpoints (split-horizon).

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

  - Hello and join/prune messages from SAPs/SDP-bindings are forwarded to all EVPN-MPLS destinations.
  - When a hello message is received from one PIM neighbor on an EVPN-MPLS destination, the single interface representing all EVPN-MPLS destinations has that neighbor.
    - Individual destinations appear in the MFIB, but the information for each EVPN-MPLS destination entry is identical.
- If a Point-to-Multipoint (P2MP) mLDP provider tunnel is configured:
  - If the PE is the root node of a P2MP LSP that is up, PIM messages and IP multicast traffic are only forwarded over the P2MP Label Switched Path (LSP) instead of being sent to the EVPN-MPLS endpoints. Therefore, the P2MP leaves must match the EVPN-MPLS endpoints, in this case, PE-2, PE-3, and PE-4.
  - If the PE is a leaf node of a P2MP LSP, it sends PIM messages and IP multicast traffic over its EVPN-MPLS endpoints.
  - The PEs can expect to receive IP multicast traffic and PIM messages from leaf nodes over their EVPN-MPLS endpoints, or over the P2MP LSPs for traffic from root nodes.
- PIM snooping is supported in inter-AS model B and inter-AS model C, as for IGMP snooping.
- All-active and single-active EVPN multi-homing are supported.
- Multi-chassis Synchronization (MCS) of PIM snooping state is supported on SAPs and spoke-SDPs in dual-homing.
  - The active (Designated Forwarder (DF)) PE sends the PIM states to the backup non-DF (NDF) PE.
  - In case of failure, the backup PE has the PIM states already, and the multicast traffic path can be re-established fast without any need to wait for PIM states to be snooped.
  - A sync-tag is configured on the ports or SDPs that need to be synchronized on both PEs.
  - MCS PIM snooping is restricted to two peers, even though MCS supports more peers for other types of information. An error is raised when attempting to configure a sync-tag on the same port or SDP to more than one peer.
- PIM snooping is supported for both IPv4 and IPv6 multicast. PIM snooping for IPv6 uses MAC-based forwarding by default, and can be configured to use (S,G)-based forwarding.
- PIM snooping is transparent to the underlying tunnel. PIM snooping works with RSVP, LDP, SR-ISIS, SR-OSPF, SR-TE, BGP, and MPLSoUDP.
- PIM snooping is not supported with routed VPLS with EVPN-MPLS, and its configuration is blocked.

Configuration

Figure 242: Example Topology shows the example topology. Source S-5 sends multicast streams to CE-5, which forwards those only after a PIM join message has been received. An mLDP P2MP LSP is used to distribute the multicast from the root node PE-1 to the other PEs. All CEs have PIM enabled and the receiving CEs (CE-6, CE-7, and CE-8) have IGMP configured on the interface toward the receivers (H-6, H-7, and H-8). EVPN-MPLS VPLS 1 is configured on the PEs. Initially, PIM snooping is disabled in the VPLS. Receiver H-8 joins multicast group 232.1.1.1 from source S-5.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 242: Example Topology

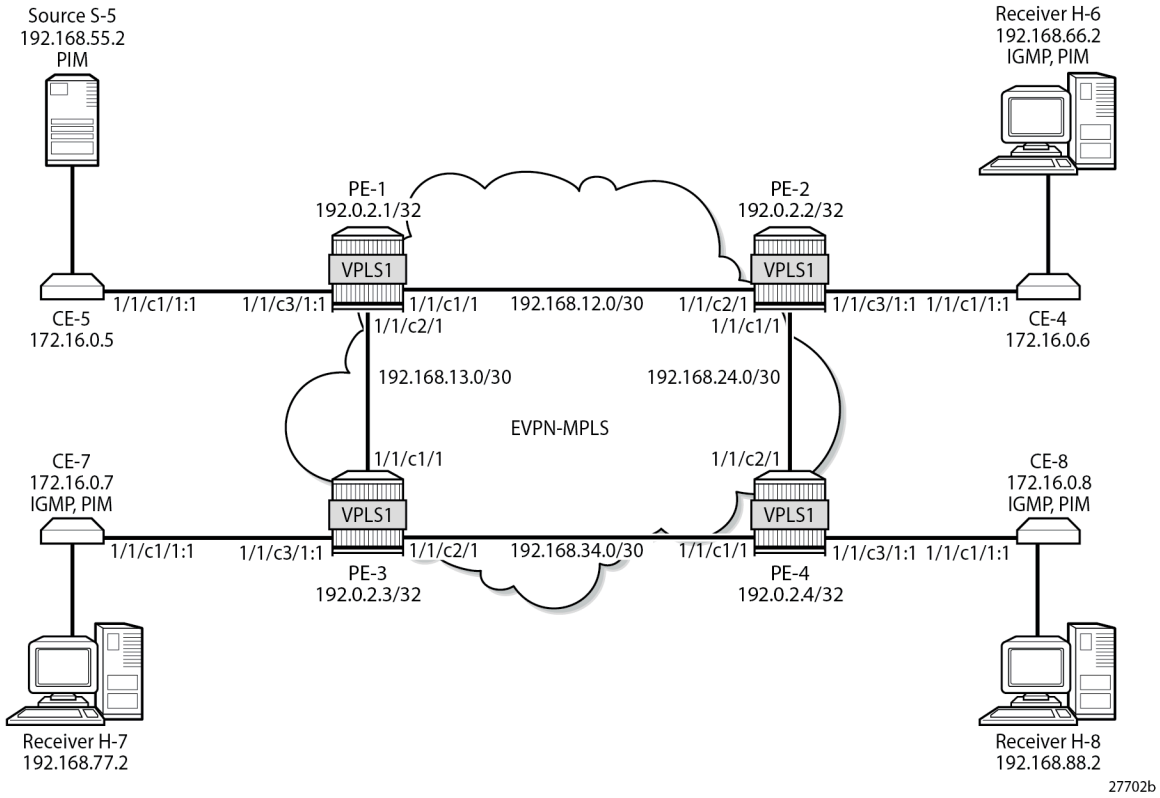

The initial configuration includes the following:

- Cards, MDAs
- Ports
  - Ports between PEs are network ports with null encapsulation
  - Ports between CEs and PEs are hybrid ports with dot1q encapsulation
- IS-IS as IGP between the PEs (alternatively, OSPF can be used)
- LDP between the PEs
- BGP with address family EVPN between the PEs. PE-2 is the route reflector (RR). The BGP configuration on RR PE-2 is as follows:

```
On PE-2:
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            rapid-withdrawal true
            ebgp-default-reject-policy {
                import false
                export false
            }
            rapid-update {
                evpn true
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
            group "INTERNAL" {
                type internal
                family {
                    evpn true
                }
                cluster {
                    cluster-id 192.0.2.2
                }
            }
            neighbor "192.0.2.1" {
                group "INTERNAL"
            }
            neighbor "192.0.2.3" {
                group "INTERNAL"
            }
            neighbor "192.0.2.4" {
                group "INTERNAL"
            }
        }
    }
}
```

EVPN-MPLS VPLS without PIM Snooping

VPLS 1 is configured with EVPN-MPLS in the PEs. By default, PIM snooping is disabled. PE-1 is configured as **root-and-leaf** node for the P2MP mLDP multicast tree, while the other three PEs have the default **no root-and-leaf** configured, so they are leaf-only nodes. The configuration of VPLS 1 on PE-1 is as follows:

```
On PE-1:
configure {
    service {
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 { }
            bgp-evpn {
                evi 1
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            sap 1/1/c3/1:1 { }
            provider-tunnel {
                inclusive {
                    admin-state enable
                    owner bgp-evpn-mpls
                    root-and-leaf true
                    mldp
                }
            }
        }
    }
}
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

A P2MP mLDP multicast tree is created from root node PE-1 to the leaf nodes. On the root node PE-1, an SDP of type **VplsPmsi** is auto-created:

```
[/]
A:admin@PE-1# show service id 1 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 1                   Vpn Id            : 0
Service Type      : VPLS
---snip---
-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                               Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/c3/1:1                           q-tag        8936    8936    Up   Up
sdp:32767:4294967294 SB(not applicable)  VplsPmsi     9782    9782    Up   Up
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

The following inclusive provider tunnel is created on root node PE-1:

```
[/]
A:admin@PE-1# show service id 1 provider-tunnel

===============================================================================
Service Provider Tunnel Information
===============================================================================
Type               : inclusive          Root and Leaf      : enabled
Admin State        : enabled            Data Delay Intvl   : 15 secs
PMSI Type          : ldp                LSP Template       :
Remain Delay Intvl : 0 secs             LSP Name used      : 8193
PMSI Owner         : bgpEvpnMpls
Oper State         : up                 Root Bind Id       : 32767
-------------------------------------------------------------------------------
Type               : selective          Wildcard SPMSI     : disabled
Admin State        : disabled           Data Delay Intvl   : 3 secs
PMSI Type          : none               Max P2MP SPMSI     : 10
PMSI Owner         : none
===============================================================================
```

When a P2MP mLDP provider tunnel is configured, the root node forwards PIM messages and IP multicast traffic over the provider tunnel instead of over the EVPN-MPLS endpoints. However, the leaf nodes of a P2MP mLDP provider tunnel send PIM messages and IP multicast traffic over the EVPN-MPLS endpoints.

The following P2MP mLDP bindings are active on root node PE-1: one toward PE-2 via port 1/1/c1/1 and one toward PE-3 via port 1/1/c2/1.

```
[/]
A:admin@PE-1# show router ldp bindings active p2mp opaque-type generic ipv4

===============================================================================
LDP Bindings (IPv4 LSR ID 192.0.2.1)
             (IPv6 LSR ID ::)
===============================================================================
Label Status:
        U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
        e - Label ELC
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
FEC Flags:
        LF - Lower FEC, UF - Upper FEC, M - Community Mismatch,
        BA - ASBR Backup FEC
===============================================================================
LDP Generic IPv4 P2MP Bindings (Active)
===============================================================================
P2MP-Id                                 Interface
RootAddr                                Op
IngLbl                                  EgrLbl
EgrNH                                   EgrIf/LspId
-------------------------------------------------------------------------------
8193                                    73728
192.0.2.1                               Push
  --                                    524281
192.168.12.2                            1/1/c1/1

8193                                    73728
192.0.2.1                               Push
  --                                    524281
192.168.13.2                            1/1/c2/1

-------------------------------------------------------------------------------
No. of Generic IPv4 P2MP Active Bindings: 2
===============================================================================
```

The following P2MP mLDP bindings are active on PE-2. PE-2 is a leaf node (pop operation) and a transit node for traffic toward PE-4 (swap operation):

```
[/]
A:admin@PE-2# show router ldp bindings active p2mp opaque-type generic ipv4

===============================================================================
---snip---
===============================================================================
LDP Generic IPv4 P2MP Bindings (Active)
===============================================================================
P2MP-Id                                 Interface
RootAddr                                Op
IngLbl                                  EgrLbl
EgrNH                                   EgrIf/LspId
-------------------------------------------------------------------------------
8193                                    73728
192.0.2.1                               Pop
524281                                    --
  --                                      --

8193                                    73728
192.0.2.1                               Swap
524281                                  524281
192.168.24.2                            1/1/c1/1

-------------------------------------------------------------------------------
No. of Generic IPv4 P2MP Active Bindings: 2
===============================================================================
```

PE-3 and PE-4 are leaf nodes, so there is a pop operation. The active P2MP LDP binding on PE-4 is the following. A similar P2MP LDP binding occurs on PE-3.

```
[/]
A:admin@PE-4# show router ldp bindings active p2mp opaque-type generic ipv4

===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
---snip---
===============================================================================
LDP Generic IPv4 P2MP Bindings (Active)
===============================================================================
P2MP-Id                                     Interface
RootAddr                                    Op
IngLbl                                      EgrLbl
EgrNH                                       EgrIf/LspId
-------------------------------------------------------------------------------
8193                                        73728
192.0.2.1                                   Pop
524281                                        --
  --                                          --

-------------------------------------------------------------------------------
No. of Generic IPv4 P2MP Active Bindings: 1
===============================================================================
```

Figure 243: P2MP mLDP Multicast Tree shows the mLDP multicast tree. Multicast traffic from source S-5 uses the mLDP multicast tree from PE-1 to both PE-2 and PE-3. PE-2 is a transit node for multicast traffic to PE-4, and also a leaf node. PE-3 and PE-4 are leaf nodes.

Figure 243: P2MP mLDP Multicast Tree

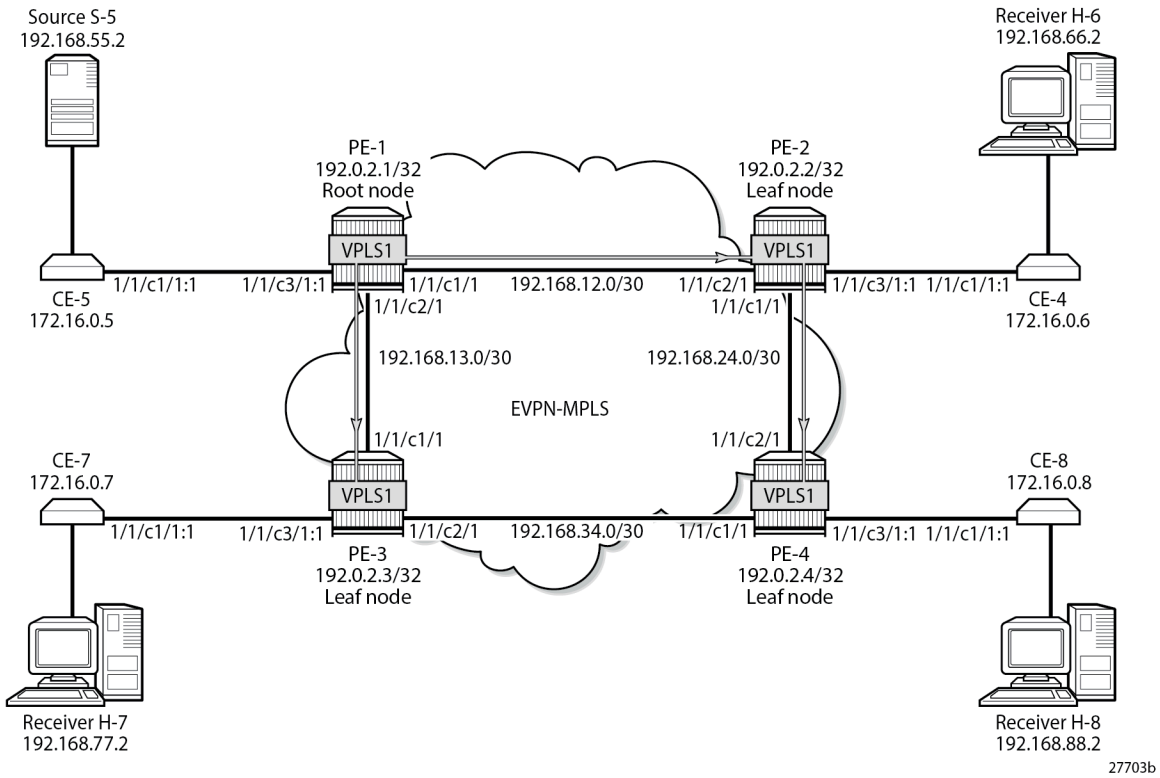

CE-6, CE-7, and CE-8 have IGMP enabled on the interface toward the receiver and PIM enabled on all interfaces. The configuration on CE-8 is as follows:

```
On CE-8:
configure {
    router "Base" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        interface "int-CE-8-H-8" {
            port 1/1/c2/1
            ipv4 {
                primary {
                    address 192.168.88.1
                    prefix-length 24
                }
            }
        }
        interface "int-CE-8-PE-4" {
            port 1/1/c1/1:1
            ipv4 {
                primary {
                    address 172.16.0.8
                    prefix-length 16
                }
            }
        }
        interface "system" {
            ipv4 {
                primary {
                    address 192.0.2.8
                    prefix-length 32
                }
            }
        }
        static-routes {
            route 192.168.55.0/30 route-type unicast {
                next-hop "172.16.0.5" {
                    admin-state enable
                }
            }
        }
        pim {
            apply-to all
        }
        igmp {
            interface "int-CE-8-H-8" { }
        }
    }
}
```

The static route is required on the receiving CEs for the PIM join/prune messages to reach the multicast source S-5 with IP address 192.168.55.2; only IP subnet 172.16.0.0/16 can be reached via the VPLS.

CE-5 has PIM enabled and static routes configured to reach the receiving hosts, as follows:

```
On CE-5:
configure {
    router "Base" {
        interface "int-CE-5-PE-1" {
            port 1/1/c1/1:1
            ipv4 {
                primary {
                    address 172.16.0.5
                    prefix-length 16
                }
            }
        }
        interface "int-CE-5-S-5" {
            port 1/1/c3/1
            ipv4 {
                primary {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    address 192.168.55.1
                    prefix-length 30
                }
            }
        }
        interface "system" {
            ipv4 {
                primary {
                    address 192.0.2.5
                    prefix-length 32
                }
            }
        }
        static-routes {
            route 192.168.66.0/24 route-type unicast {
                next-hop "172.16.0.6" {
                    admin-state enable
                }
            }
            route 192.168.77.0/24 route-type unicast {
                next-hop "172.16.0.7" {
                    admin-state enable
                }
            }
            route 192.168.88.0/24 route-type unicast {
                next-hop "172.16.0.8" {
                    admin-state enable
                }
            }
        }
        pim {
            apply-to all
        }
    }
}
```

The PIM neighbors of CE-5 are the receiving CEs: CE-6, CE-7, and CE-8, as follows:

```
[/]
A:admin@CE-5# show router pim neighbor

===============================================================================
PIM Neighbor ipv4
===============================================================================
Interface               Nbr DR Prty     Up Time       Expiry Time     Hold Time
   Nbr Address
-------------------------------------------------------------------------------
int-CE-5-PE-1           1               0d 00:00:39   0d 00:01:38     105
   172.16.0.6
int-CE-5-PE-1           1               0d 00:00:28   0d 00:01:30     105
   172.16.0.7
int-CE-5-PE-1           1               0d 00:00:14   0d 00:01:32     105
   172.16.0.8
-------------------------------------------------------------------------------
Neighbors : 3
===============================================================================
```

Figure 244: H-8 Joins Group (192.168.55.2, 232.1.1.1) and PIM Snooping is Disabled shows that receiver H-8 sends an IGMP report to CE-8 and CE-8 sends a PIM join message to CE-5 via PE-4. PE-4 floods the PIM join message to all PEs, and the message is not snooped by any intermediate PE.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 244: H-8 Joins Group (192.168.55.2, 232.1.1.1) and PIM Snooping is Disabled

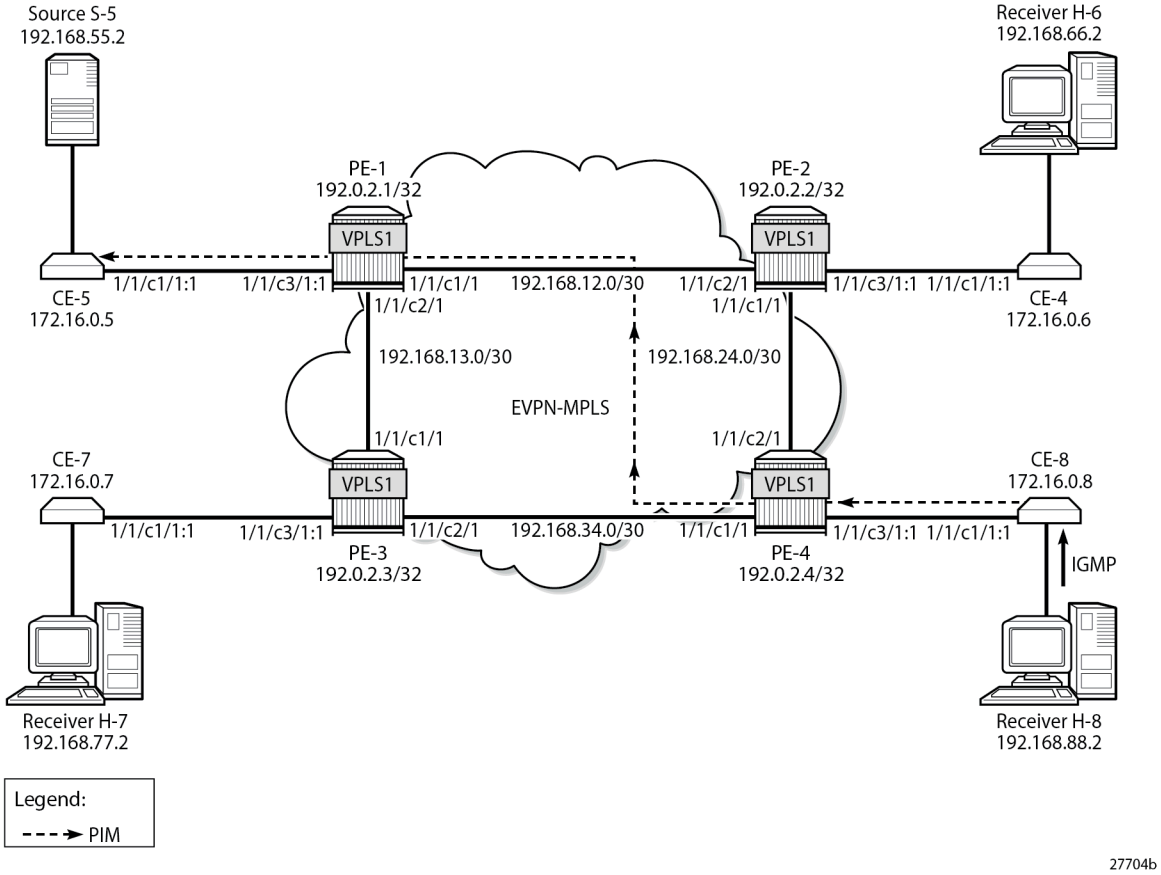

Alternatively, a static multicast group can be configured on IGMP interface int-CE-8-H-8 for multicast group (192.168.55.2, 232.1.1.1), as follows:

```
On CE-8:
configure {
    router "Base" {
        igmp {
            interface "int-CE-8-H-8" {
                ssm-translate {
                    group-range start 232.0.0.0 end 232.255.255.255 {
                        source 192.168.55.2 { }
                    }
                }
                static {
                    group 232.1.1.1 {
                        source 192.168.55.2 { }
                    }
                }
            }
        }
    }
}
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

CE-8 sends the following PIM join message for multicast group (192.168.55.2, 232.1.1.1) to upstream IP address 172.16.0.5 on CE-5:

```
1 2023/08/10 22:51:38.087 CEST MINOR: DEBUG #2001 Base PIM[Instance 1 Base]
"PIM[Instance 1 Base]: Join/Prune
[000 00:17:08.130] PIM-TX ifId 3 ifName int-CE-8-PE-4 0.0.0.0 -> 224.0.0.13 Length: 34
PIM Version: 2 Msg Type: Join/Prune Checksum: 0x4828
Upstream Nbr IP : 172.16.0.5 Resvd: 0x0, Num Groups 1, HoldTime 210
Group: 232.1.1.1/32 Num Joined Srcs: 1, Num Pruned Srcs: 0
Joined Srcs:
192.168.55.2/32 Flag S  <S,G>

"
```

Multicast stream 232.1.1.1 is sent from source S-5 to CE-5. When CE-5 has received the PIM join message, it floods the multicast stream to PE-1. Root node PE-1 sends the multicast stream to both PE-2 and PE-3. PE-2 forwards the multicast stream to PE-4 and to CE-6; PE-3 forwards the stream to CE-7, and PE-4 forwards to CE-8. The following PIM group for group address 232.1.1.1 is joined on CE-8:

```
[/]
A:admin@CE-8# show router pim group detail

===============================================================================
PIM Source Group ipv4
===============================================================================
Group Address      : 232.1.1.1
Source Address     : 192.168.55.2
RP Address         : 0
Advt Router        :
Flags              :                    Type              : (S,G)
Mode               : sparse
MRIB Next Hop      : 172.16.0.5
MRIB Src Flags     : remote
Keepalive Timer    : Not Running
Up Time            : 0d 00:02:54        Resolved By       : rtable-u

Up JP State        : Joined             Up JP Expiry      : 0d 00:00:05
Up JP Rpt          : Not Joined StarG   Up JP Rpt Override : 0d 00:00:00

Register State     : No Info
Reg From Anycast RP: No

Rpf Neighbor       : 172.16.0.5
Incoming Intf      : int-CE-8-PE-4
Outgoing Intf List : int-CE-8-H-8

Curr Fwding Rate   : 9751.560 kbps
Forwarded Packets  : 71591              Discarded Packets  : 0
Forwarded Octets   : 106097862          RPF Mismatches     : 0
Spt threshold      : 0 kbps             ECMP opt threshold : 7
Admin bandwidth    : 1 kbps
-------------------------------------------------------------------------------
Groups : 1
===============================================================================
```

CE-8 forwards the multicast stream to outgoing interface int-CE-8-H-8 toward receiver H-8, while CE-6 and CE-7 drop the traffic.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The following port statistics show that the incoming traffic on port 1/1/c3/1 on PE-1 is forwarded to port 1/1/c1/1 to PE-2 and to port 1/1/c2/1 to PE-3:

```
[/]
A:admin@PE-1# show port 1/1/c1/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c1/1                                37                           3703
                                     27933                       42354073
===============================================================================

[/]
A:admin@PE-1# show port 1/1/c2/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c2/1                                33                           3356
                                     27932                       42354034
===============================================================================

[/]
A:admin@PE-1# show port 1/1/c3/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c3/1                             27901                       41961676
                                         4                            304
===============================================================================
```

Besides the multicast traffic, signaling messages (such as IS-IS or BGP) are sent, which explains the other counters on the ports being different from zero.

A similar result occurs on PE-2, where incoming traffic from PE-1 is forwarded to PE-4 and to CE-6.

The following port statistics on CE-6 show that the incoming traffic on port 1/1/c1/1 from PE-2 is not forwarded to port 1/1/c2/1 to H-6:

```
[/]
A:admin@CE-6# show port 1/1/c1/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c1/1                             27946                       42025072
                                         1                             76
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
[/]
A:admin@CE-6# show port 1/1/c2/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                          Ingress Packets                 Ingress Octets
Id                             Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c2/1                                    0                               0
                                            2                             136
===============================================================================
```

Without PIM snooping, multicast streams are forwarded to CEs that drop them, which wastes resources. Figure 245: Multicast Stream (192.168.55.2, 232.1.1.1) with PIM Snooping Disabled shows the multicast data streams with receiver H-8 joined and PIM snooping disabled.

Figure 245: Multicast Stream (192.168.55.2, 232.1.1.1) with PIM Snooping Disabled

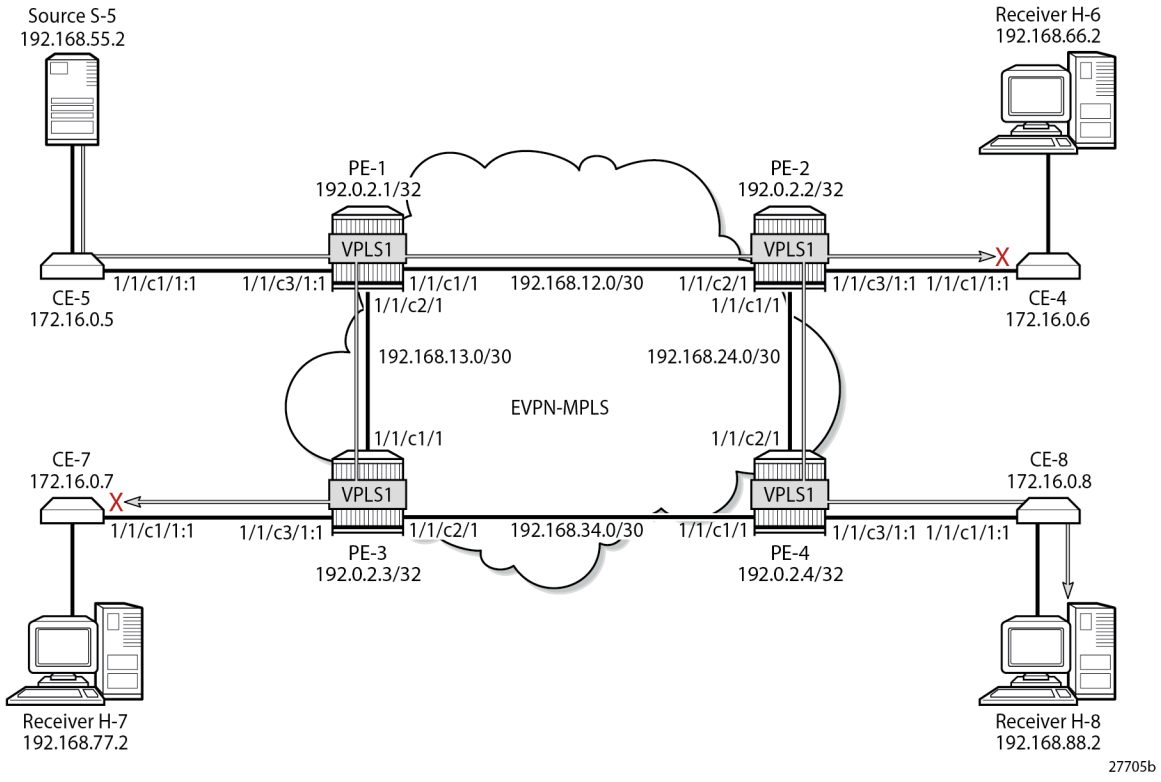

EVPN-MPLS VPLS with PIM Snooping Enabled

PIM snooping is enabled on all PEs as follows:

```
configure {
    service {
        vpls "VPLS 1" {
            pim-snooping { }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
    }
}
```

The default mode for PIM snooping is proxy mode, which allows the intermediate PEs to terminate the incoming PIM join or prune messages and create their own PIM join or prune message to be sent toward CE-5, as shown in Figure 246: H-8 Joins (192.168.55.2, 232.1.1.1) and PIM Snooping is Enabled in Proxy Mode:

Figure 246: H-8 Joins (192.168.55.2, 232.1.1.1) and PIM Snooping is Enabled in Proxy Mode

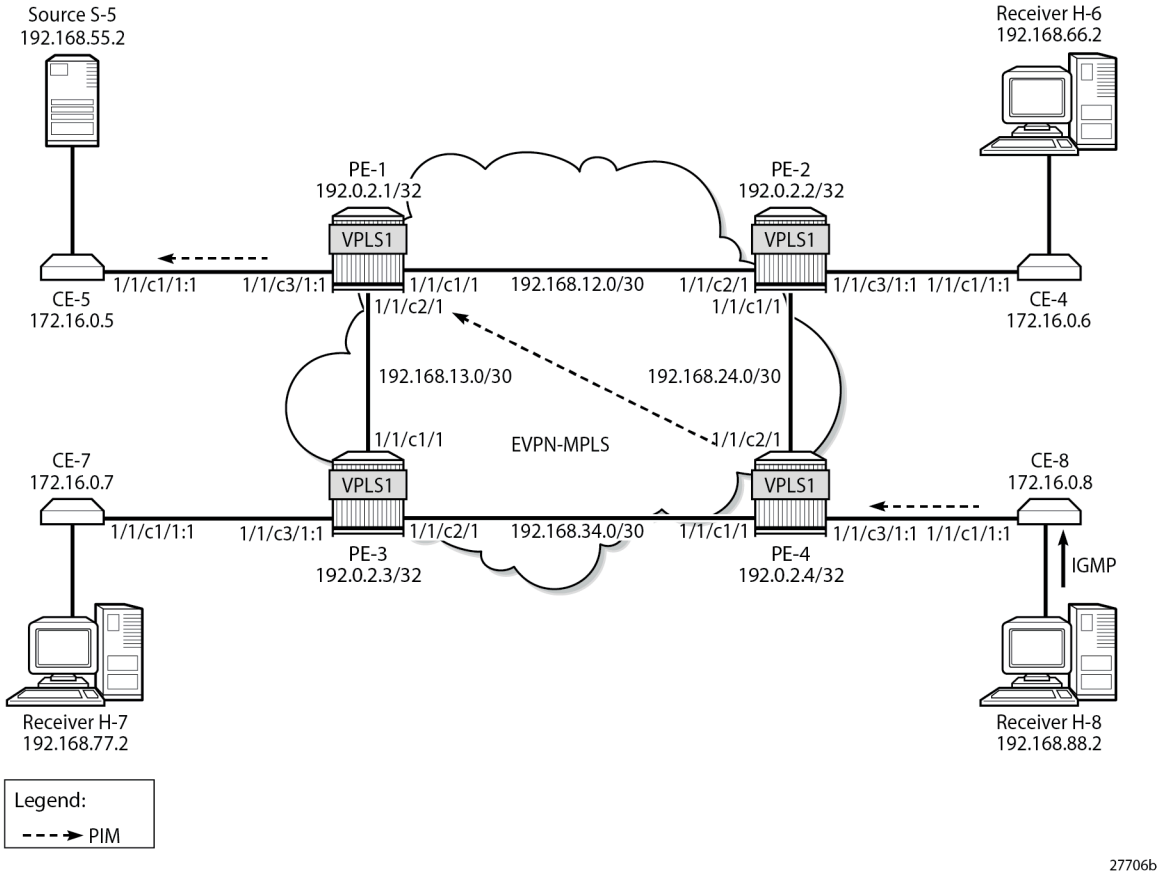

PE-4 receives the following PIM join message for multicast group (192.168.55.2, 232.1.1.1) from CE-8 to CE-5 on SAP 1/1/c3/1:1:

```
17 2023/08/10 22:59:27.162 CEST MINOR: DEBUG #2001 Base PIM[vpls 1 ]
"PIM[vpls 1 ]: Join/Prune
[000 00:24:58.740] PIM-RX ifId 1 ifName SAP:1/1/c3/1:1 172.16.0.8 -> 224.0.0.13 Length: 34
PIM Version: 2 Msg Type: Join/Prune Checksum: 0x4828
Upstream Nbr IP : 172.16.0.5 Resvd: 0x0, Num Groups 1, HoldTime 210
Group: 232.1.1.1/32 Num Joined Srcs: 1, Num Pruned Srcs: 0
Joined Srcs:
192.168.55.2/32 Flag S  <S,G>

"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

PE-4 sends the following PIM join message for multicast group (192.168.55.2, 232.1.1.1) to CE-5 on interface EVPN-MPLS:

```
18 2023/08/10 22:59:27.162 CEST MINOR: DEBUG #2001 Base PIM[vpls 1 ]
"PIM[vpls 1 ]: Join/Prune
[000 00:24:58.740] PIM-TX ifId 1071394 ifName EVPN-MPLS 0.0.0.0 -> 224.0.0.13 Length: 34
PIM Version: 2 Msg Type: Join/Prune Checksum: 0x4828
Upstream Nbr IP : 172.16.0.5 Resvd: 0x0, Num Groups 1, HoldTime 210
Group: 232.1.1.1/32 Num Joined Srcs: 1, Num Pruned Srcs: 0
Joined Srcs:
192.168.55.2/32 Flag S  <S,G>

"
```

In a similar way, PE-1 terminates this PIM join message and sends the following PIM join message for multicast group (192.168.55.2, 232.1.1.1) to CE-5 on SAP 1/1/c3/1:1.

```
19 2023/08/10 22:59:27.159 CEST MINOR: DEBUG #2001 Base PIM[vpls 1 ]
"PIM[vpls 1 ]: Join/Prune
[000 00:25:12.500] PIM-TX ifId 1 ifName SAP:1/1/c3/1:1 0.0.0.0 -> 224.0.0.13 Length: 34
PIM Version: 2 Msg Type: Join/Prune Checksum: 0x4828
Upstream Nbr IP : 172.16.0.5 Resvd: 0x0, Num Groups 1, HoldTime 210
Group: 232.1.1.1/32 Num Joined Srcs: 1, Num Pruned Srcs: 0
Joined Srcs:
192.168.55.2/32 Flag S  <S,G>

"
```

The following command shows the status of PIM snooping in VPLS 1 on PE-1:

```
[/]
A:admin@PE-1# show service id 1 pim-snooping status

===============================================================================
PIM Snooping Status ipv4
===============================================================================
Admin State                       : Up
Oper State                        : Up
Mode Admin                        : Proxy
Mode Oper                         : Proxy
Hold Time                         : 90
Designated Router                 : 172.16.0.8
J/P Tracking                      : Inactive
Up Time                           : 0d 00:01:57
Group Policy                      : None
===============================================================================
```

The following PIM snooping statistics show the number of received and transmitted PIM messages, and the source group statistics: one (S,G) group is joined and no (*,G) group.

```
[/]
A:admin@PE-1# show service id 1 pim-snooping statistics

===================================================================
PIM Snooping Statistics ipv4
===================================================================
Message Type        Received       Transmitted    Rx Errors
-------------------------------------------------------------------
Hello               34             -              0
Join Prune          2              2              0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Total Packets       36             2
-------------------------------------------------------------------------------
General Statistics
-------------------------------------------------------------------------------
Rx Neighbor Unknown              : 0
Rx Bad Checksum Discard          : 0
Rx Bad Encoding                  : 0
Rx Bad Version Discard           : 0
Join Policy Drops                : 0
-------------------------------------------------------------------------------
Source Group Statistics
-------------------------------------------------------------------------------
(S,G)                            : 1
(*,G)                            : 0
===============================================================
```

PE-4 has four neighbors for PIM snooping: the local SAP toward CE-8 and the EVPN-MPLS destinations toward the other CEs, as follows:

```
[/]
A:admin@PE-4# show service id 1 pim-snooping neighbor

===============================================================================
PIM Snooping Neighbors ipv4
===============================================================================
Port Id                 Nbr DR Prty     Up Time       Expiry Time     Hold Time
  Nbr Address
-------------------------------------------------------------------------------
SAP:1/1/c3/1:1          1               0d 00:01:46   0d 00:01:28     105
  172.16.0.8
EVPN-MPLS               1               0d 00:01:31   0d 00:01:43     105
  172.16.0.5
EVPN-MPLS               1               0d 00:01:41   0d 00:01:34     105
  172.16.0.6
EVPN-MPLS               1               0d 00:01:29   0d 00:01:15     105
  172.16.0.7
-------------------------------------------------------------------------------
Neighbors : 4
===============================================================================
```

The EVPN-MPLS destinations appear as a single entry with port ID "EVPN-MPLS" in the following **show** command:

```
[/]
A:admin@PE-4# show service id 1 pim-snooping port

===============================================================================
PIM Snooping Port ipv4
===============================================================================
Port Id                                                     Opr   PW Fwding
-------------------------------------------------------------------------------
SAP:1/1/c3/1:1                                              Up    Actv
EVPN-MPLS                                                   Up    Actv
===============================================================================
```

In the MFIB output on PE-1 and PE-4, each EVPN-MPLS destination is shown individually, but the information for each EVPN-MPLS destination is identical, as follows:

```
[/]
A:admin@PE-1# show service id 1 mfib
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Multicast FIB, Service 1
===============================================================================
Source Address  Group Address         Port Id                     Svc Id   Fwd
                                                                         Blk
-------------------------------------------------------------------------------
192.168.55.2    232.1.1.1             sap:1/1/c3/1:1              Local    Fwd
                                      mpls:192.0.2.2:524282       Local    Fwd
                                      mpls:192.0.2.3:524282       Local    Fwd
                                      mpls:192.0.2.4:524282       Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
```

On PE-2 and PE-3, the MFIB has no entries, as follows:

```
[/]
A:admin@PE-2# show service id 1 mfib

===============================================================================
Multicast FIB, Service 1
===============================================================================
Source Address  Group Address         Port Id                     Svc Id   Fwd
                                                                         Blk
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Number of entries: 0
===============================================================================
```

The MFIB statistics for VPLS 1 on PE-1 show the number of matched packets and matched octets for multicast group (192.168.55.2, 232.1.1.1), as follows:

```
[/]
A:admin@PE-1# show service id 1 mfib statistics

===============================================================================
Multicast FIB Statistics, Service 1
===============================================================================
Source Address  Group Address       Matched Pkts          Matched Octets
                                                          Forwarding Rate
-------------------------------------------------------------------------------
192.168.55.2    232.1.1.1           92556                 138834000
                                                          9867.357 kbps
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
```

The following **show** command of the PIM group snooped on PE-1 shows the SAP toward the source as incoming interface, and the EVPN-MPLS interface as outgoing interface (traffic coming in from the source is not sent back to the SAP toward the source):

```
[/]
A:admin@PE-1# show service id 1 pim-snooping group 232.1.1.1 detail

===============================================================================
PIM Snooping Source Group ipv4
===============================================================================
Group Address      : 232.1.1.1
Source Address     : 192.168.55.2
Up Time            : 0d 00:01:35
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Up JP State        : Joined             Up JP Expiry       : 0d 00:00:25
Up JP Rpt          : Not Joined StarG   Up JP Rpt Override : 0d 00:00:00

RPF Neighbor       : 172.16.0.5
Incoming Intf      : SAP:1/1/c3/1:1
Outgoing Intf List : EVPN-MPLS, SAP:1/1/c3/1:1

Forwarded Packets  : 78273              Forwarded Octets   : 117409500
-------------------------------------------------------------------------------
Groups : 1
===============================================================================
```

The following identical **show** command of the PIM group snooped on PE-4 shows the EVPN-MPLS interface as incoming interface. Even though the EVPN-MPLS interface is also listed as outgoing interface, traffic coming from that interface is not forwarded on that interface (all EVPN-MPLS destinations are treated as one single EVPN-MPLS interface), so the traffic is forwarded to the SAP toward the receiving CE only.

```
[/]
A:admin@PE-4# show service id 1 pim-snooping group 232.1.1.1 detail

===============================================================================
PIM Snooping Source Group ipv4
===============================================================================
Group Address      : 232.1.1.1
Source Address     : 192.168.55.2
Up Time            : 0d 00:01:40

Up JP State        : Joined             Up JP Expiry       : 0d 00:00:20
Up JP Rpt          : Not Joined StarG   Up JP Rpt Override : 0d 00:00:00

RPF Neighbor       : 172.16.0.5
Incoming Intf      : EVPN-MPLS
Outgoing Intf List : EVPN-MPLS, SAP:1/1/c3/1:1

Forwarded Packets  : 82285              Forwarded Octets   : 123098360
-------------------------------------------------------------------------------
Groups : 1
===============================================================================
```

The following port statistics on PE-2 show that the multicast stream coming in from PE-1 on port 1/1/c2/1 is forwarded to port 1/1/c1/1 toward PE-4 only, but not to port 1/1/c3/1 toward CE-6:

```
[/]
A:admin@PE-2# show port 1/1/c1/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c1/1                                 32                           3268
                                      27693                       41994078
===============================================================================

[/]
A:admin@PE-2# show port 1/1/c2/1 statistics

===============================================================================
Port Statistics on Slot 1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c2/1                                27695                       41994284
                                           34                           3451
===============================================================================

[/]
A:admin@PE-2# show port 1/1/c3/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c3/1                                    1                             76
                                            3                            228
===============================================================================
```

In a similar way, the multicast traffic on PE-3 that comes in from PE-1 via port 1/1/c1/1 is not forwarded to any port, as follows:

```
[/]
A:admin@PE-3# show port 1/1/c1/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c1/1                                27473                       41660165
                                           31                           3195
===============================================================================

[/]
A:admin@PE-3# show port 1/1/c2/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c2/1                                   26                            2700
                                           26                            2747
===============================================================================

[/]
A:admin@PE-3# show port 1/1/c3/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c3/1                                    1                             76
                                            3                            228
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 247: Multicast Stream (192.168.55.2, 232.1.1.1) with PIM Snooping Enabled shows that the multicast stream still flows from the source S-5 to the receiver H-8, but is not forwarded to CE-6 and CE-7 when PIM snooping is enabled. The root node PE-1 sends the multicast traffic received on the SAP to all EVPN-MPLS destinations over the P2MP mLDP provider tunnel. The EVPN-MPLS interface is treated as a single interface.

Figure 247: Multicast Stream (192.168.55.2, 232.1.1.1) with PIM Snooping Enabled

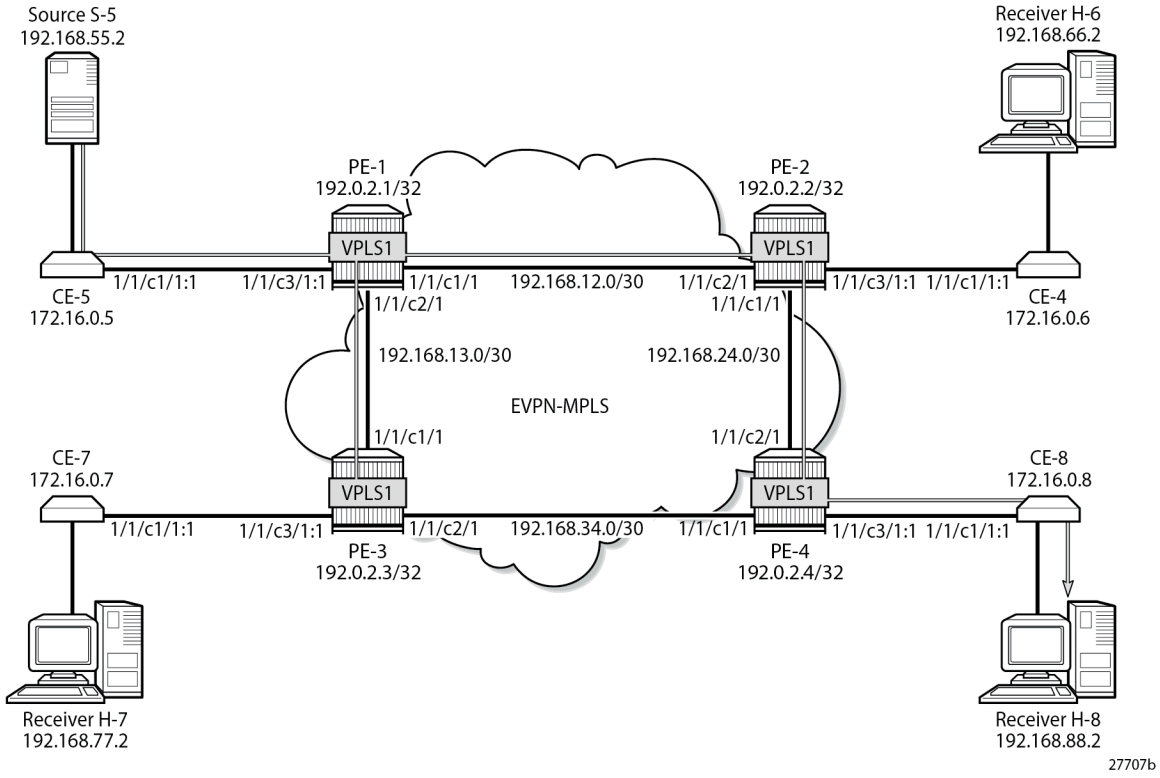

Multi-homed EVPN-MPLS VPLS without PIM Snooping

When CE-5 receives a PIM join message, it forwards the multicast stream to PE-1. All multicast traffic in VPLS 2 is sent to all receiving CEs, regardless of the received PIM join messages.

Figure 248: Example Topology with Multi-homing ESs shows the example topology with an all-active multi-homing virtual Ethernet Segment (ES) "ESI-34_2" between PE-3 and PE-4 using a LAG, and a single-active multi-homing ES "ESI-24" between PE-2 and PE-4 using SDPs.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 248: Example Topology with Multi-homing ESs

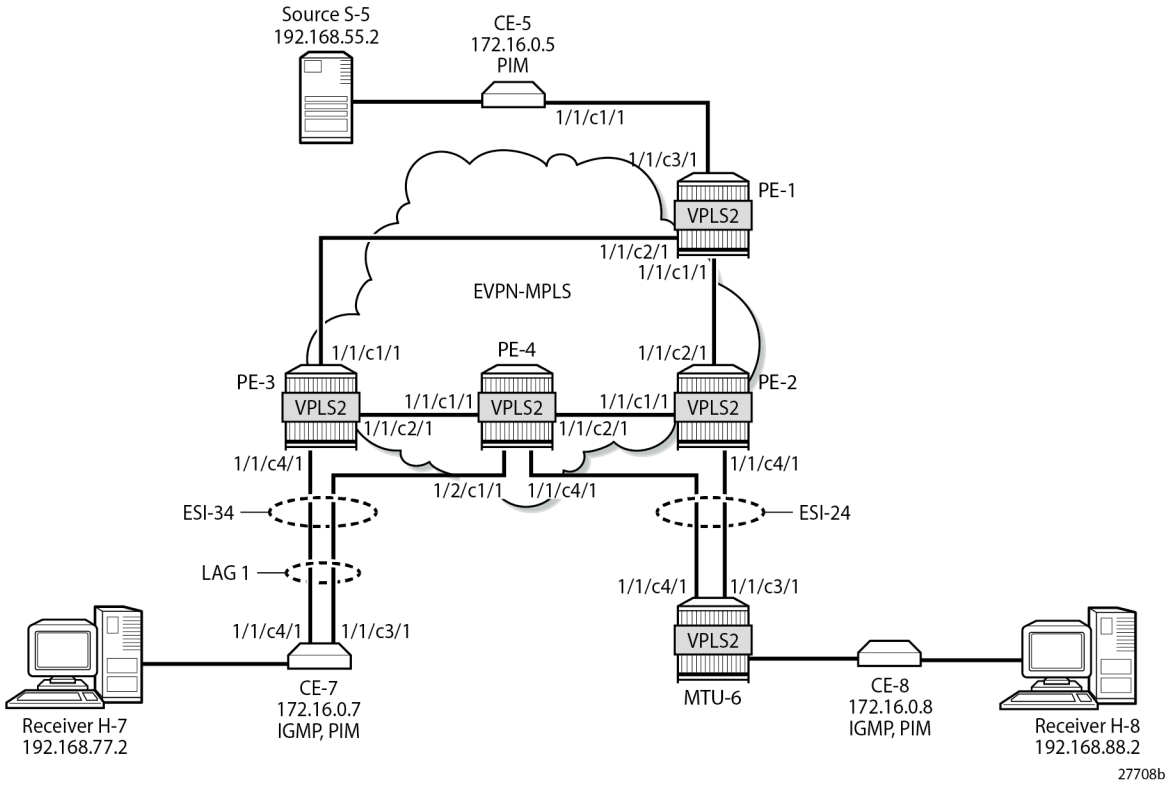

The configuration of VPLS 2 is similar to the configuration of VPLS 1 on all PEs. An identical P2MP mLDP provider tunnel is established on the PEs for VPLS 2: PE-1 is the root node, PE-2 is a leaf node and a transit node, PE-3 is a leaf node, and PE-4 is also a leaf node.

On PE-2, PE-3, and PE-4, one or more ESs are configured. The service configuration on PE-2 is as follows. An SDP is configured toward MTU-6 that is associated with a single-active multi-homing ES "ESI-24". Spoke-SDP 26:2 is associated with VPLS 2.

```
On PE-2:
configure {
    service {
        sdp 26 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.6
            }
        }
        vpls "VPLS 2" {
            admin-state enable
            service-id 2
            customer "1"
            bgp 1 { }
            bgp-evpn {
                evi 2
                mpls 1 {
                    admin-state enable
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    ingress-replication-bum-label true
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            spoke-sdp 26:2 { }
            provider-tunnel {
                inclusive {
                    admin-state enable
                    owner bgp-evpn-mpls
                    mldp
                }
            }
        }
        system {
            bgp {
                evpn {
                    ethernet-segment "ESI-24" {
                        admin-state enable
                        esi 0x01000000002400000001
                        multi-homing-mode single-active
                        df-election {
                            es-activation-timer 3
                            service-carving-mode manual
                            manual {
                                preference {
                                    mode non-revertive
                                    value 10000
                                }
                            }
                        }
                        association {
                            sdp 26 { }
                        }
                    }
                }
            }
        }
    }
}
```

The same ES is configured on PE-4, together with another ES-an all-active multi-homing virtual ES that applies to VPLS 2 only (**q-tag-range 2**); see chapter Virtual Ethernet Segments. The preference for the DF election is configured manually to a value of 5000 (which is lower than the preference 10000 on PE-3); see chapter Preference-based and Non-revertive EVPN DF Election. VPLS 2 has a SAP and a spoke-SDP configured. The service configuration on PE-4 is as follows:

```
On PE-4:
configure exclusive
    service {
        sdp 46 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.6
            }
        }
        vpls "VPLS 2" {
            admin-state enable
            service-id 2
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            customer "1"
            bgp 1 { }
            bgp-evpn {
                evi 2
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            spoke-sdp 46:2 { }
            sap lag-1:2 { }
            provider-tunnel {
                inclusive {
                    admin-state enable
                    owner bgp-evpn-mpls
                    mldp
                }
            }
        }
        system {
            bgp {
                evpn {
                    ethernet-segment "ESI-24" {
                        admin-state enable
                        esi 0x01000000002400000001
                        multi-homing-mode single-active
                        df-election {
                            es-activation-timer 3
                            service-carving-mode manual
                            manual {
                                preference {
                                    mode non-revertive
                                    value 5000
                                }
                            }
                        }
                        association {
                            sdp 46 { }
                        }
                    }
                    ethernet-segment "ESI-34_2" {
                        admin-state enable
                        type virtual
                        esi 0x01000000003402000001
                        multi-homing-mode all-active
                        df-election {
                            es-activation-timer 3
                            service-carving-mode manual
                            manual {
                                preference {
                                    mode non-revertive
                                    value 5000
                                }
                            }
                        }
                        association {
                            lag "lag-1" {
                                virtual-ranges {
                                    dot1q {
                                        q-tag 2 {
                                            end 2
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                                        }
                                    }
                                }
                            }
                        }
                    }
                }
            }
        }
    }
}
```

The service configuration on PE-3 includes the same all-active multi-homing virtual ES with preference 10000, as follows:

```
On PE-3:
configure {
    service {
        vpls "VPLS 2" {
            admin-state enable
            service-id 2
            customer "1"
            bgp 1 { }
            bgp-evpn {
                evi 2
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            sap lag-1:2 { }
            provider-tunnel {
                inclusive {
                    admin-state enable
                    owner bgp-evpn-mpls
                    mldp
                }
            }
        }
        system {
            bgp {
                evpn {
                    ethernet-segment "ESI-34_2" {
                        admin-state enable
                        type virtual
                        esi 0x0100000000003402000001
                        multi-homing-mode all-active
                        df-election {
                            es-activation-timer 3
                            service-carving-mode manual
                            manual {
                                preference {
                                    mode non-revertive
                                    value 10000
                                }
                            }
                        }
                        association {
                            lag "lag-1" {
                                virtual-ranges {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                                    dot1q {
                                        q-tag 2 {
                                            end 2
                                        }
                                    }
                                }
                            }
                        }
                    }
                }
            }
        }
    }
}
```

The following is the service configuration on MTU-6:

```
On MTU-6:
configure {
    service {
        sdp 62 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.2
            }
        }
        sdp 64 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.4
            }
        }
        vpls "VPLS 2" {
            admin-state enable
            service-id 2
            customer "1"
            endpoint "x" { }
            spoke-sdp 62:2 {
                endpoint {
                    name "x"
                }
                stp {
                    admin-state disable
                }
            }
            spoke-sdp 64:2 {
                endpoint {
                    name "x"
                }
                stp {
                    admin-state disable
                }
            }
            sap 1/2/c1/1:2 { }
        }
    }
}
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

For VPLS 2, PE-2 is the DF in ES "ESI-24", as follows:

```
[/]
A:admin@PE-2# show service id 2 ethernet-segment
No sap entries

===============================================================================
SDP Ethernet-Segment Information
===============================================================================
SDP                  Eth-Seg                          Status
-------------------------------------------------------------------------------
26:2                 ESI-24                           DF
===============================================================================
No vxlan instance entries
```

PE-3 is the DF in ES "ESI-34_2", as follows:

```
[/]
A:admin@PE-3# show service id 2 ethernet-segment

===============================================================================
SAP Ethernet-Segment Information
===============================================================================
SAP                  Eth-Seg                          Status
-------------------------------------------------------------------------------
lag-1:2              ESI-34_2                         DF
===============================================================================
No sdp entries
No vxlan instance entries
```

PE-4 is NDF for both ESI-24 and ESI-34_2, as follows:

```
[/]
A:admin@PE-4# show service id 2 ethernet-segment

===============================================================================
SAP Ethernet-Segment Information
===============================================================================
SAP                  Eth-Seg                          Status
-------------------------------------------------------------------------------
lag-1:2              ESI-34_2                         NDF
===============================================================================

===============================================================================
SDP Ethernet-Segment Information
===============================================================================
SDP                  Eth-Seg                          Status
-------------------------------------------------------------------------------
46:2                 ESI-24                           NDF
===============================================================================
No vxlan instance entries
```

When H-8 sends an IGMP report to join multicast group 232.1.1.1 from source 192.168.55.2, CE-5 forwards the multicast stream after receiving the corresponding PIM join message. PE-1 forwards the multicast traffic on the P2MP mLDP tree to PE-2, PE-3, and PE-4. The DF PE-2 forwards the traffic to MTU-6, and DF PE-3 forwards it to CE-7, even though a PIM join for this group has not been received from CE-7. PE-4 is NDF, so it does not forward the traffic to MTU-6 or CE-7. MTU-6 forwards the traffic to CE-8, which sends it to H-8. CE-7 drops the multicast traffic because no attached receiver has joined the multicast group. Figure 249: EVPN-MPLS with Multi-homing – Receiver H-8 Joined shows how this multicast is forwarded when PIM snooping is disabled.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 249: EVPN-MPLS with Multi-homing – Receiver H-8 Joined

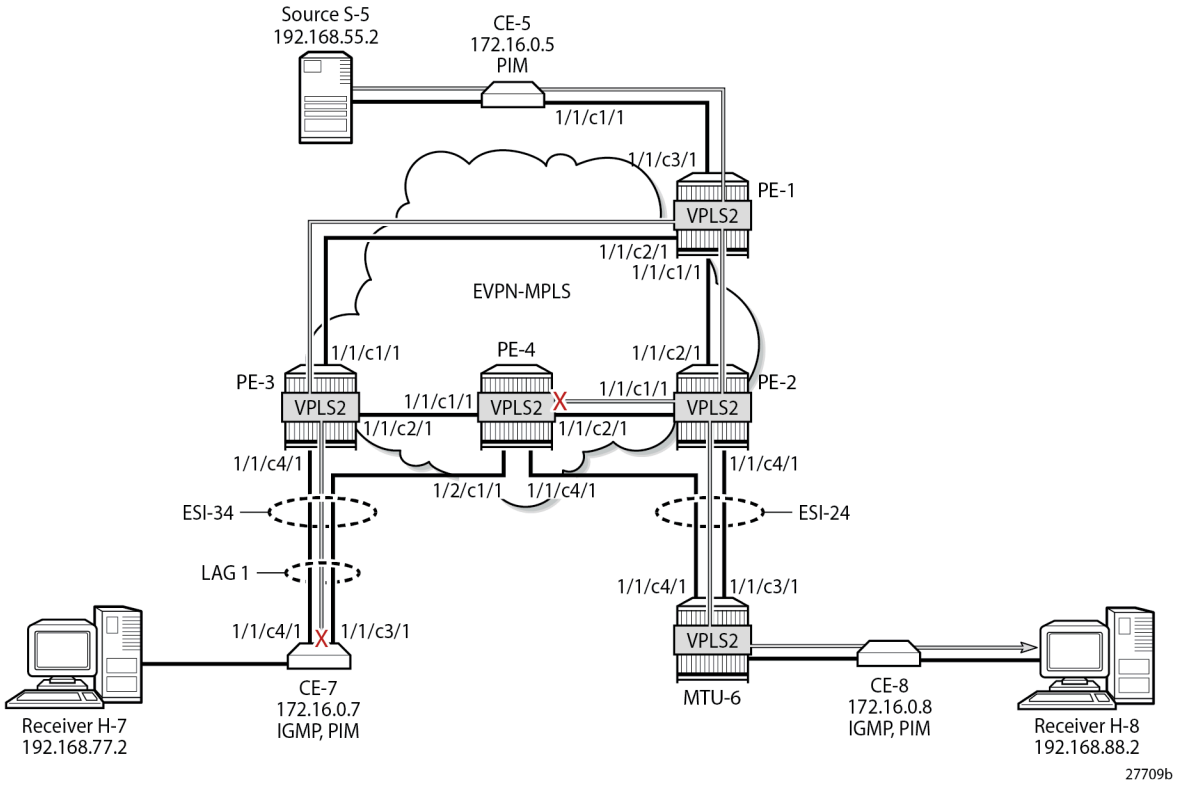

The static IGMP multicast is removed to emulate an IGMPv3 report from receiver H-8 to exclude multicast group 232.1.1.1 from source 192.168.55.2, as follows:

```
On CE-8:
configure {
    router "Base" {
        igmp {
            interface "int-CE-8-H-8" {
                static {
                    delete group 232.1.1.1
                }
            }
        }
    }
}
```

Multi-homed EVPN-MPLS VPLS with PIM Snooping

PIM snooping is enabled in VPLS 2 on all PEs, including PE-1, which is not part of an ES-with the following command:

```
configure {
    service {
        vpls "VPLS 2" {
            pim-snooping { }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
    }
}
```

All PEs have three PIM snooping neighbors: CE-5, CE-7, and CE-8. The list of PIM snooping neighbors on PE-1 is as follows:

```
[/]
A:admin@PE-1# show service id 2 pim-snooping neighbor

===============================================================================
PIM Snooping Neighbors ipv4
===============================================================================
Port Id                 Nbr DR Prty     Up Time       Expiry Time     Hold Time
  Nbr Address
-------------------------------------------------------------------------------
SAP:1/1/c3/1:2          1               0d 00:01:12   0d 00:01:33     105
  172.16.0.5
EVPN-MPLS               1               0d 00:01:14   0d 00:01:31     105
  172.16.0.7
EVPN-MPLS               1               0d 00:01:07   0d 00:01:38     105
  172.16.0.8
-------------------------------------------------------------------------------
Neighbors : 3
===============================================================================
```

On PE-2, the same PIM snooping neighbors are listed: CE-5, CE-7, and CE-8, as follows:

```
[/]
A:admin@PE-2# show service id 2 pim-snooping neighbor

===============================================================================
PIM Snooping Neighbors ipv4
===============================================================================
Port Id                 Nbr DR Prty     Up Time       Expiry Time     Hold Time
  Nbr Address
-------------------------------------------------------------------------------
SPOKE_SDP:26:2          1               0d 00:01:09   0d 00:01:35     105
  172.16.0.8
EVPN-MPLS               1               0d 00:01:15   0d 00:01:30     105
  172.16.0.5
EVPN-MPLS               1               0d 00:01:17   0d 00:01:27     105
  172.16.0.7
-------------------------------------------------------------------------------
Neighbors : 3
===============================================================================
```

PE-3 and PE-4 also have these three CEs as PIM snooping neighbors.

All-active MH EVPN-MPLS VPLS with PIM Snooping

On CE-7, the following static IGMP membership is configured on interface int-CE-7-H-7:

```
On CE-7:
configure {
    router "Base" {
        igmp {
            interface "int-CE-7-H-7" {
                ssm-translate {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    group-range start 232.0.0.0 end 232.255.255.255 {
                        source 192.168.55.2 { }
                    }
                }
                static {
                    group 232.1.1.1 {
                        source 192.168.55.2 { }
                    }
                }
            }
        }
    }
}
```

When H-7 joins the multicast group 232.1.1.1 via source 192.168.55.2, the PIM join messages are snooped by the PEs and the MFIB is built. The MFIB on PE-1 contains one entry for group address 232.1.1.1 and source address 192.168.55.2, with four port IDs: the local SAP to CE-5 and the EVPN-MPLS destinations, as follows:

```
[/]
A:admin@PE-1# show service id 2 mfib

===============================================================================
Multicast FIB, Service 2
===============================================================================
Source Address  Group Address         Port Id                     Svc Id   Fwd
                                                                        Blk
-------------------------------------------------------------------------------
192.168.55.2    232.1.1.1             sap:1/1/c3/1:2              Local    Fwd
                                      mpls:192.0.2.2:524280       Local    Fwd
                                      mpls:192.0.2.3:524280       Local    Fwd
                                      mpls:192.0.2.4:524280       Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
```

The MFIB on PE-2 is empty because no locally attached node has sent a PIM join for any multicast group:

```
[/]
A:admin@PE-2# show service id 2 mfib

===============================================================================
Multicast FIB, Service 2
===============================================================================
Source Address  Group Address         Port Id                     Svc Id   Fwd
                                                                        Blk
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Number of entries: 0
===============================================================================
```

The MFIB on PE-3 contains an entry for the (S,G) with the local SAP lag-1:2 and the EVPN-MPLS destination, as follows:

```
[/]
A:admin@PE-3# show service id 2 mfib

===============================================================================
Multicast FIB, Service 2
===============================================================================
Source Address  Group Address         Port Id                     Svc Id   Fwd
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                                                                                Blk
-------------------------------------------------------------------------------
192.168.55.2    232.1.1.1            sap:lag-1:2                  Local    Fwd
                                     mpls:192.0.2.1:524280        Local    Fwd
                                     mpls:192.0.2.2:524280        Local    Fwd
                                     mpls:192.0.2.4:524280        Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
```

Data-driven PIM state synchronization between PE-3 and PE-4 in the ESI-34_2 results in the following MFIB entry on PE-4:

```
[/]
A:admin@PE-4# show service id 2 mfib

===============================================================================
Multicast FIB, Service 2
===============================================================================
Source Address  Group Address         Port Id                      Svc Id   Fwd
                                                                           Blk
-------------------------------------------------------------------------------
192.168.55.2    232.1.1.1            sap:lag-1:2                  Local    Fwd
                                     mpls:192.0.2.1:524280        Local    Fwd
                                     mpls:192.0.2.2:524280        Local    Fwd
                                     mpls:192.0.2.3:524280        Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
```

When debugging is enabled on the PEs as follows, the synchronization between peers in ES "ESI-34_2" is logged:

```
debug {
    service {
        vpls "VPLS 2" {
            pim-snooping {
                events {
                    port {
                        evpn-mpls
                    }
                    jp { }
                }
                packet {
                    packet-types {
                        jp true
                    }
                }
            }
        }
    }
}
```

For example, PE-4 sends the following PIM message to its remote peer PE-3 in ESI-34_2:

```
82 2023/08/10 23:13:20.784 CEST MINOR: DEBUG #2001 Base PIM[vpls 2 ]
"PIM[vpls 2 ]: pimVplsFwdJPToEvpn
Forwarding to remote peer on bgp-evpn ethernet-segment ESI-34_2"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

PE-3 receives the following PIM message from its remote peer PE-4 in ESI-34_2:

```
74 2023/08/10 23:13:20.786 CEST MINOR: DEBUG #2001 Base PIM[vpls 2 ]
"PIM[vpls 2 ]: pimProcessPdu
Received from remote peer on bgp-evpn ethernet-segment ESI-34_2, will be applied on lag-1:2
"
```

On PE-1, the PIM snooping group (192.168.55.2, 232.1.1.1) has incoming interface SAP 1/1/c3/1:2 toward CE-5 and the EVPN-MPLS interface as outgoing interface, as follows:

```
[/]
A:admin@PE-1# show service id 2 pim-snooping group detail

===============================================================================
PIM Snooping Source Group ipv4
===============================================================================
Group Address      : 232.1.1.1
Source Address     : 192.168.55.2
Up Time            : 0d 00:01:10

Up JP State        : Joined             Up JP Expiry       : 0d 00:00:50
Up JP Rpt          : Not Joined StarG   Up JP Rpt Override : 0d 00:00:00

RPF Neighbor       : 172.16.0.5
Incoming Intf      : SAP:1/1/c3/1:2
Outgoing Intf List : EVPN-MPLS, SAP:1/1/c3/1:2

Forwarded Packets  : 57863              Forwarded Octets   : 86794500
-------------------------------------------------------------------------------
Groups : 1
===============================================================================
```

On PE-2, no PIM join messages are received and no groups are listed, as follows:

```
[/]
A:admin@PE-2# show service id 2 pim-snooping group detail

===============================================================================
PIM Snooping Source Group ipv4
===============================================================================
No Matching Entries
===============================================================================
```

On PE-3, the same PIM snooping group has the EVPN-MPLS as incoming interface and the SAP lag-1:2 as outgoing interface. The split-horizon mechanism ensures that the multicast traffic that enters through the EVPN-MPLS interface is not forwarded on the EVPN-MPLS interface, which is regarded as a single interface.

```
[/]
A:admin@PE-3# show service id 2 pim-snooping group detail

===============================================================================
PIM Snooping Source Group ipv4
===============================================================================
Group Address      : 232.1.1.1
Source Address     : 192.168.55.2
Up Time            : 0d 00:01:13

Up JP State        : Joined             Up JP Expiry       : 0d 00:00:02
Up JP Rpt          : Not Joined StarG   Up JP Rpt Override : 0d 00:00:00
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
RPF Neighbor       : 172.16.0.5
Incoming Intf      : EVPN-MPLS
Outgoing Intf List : EVPN-MPLS, SAP:lag-1:2

Forwarded Packets  : 60286             Forwarded Octets   : 90187856
-------------------------------------------------------------------------------
Groups : 1
===============================================================================
```

On PE-4, the same PIM snooping information is available, because of the data-driven PIM state synchronization between PE-3 and PE-4 in ESI-34_2, as follows:

```
[/]
A:admin@PE-4# show service id 2 pim-snooping group detail

===============================================================================
PIM Snooping Source Group ipv4
===============================================================================
Group Address      : 232.1.1.1
Source Address     : 192.168.55.2
Up Time            : 0d 00:01:14

Up JP State        : Joined            Up JP Expiry       : 0d 00:00:59
Up JP Rpt          : Not Joined StarG  Up JP Rpt Override : 0d 00:00:00

RPF Neighbor       : 172.16.0.5
Incoming Intf      : EVPN-MPLS
Outgoing Intf List : EVPN-MPLS, SAP:lag-1:2

Forwarded Packets  : 61554             Forwarded Octets   : 92084784
-------------------------------------------------------------------------------
Groups : 1
===============================================================================
```

Figure 250: EVPN-MPLS with All-active Multi-homing and PIM Snooping Enabled – Receiver H-7 Joined shows how the multicast traffic is forwarded when H-7 joins the multicast group and PIM snooping is enabled. DF PE-3 forwards the traffic toward CE-7. The multicast stream also reaches PE-2 and PE-4, where it is dropped.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 250: EVPN-MPLS with All-active Multi-homing and PIM Snooping Enabled – Receiver H-7 Joined

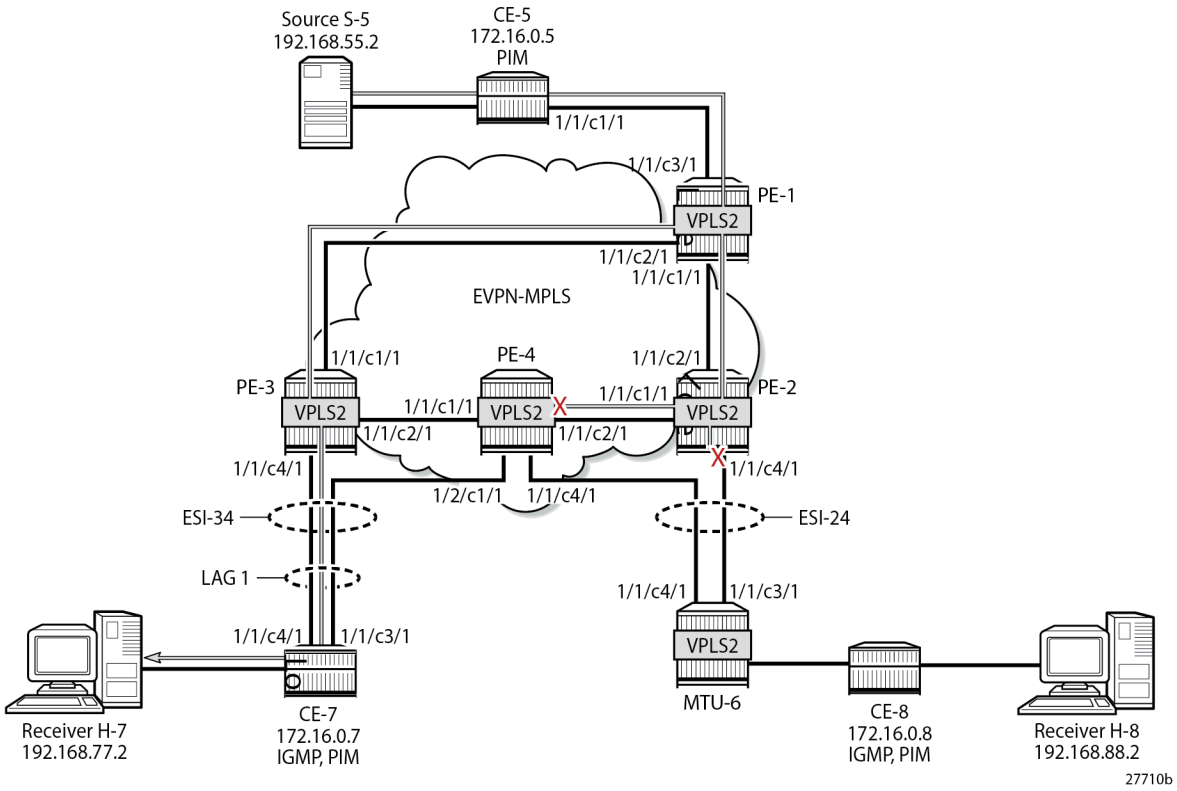

H-7 leaves the multicast group and H-8 joins it instead.

Single-active MH EVPN-MPLS VPLS with PIM Snooping

When H-8 joins the multicast group and PIM snooping is enabled, only DF PE-2 forwards traffic from the EVPN-MPLS toward a receiver. PE-3 does not forward traffic to CE-7 because no PIM join message was received from CE-7. Figure 251: EVPN-MPLS with Single-active Multi-homing and PIM Snooping Enabled – Receiver H-8 Joined shows how the multicast traffic is forwarded when H-8 joins the multicast group and PIM snooping is enabled.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 251: EVPN-MPLS with Single-active Multi-homing and PIM Snooping Enabled – Receiver H-8 Joined

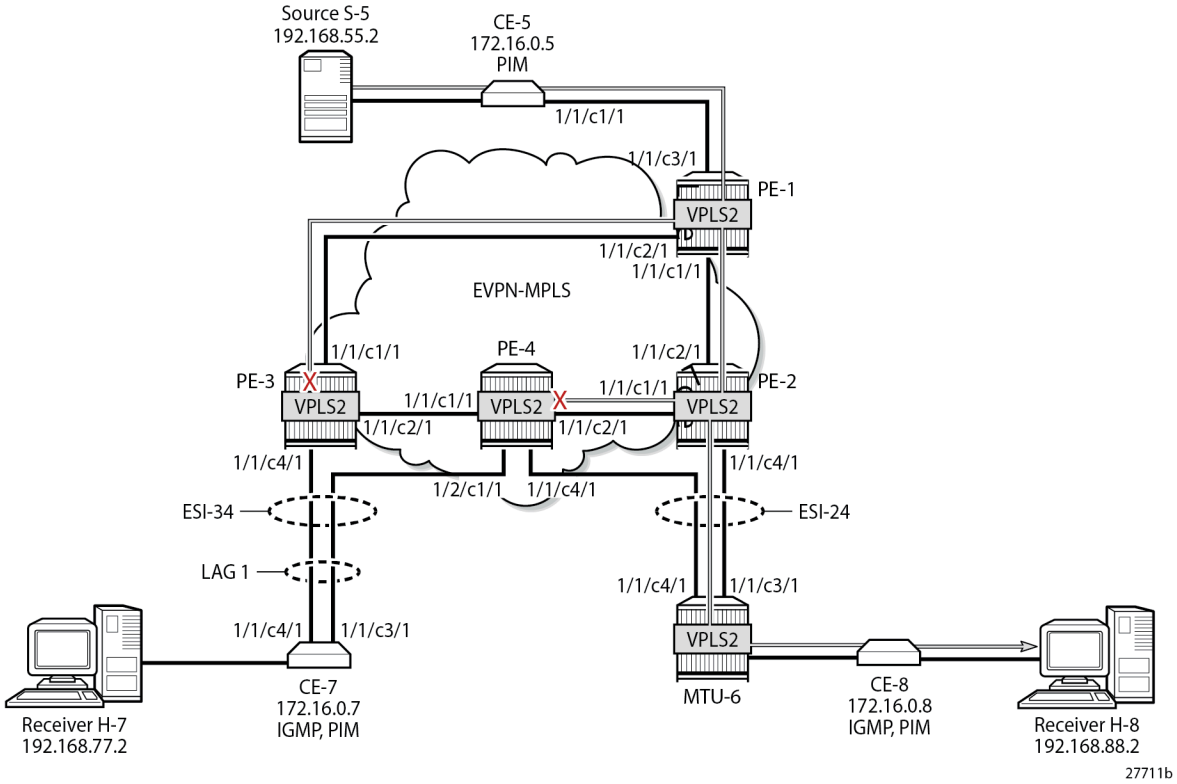

On PE-1, the MFIB looks the same as in the preceding case, as follows:

```
[/]
A:admin@PE-1# show service id 2 mfib

===============================================================================
Multicast FIB, Service 2
===============================================================================
Source Address  Group Address         Port Id                      Svc Id   Fwd
                                                                          Blk
-------------------------------------------------------------------------------
192.168.55.2    232.1.1.1             sap:1/1/c3/1:2               Local    Fwd
                                      mpls:192.0.2.2:524280        Local    Fwd
                                      mpls:192.0.2.3:524280        Local    Fwd
                                      mpls:192.0.2.4:524280        Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
```

On PE-2, the MFIB contains an entry for source address 192.168.55.2 and group address 232.1.1.1 with spoke-SDP 26:2 and the EVPN-MPLS destinations to the other PEs, as follows:

```
[/]
A:admin@PE-2# show service id 2 mfib

===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Multicast FIB, Service 2
===============================================================================
Source Address  Group Address         Port Id                     Svc Id   Fwd
                                                                        Blk
-------------------------------------------------------------------------------
192.168.55.2    232.1.1.1             sdp:26:2                    Local    Fwd
                                      mpls:192.0.2.1:524280       Local    Fwd
                                      mpls:192.0.2.3:524280       Local    Fwd
                                      mpls:192.0.2.4:524280       Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
```

The MFIB on PE-3 is empty, because multicast traffic toward H-8 is not sent via PE-3, as follows:

```
[/]
A:admin@PE-3# show service id 2 mfib

===============================================================================
Multicast FIB, Service 2
===============================================================================
Source Address  Group Address         Port Id                     Svc Id   Fwd
                                                                        Blk
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Number of entries: 0
===============================================================================
```

The data-driven PIM state synchronization ensures that DF PE-2 sends updates to NDF PE-4. With debugging enabled, the following debug message is displayed at PE-2:

```
205 2023/08/10 23:15:21.053 CEST MINOR: DEBUG #2001 Base PIM[vpls 2 ]
"PIM[vpls 2 ]: pimVplsFwdJPToEvpn
Forwarding to remote peer on bgp-evpn ethernet-segment ESI-24"
```

The following debug message is displayed at PE-4:

```
122 2023/08/10 23:15:21.053 CEST MINOR: DEBUG #2001 Base PIM[vpls 2 ]
"PIM[vpls 2 ]: pimProcessPdu
Received from remote peer on bgp-evpn ethernet-segment ESI-24, will be applied on 46:2
"
```

As a result, the MFIB on PE-4 is not empty, as follows:

```
[/]
A:admin@PE-4# show service id 2 mfib

===============================================================================
Multicast FIB, Service 2
===============================================================================
Source Address  Group Address         Port Id                     Svc Id   Fwd
                                                                        Blk
-------------------------------------------------------------------------------
192.168.55.2    232.1.1.1             sdp:46:2                    Local    Fwd
                                      mpls:192.0.2.1:524280       Local    Fwd
                                      mpls:192.0.2.2:524280       Local    Fwd
                                      mpls:192.0.2.3:524280       Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

On PE-1, the PIM snooping group (192.168.55.2, 232.1.1.1) has incoming interface SAP 1/1/c3/1:2 toward CE-5 and the EVPN-MPLS interface as outgoing interface, as follows:

```
[/]
A:admin@PE-1# show service id 2 pim-snooping group detail

===============================================================================
PIM Snooping Source Group ipv4
===============================================================================
Group Address      : 232.1.1.1
Source Address     : 192.168.55.2
Up Time            : 0d 00:01:09

Up JP State        : Joined             Up JP Expiry       : 0d 00:00:50
Up JP Rpt          : Not Joined StarG   Up JP Rpt Override : 0d 00:00:00

RPF Neighbor       : 172.16.0.5
Incoming Intf      : SAP:1/1/c3/1:2
Outgoing Intf List : EVPN-MPLS, SAP:1/1/c3/1:2

Forwarded Packets  : 57409              Forwarded Octets   : 86113500
-------------------------------------------------------------------------------
Groups : 1
===============================================================================
```

On PE-2, the same PIM snooping group has the EVPN-MPLS as incoming interface and the spoke-SDP 26:2 as outgoing interface. Again, the split-horizon mechanism ensures that the multicast traffic that enters through the EVPN-MPLS interface is not forwarded on the EVPN-MPLS interface, which is regarded as a single interface.

```
[/]
A:admin@PE-2# show service id 2 pim-snooping group detail

===============================================================================
PIM Snooping Source Group ipv4
===============================================================================
Group Address      : 232.1.1.1
Source Address     : 192.168.55.2
Up Time            : 0d 00:01:12

Up JP State        : Joined             Up JP Expiry       : 0d 00:01:14
Up JP Rpt          : Not Joined StarG   Up JP Rpt Override : 0d 00:00:00

RPF Neighbor       : 172.16.0.5
Incoming Intf      : EVPN-MPLS
Outgoing Intf List : EVPN-MPLS, SPOKE_SDP:26:2

Forwarded Packets  : 59634              Forwarded Octets   : 89212464
-------------------------------------------------------------------------------
Groups : 1
===============================================================================
```

On PE-3, no PIM join messages are received and no groups are listed, as follows:

```
[/]
A:admin@PE-3# show service id 2 pim-snooping group detail

===============================================================================
PIM Snooping Source Group ipv4
===============================================================================
No Matching Entries
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
```

On PE-4, the same PIM snooping information is available, because of the data-driven PIM state synchronization between PE-2 and PE-4 in ESI-24, as follows. The incoming interface is the EVPN-MPLS interface and the outgoing interface is spoke-SDP 46:2.

```
[/]
A:admin@PE-4# show service id 2 pim-snooping group detail

===============================================================================
PIM Snooping Source Group ipv4
===============================================================================
Group Address      : 232.1.1.1
Source Address     : 192.168.55.2
Up Time            : 0d 00:01:15

Up JP State        : Joined             Up JP Expiry       : 0d 00:00:52
Up JP Rpt          : Not Joined StarG   Up JP Rpt Override : 0d 00:00:00

RPF Neighbor       : 172.16.0.5
Incoming Intf      : EVPN-MPLS
Outgoing Intf List : EVPN-MPLS, SPOKE_SDP:46:2

Forwarded Packets  : 62461              Forwarded Octets   : 93441656
-------------------------------------------------------------------------------
Groups : 1
===============================================================================
```

PIM state synchronization is data-driven, so the PIM states are not stored in a database. Therefore, the ESs must be configured as **non-revertive** to avoid reverting back to the preferred PE while this PE is unaware of the PIM states.

PIM Snooping with Multi-chassis Synchronization

Data-driven PIM state synchronization is supported in SR OS Release 15.0.R4, and later. The ES must be configured as non-revertive, so that after a failover, the new DF remains the DF even when the original DF is operational again. When data-driven PIM state synchronization cannot be used, for example, when the service carving is configured in auto mode, or when the SR OS Release is an earlier release of 15.0, Multi-chassis synchronization (MCS) can be configured for a faster failover. MCS of the PIM snooping state on SAPs and spoke-SDPs is supported between an active and a standby PE and the PIM states are stored in a synchronization database. This can be configured in case of single-active multi-homing (MH), for example on PE-2 for peer PE-4, with PIM snooping on spoke-SDPs, as follows:

```
On PE-2:
configure {
    redundancy {
        multi-chassis {
            peer 192.0.2.4 {
                admin-state enable
                sync {
                    admin-state enable
                    pim-snooping {
                        spoke-sdps true
                    }
                    tags {
                        sdp 26 {
                            range start 2 end 2 {
                                sync-tag "syncSA"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                            }
                        }
                    }
                }
            }
        }
    }
}
```

On PE-4, MCS is configured for peer PE-2, as follows:

```
On PE-4:
configure {
    redundancy {
        multi-chassis {
            peer 192.0.2.2 {
                admin-state enable
                sync {
                    admin-state enable
                    pim-snooping {
                        spoke-sdps true
                    }
                    tags {
                        sdp 46 {
                            range start 2 end 2 {
                                sync-tag "syncSA"
                            }
                        }
                    }
                }
            }
        }
    }
}
```

When H-8 joins the multicast group, the following entries are in the MCS synchronization database of the PEs. The MCS sync-database on PE-2 shows the PIM snooping entries on the spoke-SDP 26:2 of the single-active MH ESI-24, as follows:

```
[/]
A:admin@PE-2# tools dump redundancy multi-chassis sync-database detail

If no entries are present for an application, no detail will be displayed.

FLAGS LEGEND: ld - local delete; da - delete alarm; pd - pending global delete;
              oal - omcr alarmed; ost - omcr standby

Peer Ip 192.0.2.4


Application pim-snooping-sdp
Sdp-id             Client Key
 SyncTag                          DLen  Flags           timeStamp
  deleteReason code and description                  #ShRec
-------------------------------------------------------------------------------
26:2               Adj 172.16.0.8
 syncSA                           72    -- -- -- --- --- 08/10/2023 23:24:27
  0x0                                                   0
26:2               IfSG SG 192.168.55.2 232.1.1.1
 syncSA                           69    -- -- -- --- --- 08/10/2023 23:24:21
  0x0                                                   0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
The following totals are for:
 peer ip ALL, port/lag/sdp ALL, sync-tag ALL, application ALL
Valid Entries:                   2
Locally Deleted Entries:         0
Locally Deleted Alarmed Entries: 0
Pending Global Delete Entries:   0
Omcr Alarmed Entries:            0
Omcr Standby Entries:            0
Associated Shared Records (ALL): 0
Associated Shared Records (LD):  0
```

The MCS sync-database on PE-4 is similar, with SDP ID 46:2 instead of 26:2.

On PE-4, the MFIB is populated as follows:

```
[/]
A:admin@PE-4# show service id 2 mfib

===============================================================================
Multicast FIB, Service 2
===============================================================================
Source Address  Group Address         Port Id                     Svc Id   Fwd
                                                                         Blk
-------------------------------------------------------------------------------
192.168.55.2    232.1.1.1             sdp:46:2                    Local    Fwd
                                      mpls:192.0.2.1:524280       Local    Fwd
                                      mpls:192.0.2.2:524280       Local    Fwd
                                      mpls:192.0.2.3:524280       Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
```

The PIM snooping group information on PE-4 shows the EVPN-MPLS as incoming interface and the spoke-SDP as outgoing interface, as follows. The split-horizon mechanism does not allow forwarding traffic from the EVPN-MPLS back to the EVPN-MPLS.

```
[/]
A:admin@PE-4# show service id 2 pim-snooping group detail

===============================================================================
PIM Snooping Source Group ipv4
===============================================================================
Group Address      : 232.1.1.1
Source Address     : 192.168.55.2
Up Time            : 0d 00:01:15

Up JP State        : Joined             Up JP Expiry       : 0d 00:00:52
Up JP Rpt          : Not Joined StarG   Up JP Rpt Override : 0d 00:00:00

RPF Neighbor       : 172.16.0.5
Incoming Intf      : EVPN-MPLS
Outgoing Intf List : EVPN-MPLS, SPOKE_SDP:46:2

Forwarded Packets  : 62461              Forwarded Octets   : 93441656
-------------------------------------------------------------------------------
Groups : 1
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Failover

Figure 252: EVPN-MPLS with Multi-homing and PIM Snooping - Receivers H-7 and H-8 Joined shows the multicast traffic flow in the case where both receivers H-7 and H-8 joined multicast group 232.1.1.1 from source 192.168.55.2 and no failures have occurred. For SR OS Release 15.0.R4, and later, MCS need not be configured for faster failover in single-active MH when the ES is non-revertive.

Figure 252: EVPN-MPLS with Multi-homing and PIM Snooping - Receivers H-7 and H-8 Joined

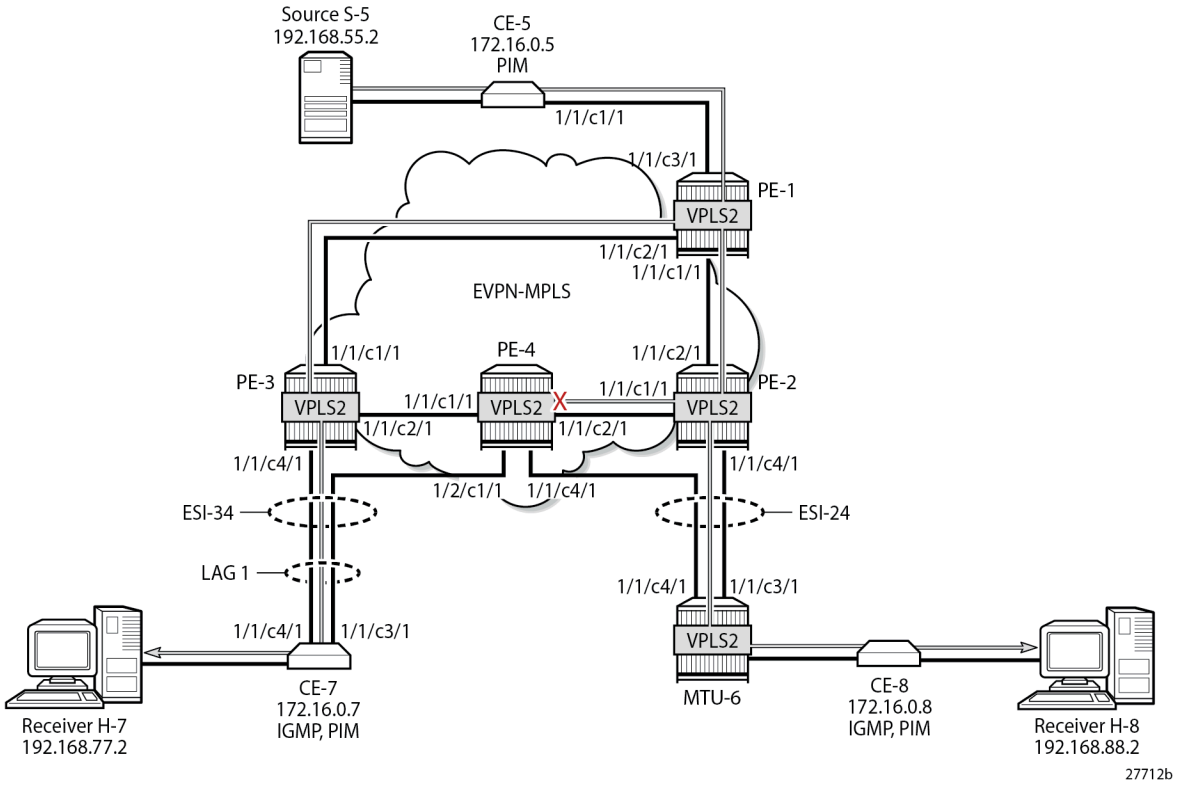

NDF PE-4 has an MFIB table with the required information for a fast failover, as follows:

```
[/]
A:admin@PE-4# show service id 2 mfib

===============================================================================
Multicast FIB, Service 2
===============================================================================
Source Address  Group Address         Port Id                      Svc Id   Fwd
                                                                          Blk
-------------------------------------------------------------------------------
192.168.55.2    232.1.1.1             sap:lag-1:2                  Local    Fwd
                                      sdp:46:2                     Local    Fwd
                                      mpls:192.0.2.1:524280        Local    Fwd
                                      mpls:192.0.2.2:524280        Local    Fwd
                                      mpls:192.0.2.3:524280        Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

In SR OS Release 15.0.R4, and later, data-driven PIM state synchronization ensures that NDF PE-4 has the following PIM snooping information for group 232.1.1.1.

```
[/]
A:admin@PE-4# show service id 2 pim-snooping group detail

===============================================================================
PIM Snooping Source Group ipv4
===============================================================================
Group Address      : 232.1.1.1
Source Address     : 192.168.55.2
Up Time            : 0d 00:03:12

Up JP State        : Joined             Up JP Expiry       : 0d 00:00:59
Up JP Rpt          : Not Joined StarG   Up JP Rpt Override : 0d 00:00:00

RPF Neighbor       : 172.16.0.5
Incoming Intf      : EVPN-MPLS
Outgoing Intf List : EVPN-MPLS, SAP:lag-1:2, SPOKE_SDP:46:2

Forwarded Packets  : 158148             Forwarded Octets   : 236589408
-------------------------------------------------------------------------------
Groups : 1
===============================================================================
```

The following failures are introduced to force a failover from PE-2 to PE-4 and from PE-3 to PE-4. On MTU-6, SDP 62 is disabled, as follows:

```
configure {
    service {
        sdp 62 {
            admin-state disable
        }
    }
}
```

On CE-7, port 1/1/c4/1 toward PE-3 is disabled, as follows:

```
configure {
    port 1/1/c4/1 {
        admin-state disable
    }
}
```

Log 99 on PE-3 shows that the DF state in ESI-34_2 changed to false:

```
172 2023/08/10 23:18:51.662 CEST MINOR: SVCMGR #2094 Base
"Ethernet Segment:ESI-34_2, EVI:2, Designated Forwarding state changed to:false"
```

PE-4 becomes the DF for both ESs, as follows:

```
[/]
A:admin@PE-4# show service id 2 ethernet-segment

===============================================================================
SAP Ethernet-Segment Information
===============================================================================
SAP                  Eth-Seg                         Status
-------------------------------------------------------------------------------
lag-1:2              ESI-34_2                        DF
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================

===============================================================================
SDP Ethernet-Segment Information
===============================================================================
SDP                  Eth-Seg                           Status
-------------------------------------------------------------------------------
46:2                 ESI-24                            DF
===============================================================================
No vxlan instance entries
```

Figure 253: EVPN-MPLS with Multi-homing and PIM Snooping - Multicast Flow after Failover shows the traffic flow after failover to new DF PE-4.

Figure 253: EVPN-MPLS with Multi-homing and PIM Snooping - Multicast Flow after Failover

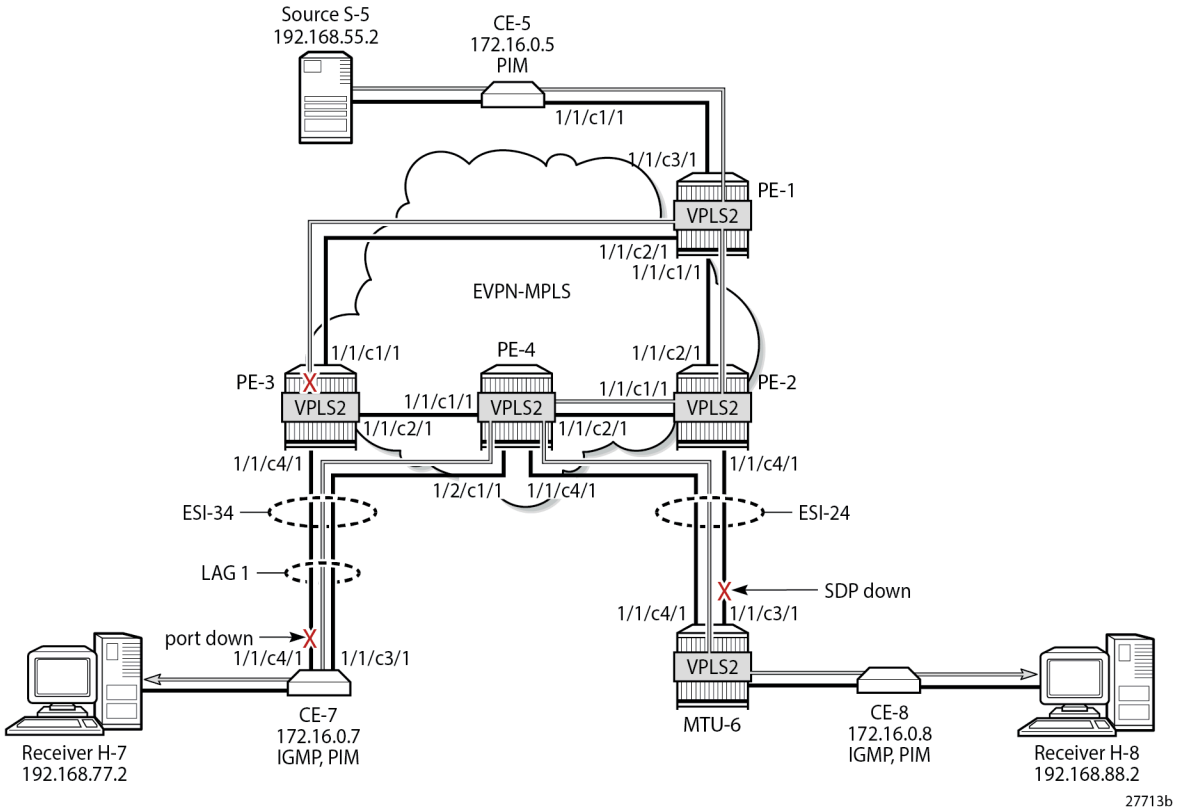

PE-1 receives the multicast on port 1/1/c3/1 and forwards it on port 1/1/c1/1 to PE-2, and on port 1/1/c2/1 to PE-3, as follows:

```
[/]
A:admin@PE-1# show port 1/1/c1/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c1/1                                  34                           3412
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                                    26284                         39855201
===============================================================================

[/]
A:admin@PE-1# show port 1/1/c2/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c2/1                                28                             2946
                                    26281                         39854977
===============================================================================

[/]
A:admin@PE-1# show port 1/1/c3/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c3/1                            26255                         39486092
                                        2                              152
===============================================================================
```

PE-2 receives the multicast stream from PE-1 on port 1/1/c2/1 and forwards it to port 1/1/c1/1 to PE-4; it does not forward to port 1/1/c4/1 because SDP 26 is down, as follows:

```
[/]
A:admin@PE-2# show port 1/1/c1/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c1/1                                36                             3664
                                    26428                         40075241
===============================================================================

[/]
A:admin@PE-2# show port 1/1/c2/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c2/1                            26430                         40075436
                                       34                             3412
===============================================================================

[/]
A:admin@PE-2# show port 1/1/c3/1 statistics

[/]
A:admin@PE-2# show port 1/1/c4/1 statistics
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c4/1                                25                           2686
                                        27                           2921
===============================================================================
```

PE-4 receives the multicast traffic on port 1/1/c2/1 and forwards it on port 1/1/c4/1 toward MTU-6, and on port 1/2/c1/1 to CE-7, as follows:

```
[/]
A:admin@PE-4# show port 1/1/c1/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c1/1                                28                           2948
                                        29                           3056
===============================================================================

[/]
A:admin@PE-4# show port 1/1/c2/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c2/1                             26274                       39841553
                                        30                           3088
===============================================================================

[/]
A:admin@PE-4# show port 1/1/c4/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c4/1                                26                           2814
                                     26270                       40051218
===============================================================================

[/]
A:admin@PE-4# show port 1/2/c1/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/2/c1/1                                33                           4172
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                                        26278                       39475224
===============================================================================
```

MTU-6 forwards the traffic to CE-8, which forwards it to H-8. CE-7 forwards the traffic to H-7. PE-3 drops the multicast traffic because LAG-1 is down because of the failure that was introduced at CE-7 (port disabled).

Conclusion

PIM snooping in EVPN-MPLS services results in a more efficient use of network resources because multicast traffic no longer needs to be flooded. PIM snooping can be used in EVPN-MPLS services with all-active and single-active multi-homing with data-driven PIM state synchronization. Alternatively, MCS synchronization of the PIM snooping state on SAPs and spoke-SDPs is supported with single-active MH.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

PIM Snooping for IPv4 in PBB-EVPN Services

This chapter describes PIM Snooping for IPv4 in PBB-EVPN Services.

Topics in this chapter include:

Applicability

This chapter was initially written based on SR OS Release 15.0.R5, but the CLI in the current edition corresponds to SR OS Release 23.7.R1. Protocol Independent Multicast (PIM) snooping for IPv4 is supported in Provider Backbone Bridging - Ethernet Virtual Private Network (PBB-EVPN) services in SR OS Release 15.0.R1, and later. PIM snooping in single-active multi-homing (MH) mode without Ethernet Segment Identifier (ESI) label is supported in SR OS Release 15.0.R1, and later, whereas PIM snooping in single-active MH mode with ESI label is supported in SR OS Release 15.0.R4, and later. PIM snooping for IPv4 in all-active MH mode is supported in SR OS Release 15.0.R4, and later. Data-driven PIM state synchronization is supported in SR OS Release 15.0.R4, and later.

Overview

PBB-EVPN services have EVPN-MPLS enabled in the B-VPLS. PIM snooping in PBB-EVPN I-VPLS provides the following:

- PIM snooping in SAPs and SDP-bindings: PIM messages received from SAPs, SDP-bindings, or the B-VPLS are forwarded to SAPs or SDP-bindings according to the PIM snooping.
- Multicast flooding between I-VPLS and B-VPLS is the same for a PBB-EVPN B-VPLS as for a B-VPLS without EVPN. The first PIM join message received over the local B-VPLS from a B-VPLS SAP/SDP-binding or EVPN endpoint results in adding the B-VPLS SAP/SDP-binding or EVPN interface into the Multicast Forwarding Information Base (MFIB) associated with the I-VPLS context. Multicast traffic is flooded throughout the B-VPLS on a per-ISID single tree.
- When the PIM router is connected to a remote I-VPLS instance over the B-VPLS infrastructure, its location is identified by the B-VPLS SAP/SDP-binding or by the set of all EVPN endpoints on which PIM hellos are received. The location is also identified by the source BMAC address in the PBB header for the PIM hello message, which is the BMAC address associated with the B-VPLS instance on the remote PBB PE.
- The set of all EVPN endpoints in the B-VPLS is treated as a single PIM interface.
  - Hello and join/prune messages from I-VPLS SAPs/SDP-bindings are always sent to all B-VPLS PBB-EVPN destinations.
  - When a hello message is received from one B-VPLS PBB-EVPN destination PIM neighbor, the single interface representing all B-VPLS PBB-EVPN destinations will have that PIM neighbor.
  - All individual B-VPLS PBB-EVPN destinations appear in the MFIB, but the information for each B-VPLS PBB-EVPN destination entry is identical.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

  - The EVPN split-horizon logic ensures that IP multicast traffic and PIM messages received on a PBB-EVPN endpoint are not forwarded back to other PBB-EVPN endpoints.

- When a point-to-multipoint (P2MP) mLDP provider tunnel is configured in the B-VPLS, the provider tunnel only works for the default multicast list. Ingress Replication (IR) is used for the per-ISID MFIB trees. ISID policies can be configured to specify ISID ranges that will use the default multicast list. ISID policies can help reduce the per-ISID MFIB resources used.
- PIM snooping for IPv4 within a PBB-EVPN I-VPLS is supported with single-active MH and with all-active MH in the associated I-VPLS.
- Data-driven PIM state synchronization between remote peers in an all-active MH Ethernet Segment (ES) is supported.
- Multi-Chassis Synchronization (MCS) of PIM snooping state on SAPs and spoke-SDPs is supported in active/standby scenarios.

The following command enables PIM snooping in an I-VPLS:

```
configure {
    service {
        vpls "I-VPLS 1" {
            pim-snooping { }
        }
```

The default PIM snooping mode is proxy mode, which implies that the PE will terminate the PIM join/prune messages and generate its own PIM join/prune messages with the same (S,G). The advantage is that the number of PIM messages to be sent can be reduced: regardless of the number of PIM join messages received for a certain (S,G), the node only needs to send one PIM join message toward the source. PIM snooping can also use snooping mode based on the information in the received PIM hello messages; in snooping mode, the PE does not modify the PIM messages.

Configuration

Figure 254: Example Topology for PBB-EVPN without MH shows the example topology with source S-5 and receivers H-6, H-7, and H-8 attached to CEs that are connected to PEs. On the PEs, B-VPLS 100 is configured and I-VPLS 1 is associated with it. B-VPLS 100 has EVPN-MPLS enabled. An mLDP P2MP provider tunnel is used to distribute multicast traffic from PE-1 to the other PEs.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 254: Example Topology for PBB-EVPN without MH

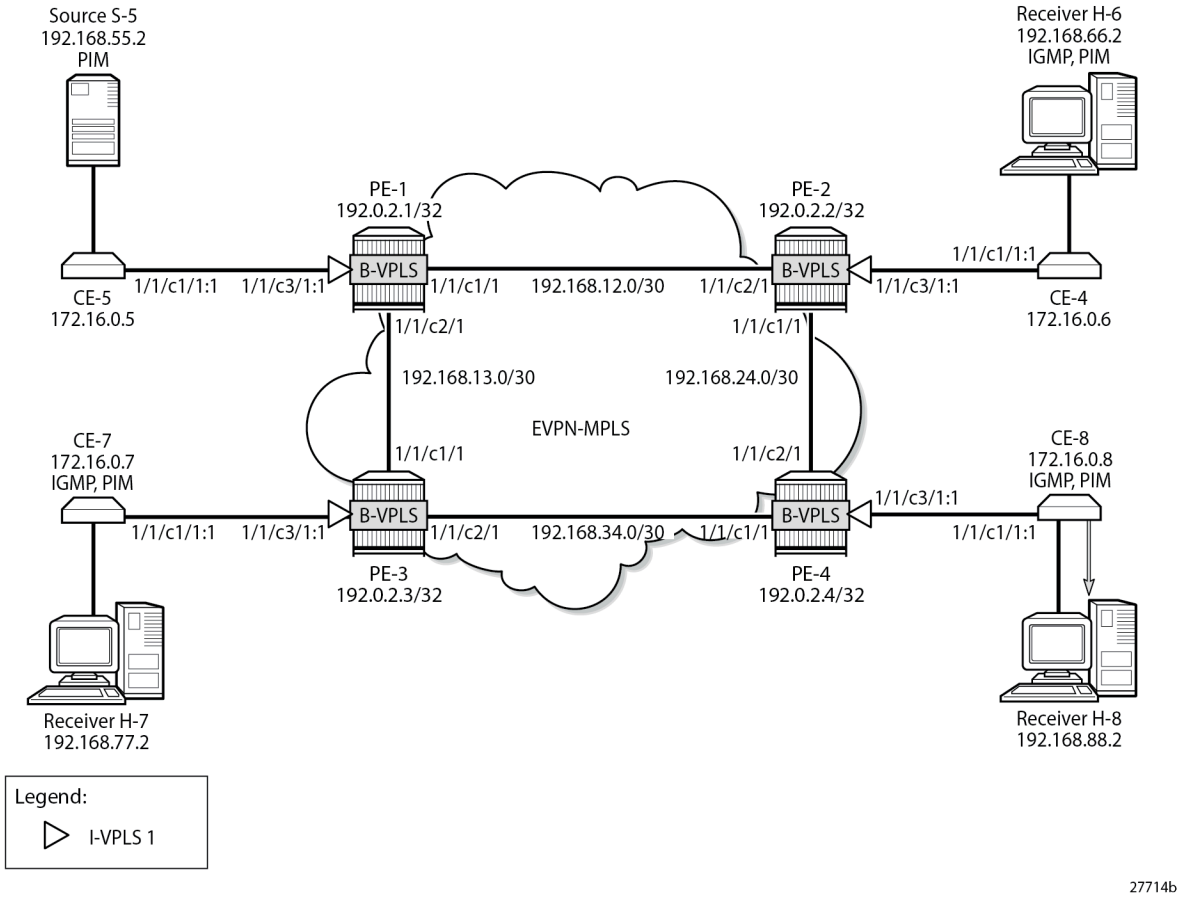

The initial configuration includes:

- Cards, MDAs, ports
- Router interfaces
- IS-IS enabled on the PEs (alternatively, OSPF can be used)
- LDP enabled on the PEs

BGP is configured on the PEs with address family EVPN, and PE-2 is configured as route reflector (RR). The BGP configuration on PE-2 is as follows:

```
On PE-2:
configure {
    router "Base" {
        bgp {
            rapid-withdrawal true
            ebgp-default-reject-policy {
                import false
                export false
            }
            rapid-update {
                evpn true
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            group "INTERNAL" {
                type internal
                family {
                    evpn true
                }
                cluster {
                    cluster-id 192.0.2.2
                }
            }
            neighbor "192.0.2.1" {
                group "INTERNAL"
            }
            neighbor "192.0.2.3" {
                group "INTERNAL"
            }
            neighbor "192.0.2.4" {
                group "INTERNAL"
            }
        }
```

PBB-EVPN without MH – No PIM Snooping

B-VPLS 100 is configured with EVPN-MPLS enabled on all PEs. Multicast LDP is configured in B-VPLS 100 with PE-1 as the P2MP tunnel root node (**root-and-leaf**) and the other PEs as leaf nodes (**no root-and-leaf** is default). An (optional) ISID policy defines that the default multicast tree -which is used by the P2MP mLDP tunnel- is used for ISIDs 1 and 2 (range 1 to 2). The configuration of B-VPLS 100 on PE-1 is as follows:

```
On PE-1:
configure {
    service {
        vpls "B-VPLS 100" {
            admin-state enable
            description "B-VPLS 100"
            service-id 100
            customer "1"
            service-mtu 2000
            pbb-type b-vpls
            pbb {
                source-bmac {
                    address 00:00:00:00:00:01
                    use-es-bmac-lsb true
                }
            }
            bgp 1 { }
            bgp-evpn {
                evi 100
                mpls 1 {
                    admin-state enable
                    split-horizon-group "CORE"
                    ingress-replication-bum-label true
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            split-horizon-group "CORE" { }
            provider-tunnel {
                inclusive {
                    admin-state enable
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    owner bgp-evpn-mpls
                    root-and-leaf true
                    mldp
                }
            }
            isid-policy {
                entry 1 {
                    advertise-local false
                    use-def-mcast true
                    range {
                        start 1
                        end 2
                    }
                }
            }
        }
```

The configuration of B-VPLS on the other PEs is similar, but without the root-and-leaf option.

In B-VPLS 100 on root node PE-1, the following mLDP provider tunnel is created with Provider Multicast Service Interface (PMSI) owner bgpEvpnMpls. PE-1 is configured as root-and-leaf node.

```
[/]
A:admin@PE-1# show service id 100 provider-tunnel

===============================================================================
Service Provider Tunnel Information
===============================================================================
Type               : inclusive          Root and Leaf      : enabled
Admin State        : enabled            Data Delay Intvl   : 15 secs
PMSI Type          : ldp                LSP Template       :
Remain Delay Intvl : 0 secs             LSP Name used      : 8193
PMSI Owner         : bgpEvpnMpls
Oper State         : up                 Root Bind Id       : 32767
-------------------------------------------------------------------------------
Type               : selective          Wildcard SPMSI     : disabled
Admin State        : disabled           Data Delay Intvl   : 3 secs
PMSI Type          : none               Max P2MP SPMSI     : 10
PMSI Owner         : none
===============================================================================
```

When the B-VPLS is created, I-VPLS 1 can be associated with it, as follows:

```
On PE-1:
configure {
    service {
        vpls "I-VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS 100" {
                    isid 1
                }
            }
            sap 1/1/c3/1:1 { }
        }
}
```

The configuration of I-VPLS 1 on the other PEs is identical.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

CE-6, CE-7, and CE-8 have IGMP enabled on the interface toward the receiver and PIM enabled on all interfaces. Source-specific multicast is used in this example. The configuration on CE-8 is as follows:

```
On CE-8:
configure {
    router "Base" {
        interface "int-CE-8-PE-4" {
            port 1/1/c1/1:1
            ipv4 {
                primary {
                    address 172.16.0.8
                    prefix-length 16
                }
            }
        }
        interface "int-CE-8-H-8" {
            port 1/1/c2/1
            ipv4 {
                primary {
                    address 192.168.88.1
                    prefix-length 24
                }
            }
        }
        interface "system" {
            ipv4 {
                primary {
                    address 192.0.2.8
                    prefix-length 32
                }
            }
        }
        static-routes {
            route 192.168.55.0/30 route-type unicast {
                next-hop "172.16.0.5" {
                    admin-state enable
                }
            }
        }
        igmp {
            interface "int-CE-8-H-8" { }
        }
        pim
            apply-to all
        }
    }
```

The static route is required on the receiving CEs for the PIM join/prune messages to reach the multicast source S-5 with IP address 192.168.55.2; only IP subnet 172.16.0.0/16 can be reached via the VPLS.

CE-5 has PIM enabled and static routes configured to reach the receiving hosts, as follows:

```
On CE-5:
configure {
    router "Base" {
        interface "int-CE-5-PE-1" {
            port 1/1/c1/1:1
            ipv4 {
                primary {
                    address 172.16.0.5
                    prefix-length 16
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
        interface "int-CE-5-S-5" {
            port 1/1/c3/1
            ipv4 {
                primary {
                    address 192.168.55.1
                    prefix-length 30
                }
            }
        }
        interface "system" {
            ipv4 {
                primary {
                    address 192.0.2.5
                    prefix-length 32
                }
            }
        }
        static-routes {
            route 192.168.66.0/24 route-type unicast {
                next-hop "172.16.0.6" {
                    admin-state enable
                }
            }
            route 192.168.77.0/24 route-type unicast {
                next-hop "172.16.0.7" {
                    admin-state enable
                }
            }
            route 192.168.88.0/24 route-type unicast {
                next-hop "172.16.0.8" {
                    admin-state enable
                }
            }
        }
        pim {
            apply-to all
        }
    }
```

When receiver H-8 sends an IGMP report to join multicast group (S,G), CE-8 sends a PIM join message to CE-5. This PIM join message is flooded by the PEs. When CE-5 receives the PIM join message, it forwards the multicast stream to receiver H-8. PIM snooping is disabled by default and the MFIB on each of the PEs remains empty, so the multicast stream is not only sent to CE-8, but also to CE-6 and CE-7. CE-6 and CE-7 drop this stream when no receiver is active, while CE-8 forwards the multicast stream to receiver H-8, as shown in Figure 255: Multicast Stream to Receiver H-8 with PIM Snooping Disabled.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 255: Multicast Stream to Receiver H-8 with PIM Snooping Disabled

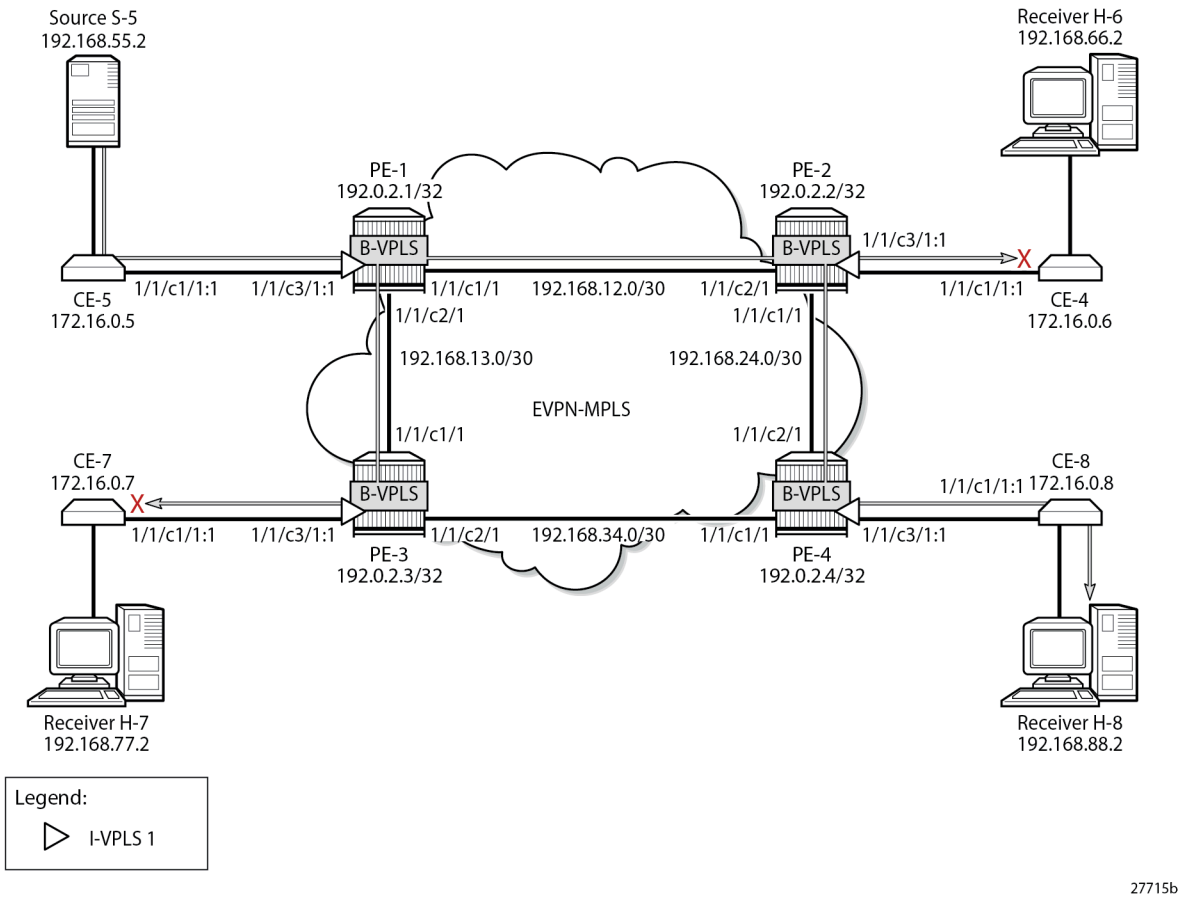

PBB-EVPN without MH – PIM Snooping for IPv4 Enabled

PIM snooping for IPv4 is enabled in I-VPLS 1 on all PEs as follows:

```
On all PEs:
configure {
    service {
        vpls "I-VPLS 1" {
            pim-snooping { }
        }
```

When PIM snooping for IPv4 is enabled, the PEs only forward the multicast traffic to those CEs that have sent PIM join messages for that multicast group. This implies that PE-2 and PE-3 do not forward traffic to the CEs; only PE-4 forwards traffic toward CE-8 and CE-8 forwards to receiver H-8, as shown in Figure 256: Multicast Stream to Receiver H-8 with PIM Snooping Enabled.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 256: Multicast Stream to Receiver H-8 with PIM Snooping Enabled

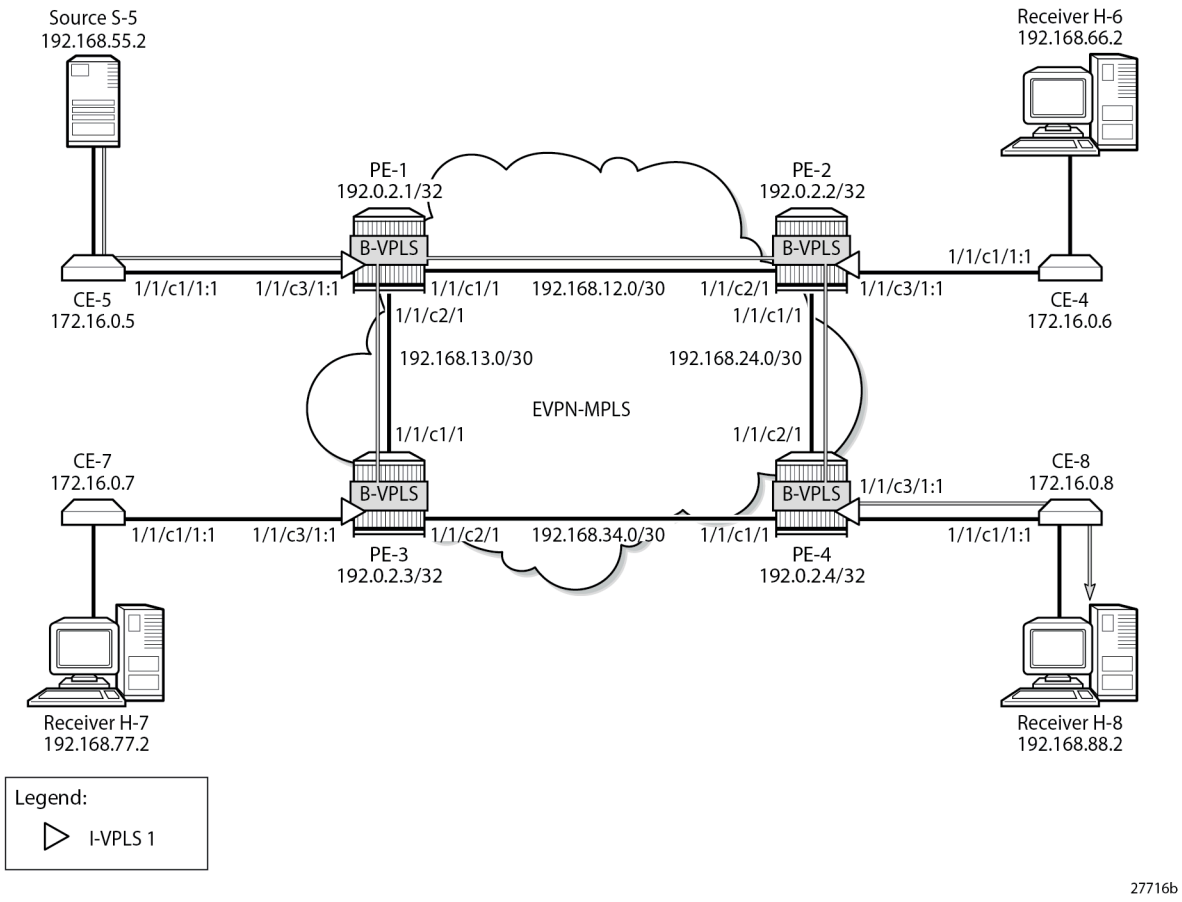

When PIM snooping for IPv4 is enabled, PE-1 has the following two PIM snooping ports: the SAP toward the source and the backbone b-EVPN-MPLS interface, which is treated as one entity for all PBB-EVPN destinations.

```
[/]
A:admin@PE-1# show service id 1 pim-snooping port

===============================================================================
PIM Snooping Port ipv4
===============================================================================
Port Id                                                     Opr   PW Fwding
-------------------------------------------------------------------------------
SAP:1/1/c3/1:1                                              Up    Actv
b-EVPN-MPLS                                                 Up    Actv
===============================================================================
```

PE-1 has the following PIM snooping neighbors: CE-5 with IP address 172.16.0.5 is attached via SAP 1/1/c3/1:1, and the other CEs are attached to the b-EVPN-MPLS. Even though this b-EVPN-MPLS is treated as one entity, individual entries are shown for each B-VPLS PBB-EVPN destination, as follows:

```
[/]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
A:admin@PE-1# show service id 1 pim-snooping neighbor

===============================================================================
PIM Snooping Neighbors ipv4
===============================================================================
Port Id                 Nbr DR Prty     Up Time       Expiry Time     Hold Time
  Nbr Address
-------------------------------------------------------------------------------
SAP:1/1/c3/1:1          1               0d 00:01:04   0d 00:01:41     105
  172.16.0.5
b-EVPN-MPLS             1               0d 00:01:08   0d 00:01:37     105
  172.16.0.6
b-EVPN-MPLS             1               0d 00:00:53   0d 00:01:22     105
  172.16.0.7
b-EVPN-MPLS             1               0d 00:01:09   0d 00:01:36     105
  172.16.0.8
-------------------------------------------------------------------------------
Neighbors : 4
===============================================================================
```

Receiver H-8 joins the multicast stream and the PIM group with group address 232.1.1.1, and source address 192.168.55.2 is shown on CE-8 with incoming interface toward PE-4 and outgoing interface toward H-8. The Reverse Path Forwarding (RPF) neighbor is CE-5 with IP address 172.16.0.5, as follows:

```
[/]
A:admin@CE-8# show router pim group detail

===============================================================================
PIM Source Group ipv4
===============================================================================
Group Address      : 232.1.1.1
Source Address     : 192.168.55.2
RP Address         : 0
Advt Router        :
Flags              :                    Type              : (S,G)
Mode               : sparse
MRIB Next Hop      : 172.16.0.5
MRIB Src Flags     : remote
Keepalive Timer    : Not Running
Up Time            : 0d 00:00:29        Resolved By       : rtable-u

Up JP State        : Joined             Up JP Expiry      : 0d 00:00:30
Up JP Rpt          : Not Joined StarG   Up JP Rpt Override : 0d 00:00:00

Register State     : No Info
Reg From Anycast RP: No

Rpf Neighbor       : 172.16.0.5
Incoming Intf      : int-CE-8-PE-4
Outgoing Intf List : int-CE-8-H-8

Curr Fwding Rate   : 9745.632 kbps
Forwarded Packets  : 23848              Discarded Packets  : 0
Forwarded Octets   : 35342736           RPF Mismatches     : 0
Spt threshold      : 0 kbps             ECMP opt threshold : 7
Admin bandwidth    : 1 kbps
-------------------------------------------------------------------------------
Groups : 1
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

With PIM snooping for IPv4 enabled, and after receiving a PIM join message for multicast (192.168.55.2, 232.1.1.1), the MFIB on PE-1 has an entry for group address 232.1.1.1 and source address 192.168.55.2, as follows. The local SAP connects to CE-5; the other port IDs correspond to the b-EVPN-MPLS interface.

```
[/]
A:admin@PE-1# show service id 1 mfib

===============================================================================
Multicast FIB, Service 1
===============================================================================
Source Address  Group Address         Port Id                     Svc Id   Fwd
                                                                        Blk
-------------------------------------------------------------------------------
192.168.55.2    232.1.1.1             sap:1/1/c3/1:1              Local    Fwd
                                      b-mpls:192.0.2.2:524282     100      Fwd
                                      b-mpls:192.0.2.3:524282     100      Fwd
                                      b-mpls:192.0.2.4:524282     100      Fwd
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
```

The MFIB on PE-4 is similar, with a local SAP connecting to CE-8 and three b-eMpls port IDs for each of the PE peers. In contrast, the MFIBs on PE-2 and PE-3 are empty, because no multicast traffic needs to be forwarded to the attached CEs, as follows:

```
[/]
A:admin@PE-2# show service id 1 mfib

===============================================================================
Multicast FIB, Service 1
===============================================================================
Source Address  Group Address         Port Id                     Svc Id   Fwd
                                                                        Blk
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Number of entries: 0
===============================================================================
```

The following MFIB statistics on PE-1 show the number of matched packets and matched octets for group address 232.1.1.1 and source address 192.168.55.2:

```
[/]
A:admin@PE-1# show service id 1 mfib statistics

===============================================================================
Multicast FIB Statistics, Service 1
===============================================================================
Source Address  Group Address       Matched Pkts          Matched Octets
                                                          Forwarding Rate
-------------------------------------------------------------------------------
192.168.55.2    232.1.1.1           82582                 123873000
                                                          61.864 kbps
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
```

The following shows that PE-2 receives the multicast packets on port 1/1/c2/1 and forwards them to PE-4 on port 1/1/c1/1. With PIM snooping for IPv4 enabled, PE-2 does not forward the traffic to CE-6 on port 1/1/c3/1 because no PIM join message was received from CE-6. Besides the multicast traffic, some signaling

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

messages (such as PIM, IS-IS, and so on) are sent on the ports, which explains why all counters have non-zero values.

```
[/]
A:admin@PE-2# show port 1/1/c1/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c1/1                                33                           3392
                                     29361                       45048369
===============================================================================

[/]
A:admin@PE-2# show port 1/1/c2/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c2/1                             29363                       45048520
                                        36                           3663
===============================================================================

[/]
A:admin@PE-2# show port 1/1/c3/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c3/1                                1                             76
                                        4                            304
===============================================================================
```

The following PIM snooping group with group address 232.1.1.1 and source address 192.168.55.2 is shown on PE-1. The incoming interface is the SAP toward CE-5 and the outgoing interface is the b-EVPN-MPLS interface. A single b-EVPN-MPLS interface is shown in the outgoing interface list, regardless of the B-VPLS PBB-EVPN destination. The split-horizon mechanism ensures that all traffic from the incoming interface SAP 1/1/c3/1:1 is only forwarded on the b-EVPN-MPLS interface, not sent back on the SAP.

```
[/]
A:admin@PE-1# show service id 1 pim-snooping group detail

===============================================================================
PIM Snooping Source Group ipv4
===============================================================================
Group Address      : 232.1.1.1
Source Address     : 192.168.55.2
Up Time            : 0d 00:01:41

Up JP State        : Joined            Up JP Expiry       : 0d 00:00:19
Up JP Rpt          : Not Joined StarG  Up JP Rpt Override : 0d 00:00:00

RPF Neighbor       : 172.16.0.5
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Incoming Intf      : SAP:1/1/c3/1:1
Outgoing Intf List : b-EVPN-MPLS, SAP:1/1/c3/1:1

Forwarded Packets  : 82582             Forwarded Octets   : 123873000
-------------------------------------------------------------------------------
Groups : 1
===============================================================================
```

On PE-2 and PE-3, there are no PIM snooping groups.

On PE-4, the PIM snooping group with group address 232.1.1.1 and source address 192.168.55.2 has the b-EVPN-MPLS interface as incoming interface and SAP 1/1/c3/1:1 toward CE-8 as outgoing interface, as follows. The split-horizon mechanism ensures that traffic received from the b-EVPN-MPLS interface is not forwarded on the b-EVPN-MPLS interface to the other PEs, so it is only forwarded on the SAP 1/1/c3/1:1 toward CE-8.

```
[/]
A:admin@PE-4# show service id 1 pim-snooping group 232.1.1.1 detail

===============================================================================
PIM Snooping Source Group ipv4
===============================================================================
Group Address      : 232.1.1.1
Source Address     : 192.168.55.2
Up Time            : 0d 00:01:46

Up JP State        : Joined            Up JP Expiry       : 0d 00:00:13
Up JP Rpt          : Not Joined StarG  Up JP Rpt Override : 0d 00:00:00

RPF Neighbor       : 172.16.0.5
Incoming Intf      : b-EVPN-MPLS
Outgoing Intf List : b-EVPN-MPLS, SAP:1/1/c3/1:1

Forwarded Packets  : 83498             Forwarded Octets   : 126415972
-------------------------------------------------------------------------------
Groups : 1
===============================================================================
```

PBB-EVPN with MH – No PIM Snooping

Figure 257: Example Topology for PBB-EVPN with MH shows the example topology with CE-7 attached to MTU-9, which is connected to both PE-3 and PE-4 via LAG lag-1. Virtual ES (vES) ESI-34_2 is configured in all-active MH mode using lag-1 for dot1q value 2. MTU-6 is connected to PE-2 and PE-4 with SDPs. These SDPs are associated with a single-active MH ES ESI-24.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 257: Example Topology for PBB-EVPN with MH

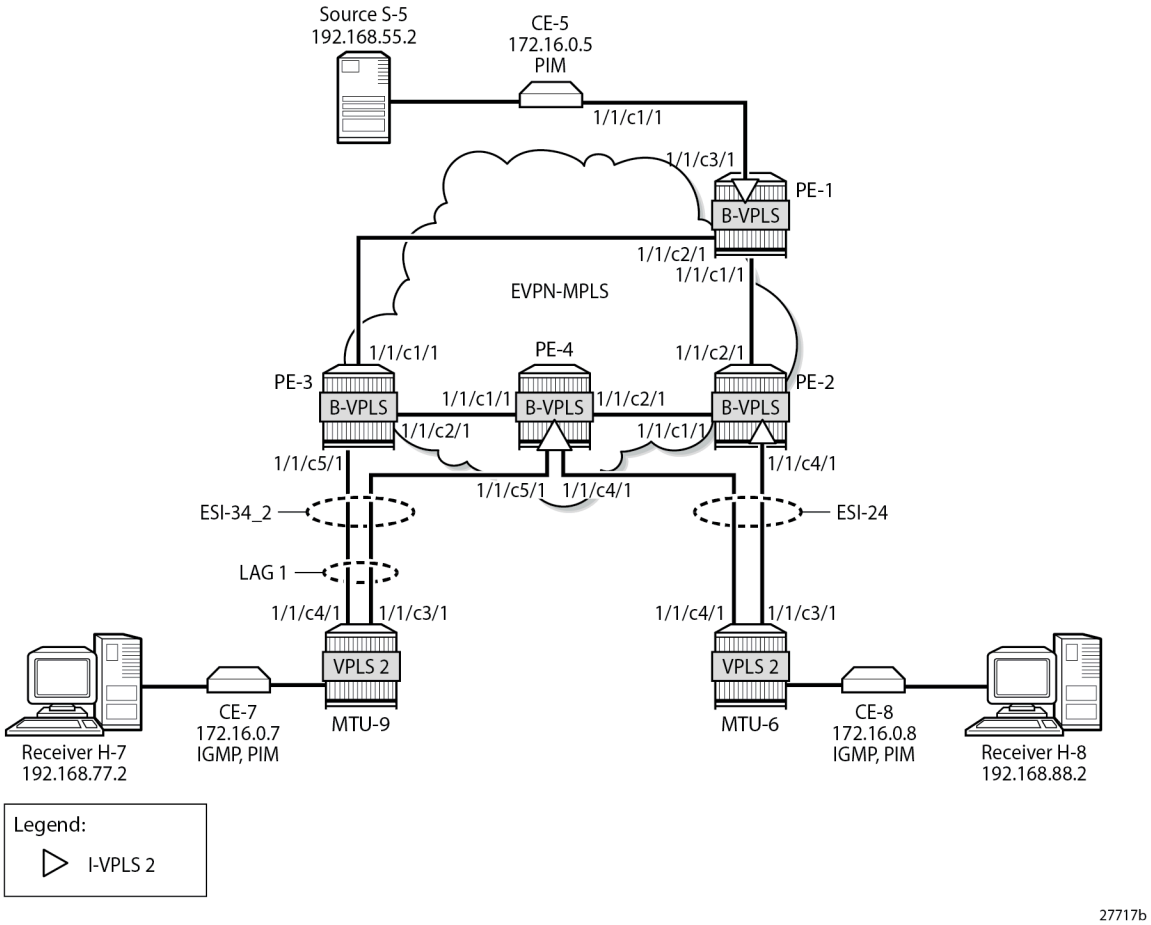

The configuration of I-VPLS 2 is similar to the preceding configuration of I-VPLS 1 on all PEs.

On PE-2, PE-3, and PE-4, one or more ESs are configured. The service configuration on PE-2 is as follows. An SDP is configured toward MTU-6 that is associated with a single-active MH ES ESI-24, that is non-restrictive -after failover, it does not restore to the initial designated forwarder (DF) if available again; see chapter Preference-based and Non-revertive EVPN DF Election. The manually configured preference is 200 on PE-2, which is higher than preference 50 at PE-3, so PE-2 is the DF when no failover has occurred. Spoke-SDP 26:2 is associated with I-VPLS 2. The B-VPLS 100 remains unchanged and is not repeated here.

```
On PE-2:
configure {
    service {
        sdp 26 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.6
            }
        }
        system {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            bgp {
                evpn {
                    ethernet-segment "ESI-24" {
                        admin-state enable
                        esi 0x01000000002400000001
                        multi-homing-mode single-active
                        df-election {
                            es-activation-timer 3
                            service-carving-mode manual
                            manual {
                                preference {
                                    mode non-revertive
                                    value 200
                                }
                            }
                        }
                        association {
                            sdp 26 { }
                        }
                        pbb {
                            source-bmac-lsb 0x2402
                        }
                    }
                }
            }
        }
        vpls "I-VPLS 2" {
            admin-state enable
            service-id 2
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS 100" {
                    isid 2
                }
            }
            spoke-sdp 26:2 { }
        }
        vpls "B-VPLS 100" {
            pbb {
                source-bmac {
                    use-es-bmac-lsb true
                }
            }
        }
    }
```

On PE-4, lag-1 is configured in access mode with dot1q encapsulation on the port to MTU-9, as follows. The LAG configuration is similar on PE-3.

```
On PE-4:
configure {
    lag "lag-1" {
        admin-state enable
        encap-type dot1q
        mode access
        lacp {
            mode active
            system-id 00:00:00:00:01:34
            administrative-key 1
        }
        port 1/1/c5/1 { }
    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Single-active ES ESI-24 is configured on PE-4, together with a virtual ES ESI-34_2, which is an all-active MH virtual ES that applies to lag-1 for I-VPLS 2 only (**q-tag-range 2**); see chapter Virtual Ethernet Segments. The preference for the DF election is configured manually to a value of 50 (which is lower than preference 200 on the remote peer in the ES). I-VPLS 2 has a SAP and a spoke-SDP configured. The service configuration on PE-4 is as follows:

```
On PE-4:
configure {
    service {
        sdp 46 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.6
            }
        }
        system {
            bgp {
                evpn {
                    ethernet-segment "ESI-24" {
                        admin-state enable
                        esi 0x01000000002400000001
                        multi-homing-mode single-active
                        df-election {
                            es-activation-timer 3
                            service-carving-mode manual
                            manual {
                                preference {
                                    mode non-revertive
                                    value 50
                                }
                            }
                        }
                        association {
                            sdp 46 { }
                        }
                        pbb {
                            source-bmac-lsb 0x2404
                        }
                    }
                    ethernet-segment "ESI-34_2" {
                        admin-state enable
                        type virtual
                        esi 0x01000000003402000001
                        multi-homing-mode all-active
                        df-election {
                            es-activation-timer 3
                            service-carving-mode manual
                            manual {
                                preference {
                                    mode non-revertive
                                    value 50
                                }
                            }
                        }
                        association {
                            lag "lag-1" {
                                virtual-ranges {
                                    dot1q {
                                        q-tag 2 {
                                            end 2
                                        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                                    }
                                }
                            }
                        }
                        pbb {
                            source-bmac-lsb 0x3434
                        }
                    }
                }
            }
        }
        vpls "I-VPLS 2" {
            admin-state enable
            service-id 2
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS 100" {
                    isid 2
                }
            }
            spoke-sdp 46:2 { }
            sap lag-1:2 { }
        }
        vpls "B-VPLS 100" {
            pbb {
                source-bmac {
                    use-es-bmac-lsb true
                }
            }
        }
    }
```

On PE-4, the source BMAC in the all-active MH ESI-34_2 is identical to the source BMAC on remote peer PE-3, but in the single-active MH ESI-24, the source BMAC must be different. Remote PEs might send traffic to the NDF PE based on the shared source BMAC, which is fine for all-active MH ESs, but not for single-active MH ESs.

The following service configuration on PE-3 includes an all-active virtual ES ESI-34_2 with preference 200, which makes PE-3 the DF for ESI-34_2 when no failover has occurred. After failover, PE-4 becomes DF, and it does not revert to PE-3 when available.

```
On PE-3:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "ESI-34_2" {
                        admin-state enable
                        type virtual
                        esi 0x01000000003402000001
                        multi-homing-mode all-active
                        df-election {
                            es-activation-timer 3
                            service-carving-mode manual
                            manual {
                                preference {
                                    mode non-revertive
                                    value 200
                                }
                            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        }
                        association {
                            lag "lag-1" {
                                virtual-ranges {
                                    dot1q {
                                        q-tag 2 {
                                            end 2
                                        }
                                    }
                                }
                            }
                        }
                        pbb {
                            source-bmac-lsb 0x3434
                        }
                    }
                }
            }
        }
        vpls "I-VPLS 2" {
            admin-state enable
            service-id 2
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS 100" {
                    isid 2
                }
            }
            sap lag-1:2 { }
        }
        vpls "B-VPLS 100" {
            pbb {
                source-bmac {
                    use-es-bmac-lsb true
                }
            }
        }
    }
```

The following is the service configuration on MTU-6:

```
On MTU-6:
configure {
    service {
        sdp 62 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.2
            }
        }
        sdp 64 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.4
            }
        }
        vpls "VPLS 2" {
            admin-state enable
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            service-id 2
            customer "1"
            endpoint "x" { }
            spoke-sdp 62:2 {
                endpoint {
                    name "x"
                }
                stp {
                    admin-state disable
                }
            }
            spoke-sdp 64:2 {
                endpoint {
                    name "x"
                }
                stp {
                    admin-state disable
                }
            }
            sap 1/2/c1/1:2 { }
        }
    }
```

The following is the LAG configuration on MTU-9:

```
On MTU-9:
configure {
    lag "lag-1" {
        admin-state enable
        encap-type dot1q
        mode access
        lacp {
            mode active
            administrative-key 32768
        }
        port 1/1/c3/1 { }
        port 1/1/c4/1 { }
    }
```

The configuration of VPLS 2 on MTU-9 is as follows:

```
On MTU-9:
configure {
    service {
        vpls "VPLS 2" {
            admin-state enable
            service-id 2
            customer "1"
            sap 1/1/c1/1:2 { }
            sap lag-1:2 { }
        }
```

For I-VPLS 2, PE-2 is the DF in ES ESI-24, as follows:

```
[/]
A:admin@PE-2# show service id 2 ethernet-segment
No sap entries

===============================================================================
SDP Ethernet-Segment Information
===============================================================================
SDP                   Eth-Seg                           Status
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
26:2                  ESI-24                          DF
===============================================================================
No vxlan instance entries
```

PE-3 is the DF in virtual ES ESI-34_2, as follows:

```
[/]
A:admin@PE-3# show service id 2 ethernet-segment

===============================================================================
SAP Ethernet-Segment Information
===============================================================================
SAP                   Eth-Seg                         Status
-------------------------------------------------------------------------------
lag-1:2               ESI-34_2                        DF
===============================================================================
No sdp entries
No vxlan instance entries
```

PE-4 is the Non-DF (NDF) for both ESI-24 and ESI-34_2, as follows:

```
[/]
A:admin@PE-4# show service id 2 ethernet-segment

===============================================================================
SAP Ethernet-Segment Information
===============================================================================
SAP                   Eth-Seg                         Status
-------------------------------------------------------------------------------
lag-1:2               ESI-34_2                        NDF
===============================================================================

===============================================================================
SDP Ethernet-Segment Information
===============================================================================
SDP                   Eth-Seg                         Status
-------------------------------------------------------------------------------
46:2                  ESI-24                          NDF
===============================================================================
No vxlan instance entries
```

When H-8 sends an IGMP report to join multicast group 232.1.1.1 from source 192.168.55.2, CE-5 forwards the multicast stream after receiving the corresponding PIM join message. PE-1 forwards the multicast traffic on the P2MP mLDP tunnel to all EVPN-MPLS destinations: PE-2, PE-3, and PE-4. PE-2 is the DF for ESI-24 and forwards the traffic to MTU-6, which forwards it to CE-8, where it is sent to the attached receiver H-8 that joined the multicast group. PE-3 is the DF for ESI-34_2 and sends the multicast stream to MTU-9, which forwards it to CE-7, where it is dropped because no attached receiver has joined the multicast group. PE-4 is the NDF for both ESs, so it does not forward the traffic to MTU-6 or MTU-9. Figure 258: EVPN-MPLS with MH - PIM Snooping Disabled – Receiver H-8 Joined shows how this multicast is forwarded when PIM snooping is disabled.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 258: EVPN-MPLS with MH - PIM Snooping Disabled – Receiver H-8 Joined

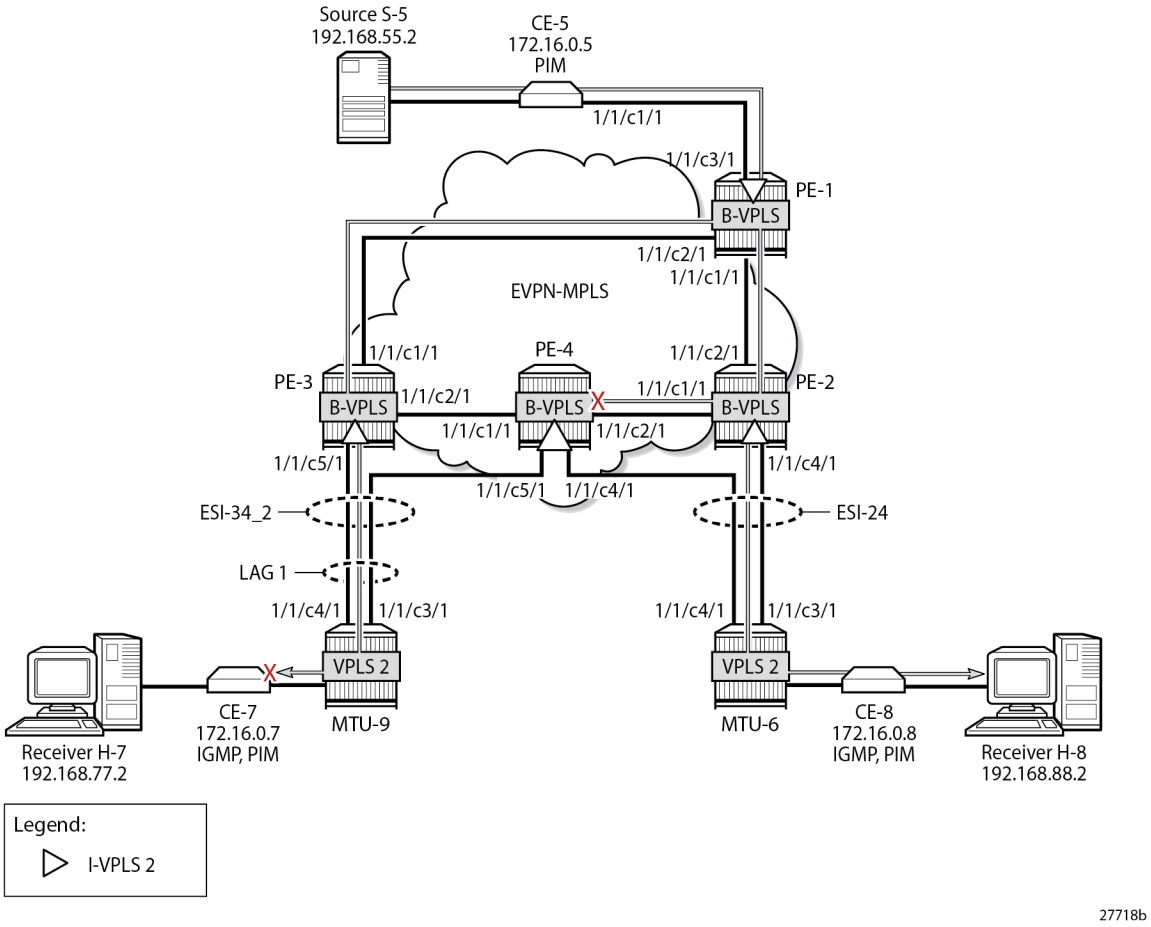

PBB-EVPN with MH – PIM Snooping for IPv4 Enabled

PIM snooping for IPv4 is enabled in I-VPLS 2 on all PEs with the following command:

```
On all PEs:
configure {
    service {
        vpls "I-VPLS 2" {
            pim-snooping { }
        }
```

All PEs have three PIM snooping neighbors: CE-5, CE-7, and CE-8. The list of PIM snooping neighbors on PE-1 is as follows:

```
[/]
A:admin@PE-1# show service id 2 pim-snooping neighbor

===============================================================================
PIM Snooping Neighbors ipv4
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Port Id                 Nbr DR Prty     Up Time       Expiry Time     Hold Time
  Nbr Address
-------------------------------------------------------------------------------
SAP:1/1/c3/1:2          1               0d 00:01:09   0d 00:01:35     105
  172.16.0.5
b-EVPN-MPLS             1               0d 00:01:12   0d 00:01:32     105
  172.16.0.7
b-EVPN-MPLS             1               0d 00:01:22   0d 00:01:23     105
  172.16.0.8
-------------------------------------------------------------------------------
Neighbors : 3
===============================================================================
```

When H-7 and H-8 join the group 232.1.1.1 via source 192.168.55.2, the PIM join messages are snooped by the PEs and the MFIB is built. The MFIB on PE-1 contains one entry for group address 232.1.1.1 and source address 192.168.55.2 with four port IDs: the local SAP to CE-5 and the B-VPLS PBB-EVPN destinations, as follows:

```
[/]
A:admin@PE-1# show service id 2 mfib

===============================================================================
Multicast FIB, Service 2
===============================================================================
Source Address  Group Address         Port Id                     Svc Id   Fwd
                                                                        Blk
-------------------------------------------------------------------------------
192.168.55.2    232.1.1.1             sap:1/1/c3/1:2              Local    Fwd
                                      b-mpls:192.0.2.2:524282     100      Fwd
                                      b-mpls:192.0.2.3:524282     100      Fwd
                                      b-mpls:192.0.2.4:524282     100      Fwd
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
```

In a similar way, the other PEs that snooped PIM messages build their MFIBs. On PE-2, the following MFIB is shown when H-8 has joined the multicast group.

```
[/]
A:admin@PE-2# show service id 2 mfib

===============================================================================
Multicast FIB, Service 2
===============================================================================
Source Address  Group Address         Port Id                     Svc Id   Fwd
                                                                        Blk
-------------------------------------------------------------------------------
192.168.55.2    232.1.1.1             sdp:26:2                    Local    Fwd
                                      b-mpls:192.0.2.1:524282     100      Fwd
                                      b-mpls:192.0.2.3:524282     100      Fwd
                                      b-mpls:192.0.2.4:524282     100      Fwd
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
```

On PE-3, the following MFIB is present when H-7 has joined the multicast group:

```
[/]
A:admin@PE-3# show service id 2 mfib

===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Multicast FIB, Service 2
===============================================================================
Source Address  Group Address         Port Id                     Svc Id   Fwd
                                                                        Blk
-------------------------------------------------------------------------------
192.168.55.2    232.1.1.1             sap:lag-1:2                 Local    Fwd
                                      b-mpls:192.0.2.1:524282     100      Fwd
                                      b-mpls:192.0.2.2:524282     100      Fwd
                                      b-mpls:192.0.2.4:524282     100      Fwd
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
```

Furthermore, data-driven PIM state synchronization between PEs in an all-active MH ES allows the NDF PE-4 to build its MFIB, even when the NDF does not forward multicast traffic to the receivers. When the NDF has the MFIB information, the failover is faster and the loss of traffic is limited. For data-driven PIM state synchronization, the source BMAC must be identical within the ES, so it only works for all-active MH in PBB-EVPN, not for single-active MH. The MFIB on PE-4 contains the SAP from the all-active MH ESI-34_2, but not the spoke-SDP from the single-active MH ESI-24, as follows:

```
[/]
A:admin@PE-4# show service id 2 mfib

===============================================================================
Multicast FIB, Service 2
===============================================================================
Source Address  Group Address         Port Id                     Svc Id   Fwd
                                                                        Blk
-------------------------------------------------------------------------------
192.168.55.2    232.1.1.1             sap:lag-1:2                 Local    Fwd
                                      b-mpls:192.0.2.1:524282     100      Fwd
                                      b-mpls:192.0.2.2:524282     100      Fwd
                                      b-mpls:192.0.2.3:524282     100      Fwd
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
```

The snooped PIM group information on PE-1 shows the SAP to CE-5 as incoming interface and the b-EVPN-MPLS interface as outgoing, as follows. The split-horizon mechanism prevents multicast traffic coming from the SAP to CE-5 from being returned.

```
[/]
A:admin@PE-1# show service id 2 pim-snooping group detail

===============================================================================
PIM Snooping Source Group ipv4
===============================================================================
Group Address      : 232.1.1.1
Source Address     : 192.168.55.2
Up Time            : 0d 00:02:50

Up JP State        : Joined            Up JP Expiry       : 0d 00:00:10
Up JP Rpt          : Not Joined StarG  Up JP Rpt Override : 0d 00:00:00

RPF Neighbor       : 172.16.0.5
Incoming Intf      : SAP:1/1/c3/1:2
Outgoing Intf List : b-EVPN-MPLS, SAP:1/1/c3/1:2

Forwarded Packets  : 140252            Forwarded Octets   : 210378000
-------------------------------------------------------------------------------
Groups : 1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
```

On PE-2, the incoming interface is the b-EVPN-MPLS interface and the outgoing interface is the spoke-SDP toward MTU-6, as follows:

```
[/]
A:admin@PE-2# show service id 2 pim-snooping group detail

===============================================================================
PIM Snooping Source Group ipv4
===============================================================================
Group Address      : 232.1.1.1
Source Address     : 192.168.55.2
Up Time            : 0d 00:01:11

Up JP State        : Joined             Up JP Expiry       : 0d 00:00:28
Up JP Rpt          : Not Joined StarG   Up JP Rpt Override : 0d 00:00:00

RPF Neighbor       : 172.16.0.5
Incoming Intf      : b-EVPN-MPLS
Outgoing Intf List : b-EVPN-MPLS, SPOKE_SDP:26:2

Forwarded Packets  : 59176              Forwarded Octets   : 89592464
-------------------------------------------------------------------------------
Groups : 1
===============================================================================
```

On PE-3, the incoming interface is the b-EVPN-MPLS interface and the outgoing interface is the SAP lag-1:2 toward MTU-9, as follows:

```
[/]
A:admin@PE-3# show service id 2 pim-snooping group detail

===============================================================================
PIM Snooping Source Group ipv4
===============================================================================
Group Address      : 232.1.1.1
Source Address     : 192.168.55.2
Up Time            : 0d 00:02:53

Up JP State        : Joined             Up JP Expiry       : 0d 00:00:57
Up JP Rpt          : Not Joined StarG   Up JP Rpt Override : 0d 00:00:00

RPF Neighbor       : 172.16.0.5
Incoming Intf      : b-EVPN-MPLS
Outgoing Intf List : b-EVPN-MPLS, SAP:lag-1:2

Forwarded Packets  : 142775             Forwarded Octets   : 216161350
-------------------------------------------------------------------------------
Groups : 1
===============================================================================
```

In case of all-active MH ES ESI-34_2, one of the PE -DF or NDF- in the ES forwards the PIM states to its remote peer and therefore, PE-4 has the same PIM snooping group information as PE-3, as follows:

```
[/]
A:admin@PE-4# show service id 2 pim-snooping group detail

===============================================================================
PIM Snooping Source Group ipv4
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Group Address      : 232.1.1.1
Source Address     : 192.168.55.2
Up Time            : 0d 00:02:55

Up JP State        : Joined             Up JP Expiry       : 0d 00:01:01
Up JP Rpt          : Not Joined StarG   Up JP Rpt Override : 0d 00:00:00

RPF Neighbor       : 172.16.0.5
Incoming Intf      : b-EVPN-MPLS
Outgoing Intf List : b-EVPN-MPLS, SAP:lag-1:2

Forwarded Packets  : 143074             Forwarded Octets   : 216614036
-------------------------------------------------------------------------------
Groups : 1
===============================================================================
```

With the PIM snooping group information available on the NDF, the traffic loss is limited when the NDF PE-4 becomes the DF after failover. Data-driven PIM state synchronization does not store PIM states in a database, so the DF election in the ES should be configured as non-revertive, to prevent that when the preferred DF is restored after a failover, the system would revert to a DF that is unaware of the PIM state.

PE-4 is also NDF in the single-active MH ES ESI-24, but it received no PIM state synchronization information from DF PE-2. Data-driven PIM state synchronization is not supported for single-active MH in PBB-EVPN services, because it is not allowed to have two PEs in a single-active MH ES using the same source BMAC, with the potential risk of traffic sent by remote PEs to the NDF PE (based on it sending to the shared source BMAC) being dropped. However, for a faster failover in single-active MH, multi-chassis synchronization (MCS) can be configured, as described in the next section.

Figure 259: EVPN-MPLS with MH and PIM Snooping – Receivers H-7 and H-8 Joined shows the multicast traffic flow when PIM snooping is enabled and both receivers H-7 and H-8 have joined the multicast group. All PEs receive the multicast traffic on the P2MP tunnel, but only DF PE-2 and DF PE-3 forward the multicast traffic to the MTUs, which forward the traffic to the CEs, where it is forwarded to the receivers.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 259: EVPN-MPLS with MH and PIM Snooping – Receivers H-7 and H-8 Joined

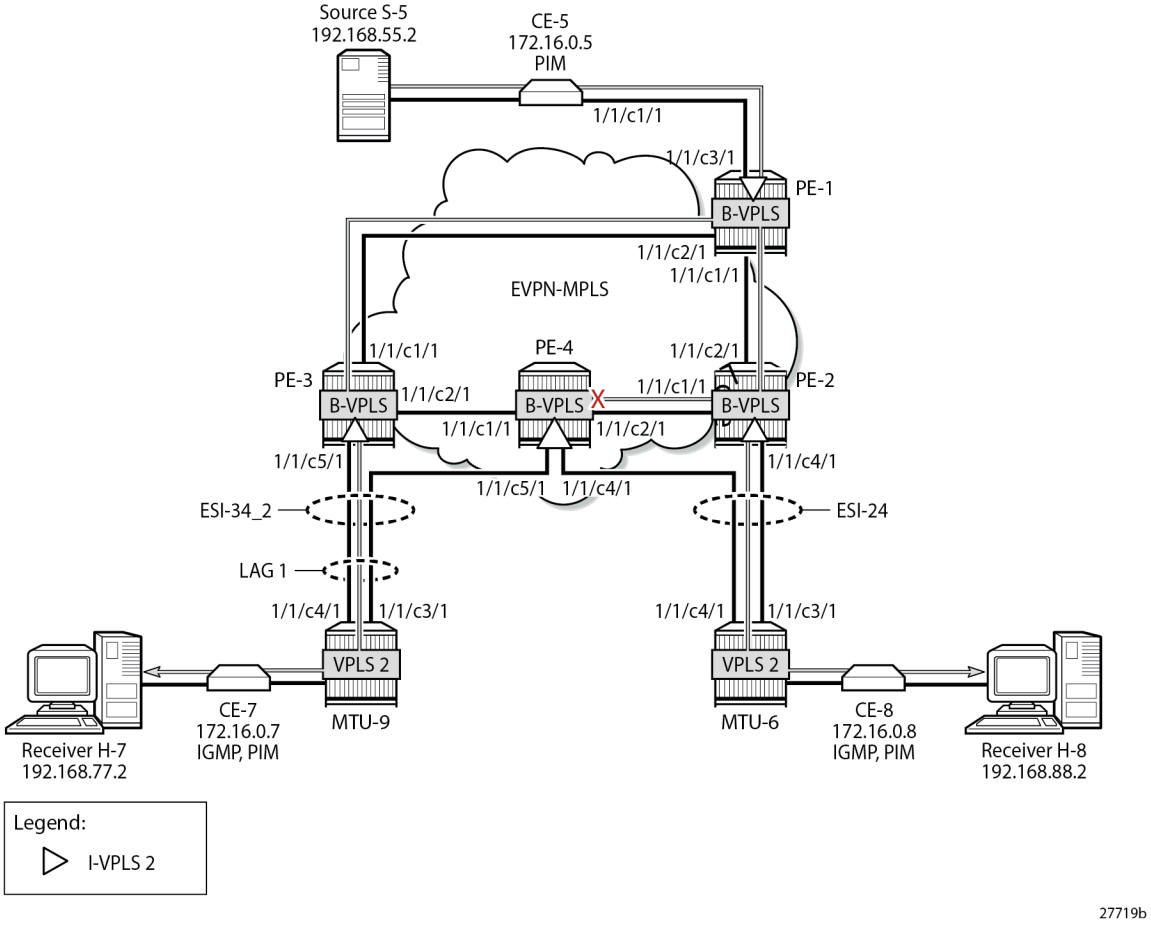

PBB-EVPN with MH – PIM Snooping for IPv4 with MCS

MCS of the IPv4 PIM snooping state for SAPs and spoke-SDPs can optionally be configured in the case of MH. MCS reduces the failover time when data-driven PIM state synchronization is not supported; for example, for single-active MH in PBB-EVPN services. The synchronization information is stored in an MCS synchronization DB. MCS is configured on PE-2, identifying the peer (PE-4), with PIM snooping for spoke-SDPs as MCS client application and the list of spoke-SDPs, as follows:

```
On PE-2:
configure {
    redundancy {
        multi-chassis {
            peer 192.0.2.4 {
                admin-state enable
                sync {
                    admin-state enable
                    pim-snooping {
                        spoke-sdps true
                    }
                    tags {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        sdp 26 {
                            range start 2 end 2 {
                                sync-tag "syncSA"
                            }
                        }
                    }
                }
            }
```

On PE-4, MCS is configured for peer PE-2, as follows:

```
On PE-4:
configure exclusive
    redundancy {
        multi-chassis {
            peer 192.0.2.2 {
                admin-state enable
                sync {
                    admin-state enable
                    pim-snooping {
                        spoke-sdps true
                    }
                    tags {
                        sdp 46 {
                            range start 2 end 2 {
                                sync-tag "syncSA"
                            }
                        }
                    }
                }
            }
```

When H-8 has joined the multicast group, the MCS sync-database on PE-2 shows the PIM snooping entries on the spoke-SDP 26:2 of the single-active MH ESI-24, as follows:

```
[/]
A:admin@PE-2# tools dump redundancy multi-chassis sync-database detail

If no entries are present for an application, no detail will be displayed.

FLAGS LEGEND: ld - local delete; da - delete alarm; pd - pending global delete;
              oal - omcr alarmed; ost - omcr standby

Peer Ip 192.0.2.4


Application pim-snooping-sdp
Sdp-id             Client Key
 SyncTag                          DLen  Flags            timeStamp
  deleteReason code and description                   #ShRec
-------------------------------------------------------------------------------
26:2               Adj 172.16.0.8
 syncSA                           72    -- -- -- --- --- 08/28/2023 22:20:43
  0x0                                                    0
26:2               IfSG SG 192.168.55.2 232.1.1.1
 syncSA                           69    -- -- -- --- --- 08/28/2023 22:20:25
  0x0                                                    0

The following totals are for:
 peer ip ALL, port/lag/sdp ALL, sync-tag ALL, application ALL
Valid Entries:                   2
Locally Deleted Entries:         0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Locally Deleted Alarmed Entries: 0
Pending Global Delete Entries:   0
Omcr Alarmed Entries:            0
Omcr Standby Entries:            0
Associated Shared Records (ALL): 0
Associated Shared Records (LD):  0
```

On PE-4, the MCS sync-database is similar, but with SDP ID 46:2 instead of 26:2.

Even though PE-4 is the NDF for both ESs, the MFIB is populated with the spoke-SDP to MTU-6, as well as the B-VPLS PBB-EVPN destinations to the other PEs, as follows:

```
[/]
A:admin@PE-4# show service id 2 mfib

===============================================================================
Multicast FIB, Service 2
===============================================================================
Source Address  Group Address         Port Id                     Svc Id   Fwd
                                                                         Blk
-------------------------------------------------------------------------------
192.168.55.2    232.1.1.1             sdp:46:2                    Local    Fwd
                                      b-mpls:192.0.2.1:524282     100      Fwd
                                      b-mpls:192.0.2.2:524282     100      Fwd
                                      b-mpls:192.0.2.3:524282     100      Fwd
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
```

The following command on PE-4 shows that the incoming PIM interface is the B-VPLS EVPN-MPLS interface and the spoke-SDP is the outgoing interface. Again, the split-horizon mechanism prevents traffic received from the B-VPLS EVPN-MPLS interface from being forwarded on the B-VPLS EVPN-MPLS interface.

```
[/]
A:admin@PE-4# show service id 2 pim-snooping group detail

===============================================================================
PIM Snooping Source Group ipv4
===============================================================================
Group Address      : 232.1.1.1
Source Address     : 192.168.55.2
Up Time            : 0d 00:02:03

Up JP State        : Joined            Up JP Expiry       : 0d 00:01:19
Up JP Rpt          : Not Joined StarG  Up JP Rpt Override : 0d 00:00:00

RPF Neighbor       : 172.16.0.5
Incoming Intf      : b-EVPN-MPLS
Outgoing Intf List : b-EVPN-MPLS, SPOKE_SDP:46:2

Forwarded Packets  : 101850            Forwarded Octets   : 154200900
-------------------------------------------------------------------------------
Groups : 1
===============================================================================
```

However, PE-4 remains the NDF for both ESs and does not forward any traffic from the B-VPLS EVPN-MPLS interface to the spoke-SDP. Figure 260: PBB-EVPN with MH and PIM Snooping – Receiver H-8 Joined shows the multicast traffic flow when PIM snooping is enabled and receiver H-8 has joined.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 260: PBB-EVPN with MH and PIM Snooping – Receiver H-8 Joined

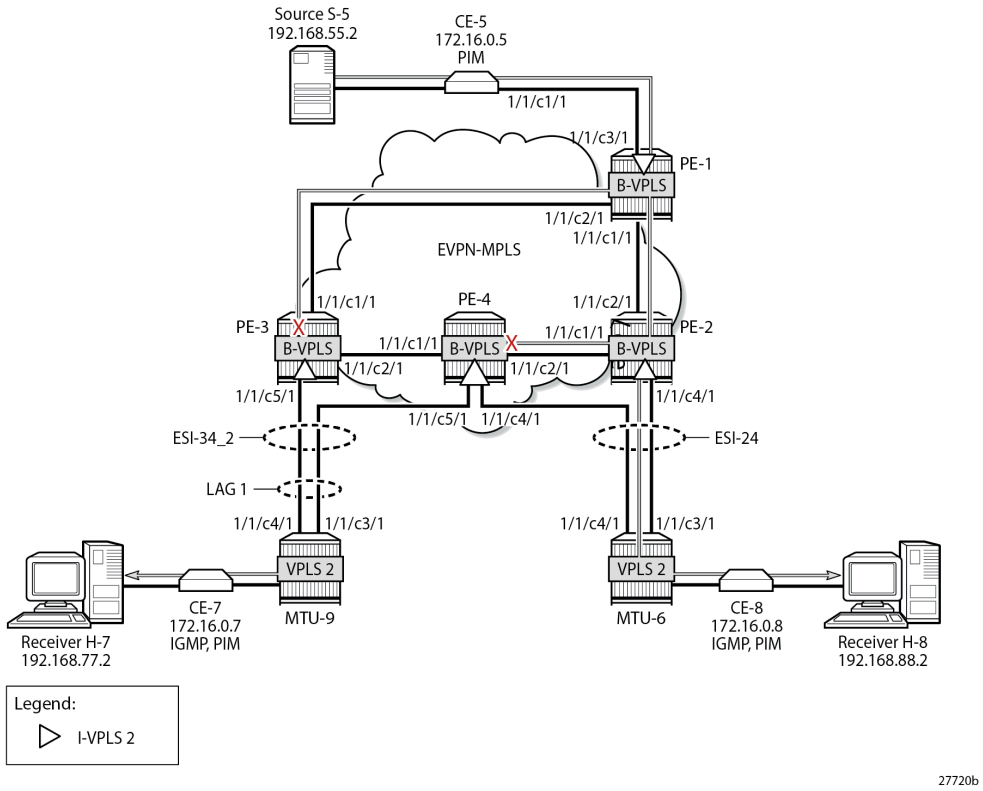

Failover

Figure 259: EVPN-MPLS with MH and PIM Snooping – Receivers H-7 and H-8 Joined showed the multicast traffic flow when both H-7 and H-8 have joined the multicast group. PE-2 is the DF for ESI-24 and PE-4 is the DF for ESI-34_2. The following failures are introduced to force a failover from PE-2 to PE-4 and from PE-3 to PE-4. Data-driven PIM state synchronization is used for all-active MH; MCS is configured for fast failover in the single-active MH ES ESI-24.

On MTU-6, SDP 62 is disabled, as follows:

```
On MTU-6:
configure {
    service {
        sdp 62 {
            admin-state disable
        }
```

On MTU-9, port 1/1/c3/1 toward PE-3 is disabled, as follows:

```
On MTU-9:
configure {
    port 1/1/c3/1 {
        admin-state disable
    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Log 99 on PE-3 shows that the DF state in ESI-34_2 changed to false:

```
200 2023/08/28 22:23:40.972 CEST MINOR: SVCMGR #2095 Base
"Ethernet Segment:ESI-34_2, ISID:2, Designated Forwarding state changed to:false"
```

PE-4 becomes the DF for both ESs, as follows:

```
[/]
A:admin@PE-4# show service id 2 ethernet-segment

===============================================================================
SAP Ethernet-Segment Information
===============================================================================
SAP                   Eth-Seg                         Status
-------------------------------------------------------------------------------
lag-1:2               ESI-34_2                        DF
===============================================================================

===============================================================================
SDP Ethernet-Segment Information
===============================================================================
SDP                   Eth-Seg                         Status
-------------------------------------------------------------------------------
46:2                  ESI-24                          DF
===============================================================================
No vxlan instance entries
```

Figure 261: EVPN-MPLS with MH and PIM Snooping – Multicast Flow after Failover shows the traffic flow after failover to the new DF, PE-4.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 261: EVPN-MPLS with MH and PIM Snooping – Multicast Flow after Failover

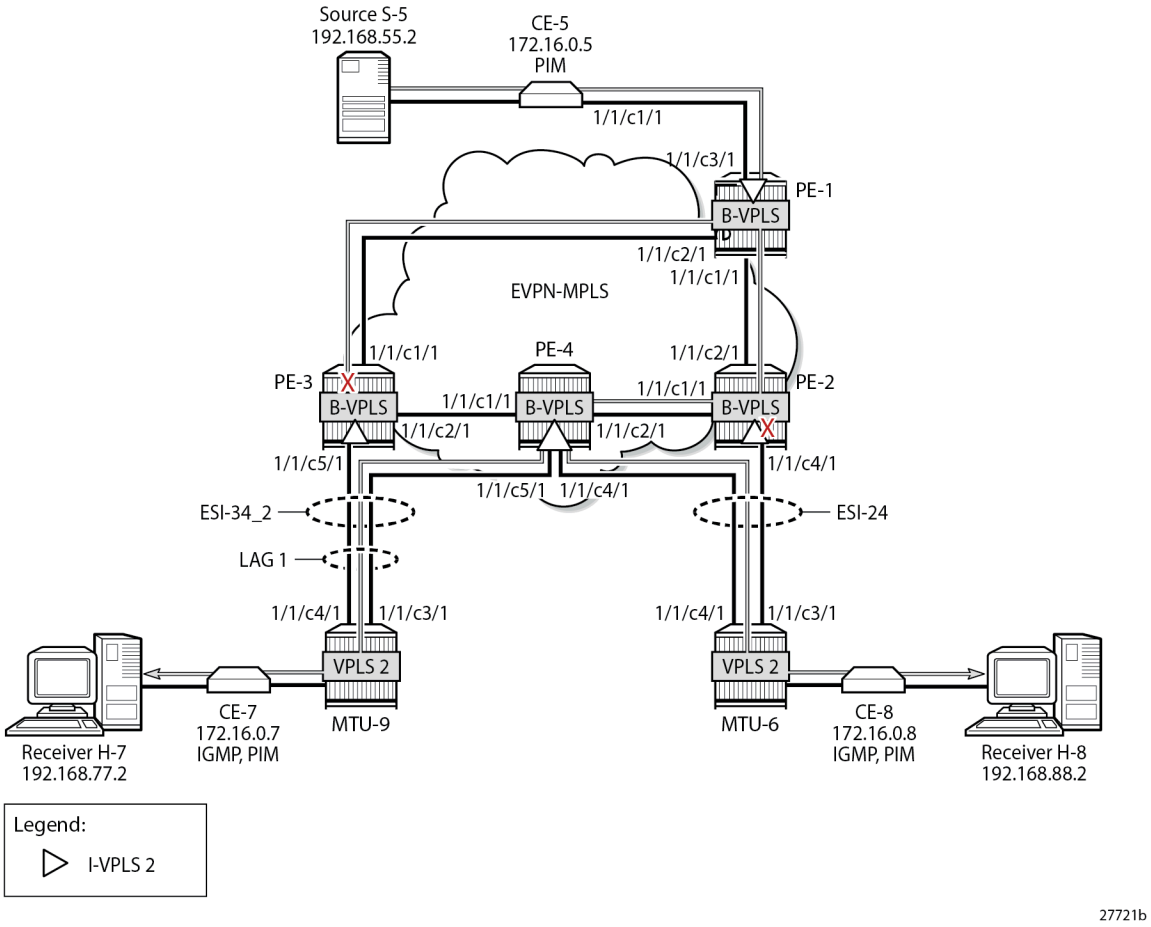

PE-2 receives the multicast stream from PE-1 on port 1/1/c2/1 and forwards it to port 1/1/c1/1 to PE-4; it does not forward to port 1/1/c4/1 because SDP 26 is down, as follows:

```
[/]
A:admin@PE-2# show port 1/1/c1/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c1/1                                 110                          19829
                                       26601                       40719727
===============================================================================

[/]
A:admin@PE-2# show port 1/1/c2/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Id                          Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c2/1                             26528                       40704071
                                        32                           3363
===============================================================================

[/]
A:admin@PE-2# show port 1/1/c3/1 statistics

[/]
A:admin@PE-2# show port 1/1/c4/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c4/1                                25                            2721
                                        26                            2869
===============================================================================
```

PE-4 receives the multicast traffic on port 1/1/c2/1 and forwards it on port 1/1/c4/1 toward MTU-6, and on port 1/1/c5/1 to MTU-9, as follows:

```
[/]
A:admin@PE-4# show port 1/1/c1/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c1/1                                26                            2779
                                        29                            3170
===============================================================================

[/]
A:admin@PE-4# show port 1/1/c2/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c2/1                             26636                       40772035
                                       110                          19843
===============================================================================

[/]
A:admin@PE-4# show port 1/1/c3/1 statistics

[/]
A:admin@PE-4# show port 1/1/c4/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
1/1/c4/1                                29                           3040
                                      26560                        40490857
===============================================================================

[/]
A:admin@PE-4# show port 1/1/c5/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                         Ingress Packets                 Ingress Octets
Id                            Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c5/1                                34                           4248
                                      26566                        39908376
===============================================================================
```

MTU-6 forwards the traffic to CE-8, which forwards it to H-8. MTU-9 forwards the traffic to CE-7, which sends it to H-7. PE-3 drops the multicast traffic because lag-1 is down because of the failure that was introduced at MTU-9 (port disabled).

Conclusion

PIM snooping reduces flooding of multicast traffic in L2 services and can be used in PBB-EVPN I-VPLSs in the same way as in I-VPLSs using B-VPLS without EVPN. PIM snooping can be used in all-active and single-active MH scenarios with data-driven state synchronization and MCS, respectively.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Preference-based and Non-revertive EVPN DF Election

This chapter provides information about Preference-based and Non-revertive EVPN DF Election.

Topics in this chapter include:

Applicability

This chapter was initially written based on SR OS Release 15.0.R3, but the MD-CLI in the current edition corresponds to SR OS Release 21.2.R2. Preference-based and non-revertive EVPN Designated Forwarder (DF) election is supported in SR OS Release 15.0.R1, and later. This mechanism works for Ethernet Segments (ESs) and virtual ESs (vESs).

Overview

RFC 7432 defines the Designated Forwarder (DF) in (PBB-)EVPN networks as the PE that will forward the following packets to a multi-homed node:

- Broadcast, Unknown unicast, and Multicast (BUM) traffic in an all-active multi-homing Ethernet Segment (ES)
- BUM and unicast in a single-active multi-homing ES

For more information about vESs, see chapter Virtual Ethernet Segments.

Figure 262: Virtual Ethernet Segments shows a topology with two vESs.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 262: Virtual Ethernet Segments

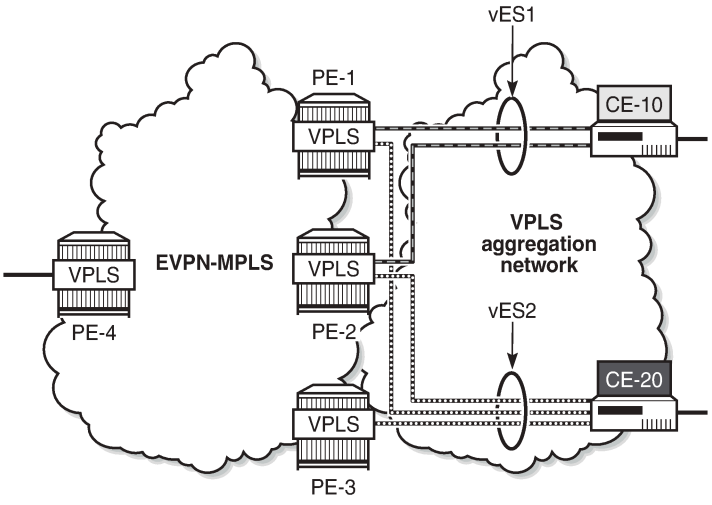

Taking the Ethernet VPN Identifier (EVI) or ISID and the number of PEs in the ES as input, the RFC 7432 service-carving algorithm elects the DF from the list of candidate PEs that advertise the ES identifier (ESI). While this algorithm provides an automated and fair DF distribution across services in the ES, it does not allow the operator to control what PE is the DF for which service. In addition, in case of a DF failure, when the former DF comes back up, a new DF switchover will cause unnecessary packet loss (this mode of operation is called revertive). SR OS implements *draft-ietf-bess-evpn-pref-df* to give more control to the operator on the DF election and avoid the revertive mode.

In SR OS, in addition to the automated service-carving, the DF election can also be controlled by configuring a preference manually. Also, it is possible to force an on-demand DF switchover without reconfiguring the PEs in the ES. Furthermore, the non-revertive option prevents an automatic switchover when a new active PE can preempt the existing DF PE. The non-revertive option avoids service impact when an ES comes back up.

Figure 263: BGP-EVPN extended community for DF election shows the BGP-EVPN extended community defined for DF election and the different values described in draft-ietf-bess-evpn-pref-df.

Figure 263: BGP-EVPN extended community for DF election

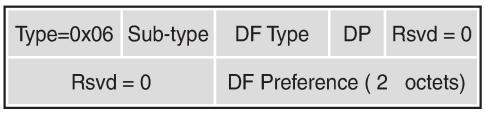

Type=0x06	Sub-type	DF Type	DP	Rsvd = 0
Rsvd = 0		DF Preference (2 octets)		

DP = Do not preempt (non-revertive)
DF = Designated forwarder
- Type 0 – Default, modulo-based DF election (RFC7432)
- Type 1 – Highest Random Weight (HRW) algorithm
- Type 2 – Preference algorithm

26787

The "Do not preempt" (DP) bit is set to enable the non-revertive option. When preference-based service carving is configured in the ES, DF type 2 is advertised along with a 2-byte preference value, which is 32767 by default.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Service carving can be configured in auto mode or manual mode. The preference can only be configured in manual mode.

```
*[ex:/configure service system bgp evpn ethernet-segment "vESI-23_1" df-election]
A:admin@PE-2# service-carving-mode ?

 service-carving-mode <keyword>
 <keyword>  - (auto|manual|off)
 Default    - auto

    Mode of service carving enabled per EVPN associated with this Ethernet segment entry
```

When manual mode is enabled, the following parameters can be configured to control which PE will be elected as DF:

```
*[ex:/configure service system bgp evpn ethernet-segment "vESI-23_1" df-election]
A:admin@PE-2# manual ?

 manual

 evi                   + Enter the evi list instance
 isid                  + Enter the isid list instance
 preference            + Enable the preference context
```

The EVI and ISID ranges configured in the service-carving context do not need to be consistent with any ranges configured for virtual ESs.

When preference is configured manually, the mode can be configured as revertive (default) or non-revertive:

```
*[ex:/configure service system bgp evpn ethernet-segment "vESI-23_1" df-election manual
 preference]
A:admin@PE-2# mode ?

 mode <keyword>
 <keyword>  - (revertive|non-revertive)
 Default    - revertive

 'mode' is: immutable

    Method used to elect the DF

    Warning: Modifying this element recreates 'configure service system bgp evpn
    ethernet-segment "vESI-23_1" df-election manual preference' automatically for the
    new value to take effect.
```

The preference-based EVPN DF election is as follows:

- By default, all SAPs and spoke-SDPs on the configured ES select the highest-preference PE as DF; however, when the EVI or ISID ranges are configured in the ES, the lowest-preference PE is selected.
- When the preference is equal, the DP bit is the tiebreaker: DP=1 wins over DP=0.
- For equal preference and DP, the PE IP address is the tiebreaker: the lowest IP address wins.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Configuration

Figure 264: Example topology with all-active and single-active vESs shows the example topology with six nodes. EVPN-MPLS is configured between the core PE nodes. All-active vESs are configured between PE-2 and PE-3 and single-active vESs are configured between PE-4 and PE-5.

Figure 264: Example topology with all-active and single-active vESs

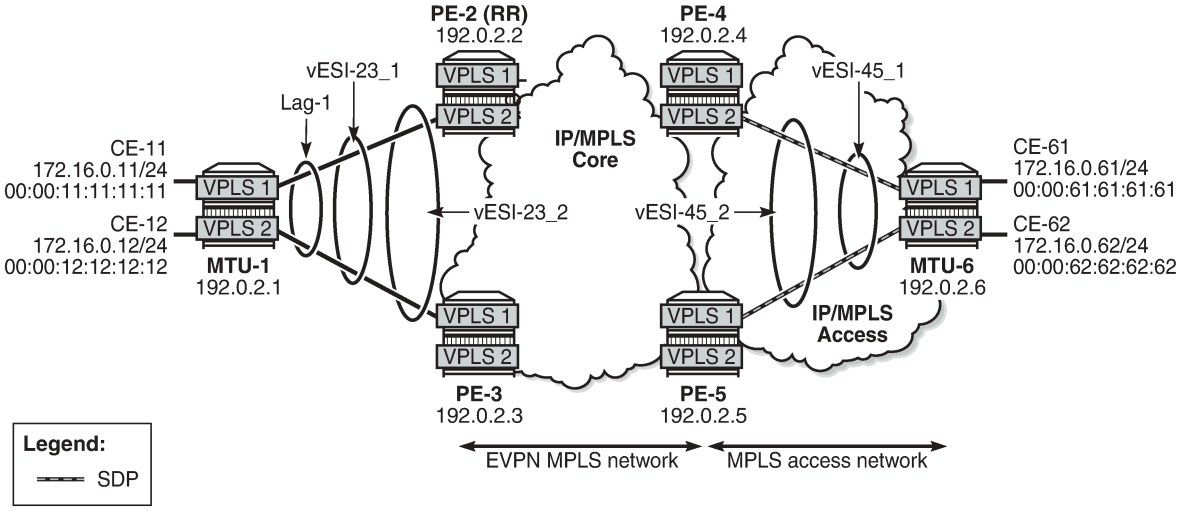

The initial configuration includes:

- Cards, MDAs, ports
- LAG 1 between MTU-1, PE-2, PE-3
- Router interfaces
- IS-IS (alternatively, OSPF could be used)
- LDP

BGP is configured on the four core PEs with PE-2 as Route Reflector (RR). The BGP configuration on RR PE-2 is as follows:

```
# on RR PE-2:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            split-horizon true
            rapid-update {
                evpn true
            }
            group "internal" {
                peer-as 64500
                family {
                    evpn true
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
                cluster {
                    cluster-id 1.1.1.1
                }
            }
            neighbor "192.0.2.3" {
                group "internal"
            }
            neighbor "192.0.2.4" {
                group "internal"
            }
            neighbor "192.0.2.5" {
                group "internal"
            }
        }
```

VPLS 1 and VPLS 2 are configured on each node. The PEs have EVPN-MPLS enabled. The configuration on PE-2 is as follows:

```
# on PE-2:
configure {
    service {
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
            }
            bgp-evpn {
                evi 1
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    ecmp 2
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            sap lag-1:1.1 {
            }
        }
        vpls "VPLS 2" {
            admin-state enable
            service-id 2
            customer "1"
            bgp 1 {
            }
            bgp-evpn {
                evi 2
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    ecmp 2
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            sap lag-1:2.1 {
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The configuration on the other PEs is similar; PE-4 and PE-5 have a spoke-SDP configured instead of a SAP. For an explanation of the configuration, see chapter EVPN for MPLS Tunnels.

Service carving: auto mode

On PE-2 and PE-3, the following all-active multi-homing vESs are configured:

```
# on PE-2, PE-3:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "vESI-23_1" {
                        admin-state enable
                        type virtual
                        esi 01:00:00:00:00:23:01:00:00:01
                        multi-homing-mode all-active
                        df-election {
                            es-activation-timer 3
                            service-carving-mode auto
                        }
                        association {
                            lag "lag-1" {
                                virtual-ranges {
                                    qinq {
                                        s-tag 1 {
                                            end 1
                                        }
                                    }
                                }
                            }
                        }
                    }
                    ethernet-segment "vESI-23_2" {
                        admin-state enable
                        type virtual
                        esi 01:00:00:00:00:23:02:00:00:01
                        multi-homing-mode all-active
                        df-election {
                            es-activation-timer 3
                            service-carving-mode auto
                        }
                        association {
                            lag "lag-1" {
                                virtual-ranges {
                                    qinq {
                                        s-tag 2 {
                                            end 2
                                        }
                                    }
                                }
                            }
                        }
                    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The service carving mode is set to **auto**, so the DF election is based on a modulo function of the EVI and the number of DF candidates. In the vES "vESI-23_1", there are two DF candidates, PE-2 and PE-3, listed in that order because PE-2 has the lower system IP address, as follows:

```
[/]
A:admin@PE-3# show service system bgp-evpn ethernet-segment name "vESI-23_1" all
                                          | match "EVI Information" post-lines 19
EVI Information
===============================================================================
EVI                 SvcId               Actv Timer Rem      DF
-------------------------------------------------------------------------------
1                   1                   0                   yes
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
-------------------------------------------------------------------------------
DF Candidate list
-------------------------------------------------------------------------------
EVI                                     DF Address
-------------------------------------------------------------------------------
1                                       192.0.2.2
1                                       192.0.2.3
-------------------------------------------------------------------------------
Number of entries: 2
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
===============================================================================
```

The first DF candidate from the list will be selected when the result of the modulo function equals 0; the second DF candidate when the result equals 1. The calculation is as follows:

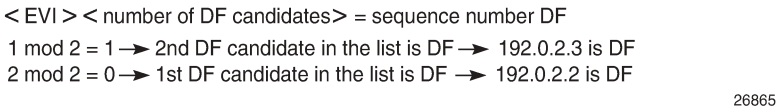

The following shows that PE-2 is not the DF for VPLS 1, but it is the DF for VPLS 2:

```
[/]
A:admin@PE-2# show service id 1 ethernet-segment

===============================================================================
SAP Ethernet-Segment Information
===============================================================================
SAP                   Eth-Seg                         Status
-------------------------------------------------------------------------------
lag-1:1.1             vESI-23_1                       NDF
===============================================================================
No sdp entries
No vxlan instance entries
```

```
[/]
A:admin@PE-2# show service id 2 ethernet-segment

===============================================================================
SAP Ethernet-Segment Information
===============================================================================
SAP                   Eth-Seg                         Status
-------------------------------------------------------------------------------
lag-1:2.1             vESI-23_2                       DF
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
No sdp entries
No vxlan instance entries
```

Instead of the preceding show commands, the following tools commands can be used:

```
[/]
A:admin@PE-2# tools dump service system bgp-evpn ethernet-segment "vESI-23_1"  evi 1 df

[04/26/2021 09:19:25] Computed DF: 192.0.2.3 (Remote) (Boot Timer Expired: Yes)


[/]
A:admin@PE-2# tools dump service system bgp-evpn ethernet-segment "vESI-23_2"  evi 2 df

[04/26/2021 09:19:25] Computed DF: 192.0.2.2 (This Node) (Boot Timer Expired: Yes)
```

Service carving: preference-based manual mode

To have more control, the vES can be configured in manual mode. The following reconfigures the vES "vESI-23_1" in manual mode, preference-based and revertive with preference value 32767 (default) on PE-2 and 5000 on PE-3, whereas vES "vESI-23_2" is preference-based and non-revertive with preference value 15000 on PE-2 and 20000 on PE-3.

An EVI range is configured for ES "vESI-23_2", but not for ES "vESI-23_1". When no EVI range is configured, the highest preference wins; for configured EVI ranges, the lowest preference wins. When there are no failures, PE-2 will be the DF for "vESI-23_1" (highest preference) and for "vESI-23_2" (lowest preference for configured EVI 2).

On PE-2, the ESs are reconfigured as follows:

```
# on PE-2:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "vESI-23_1" {
                        df-election {
                            service-carving-mode manual
                            manual {
                                preference {
                                }
                            }
                        }
                    }
                    ethernet-segment "vESI-23_2" {
                        df-election {
                            service-carving-mode manual
                            manual {
                                evi 2 {
                                    end 2
                                }
                                preference {
                                    mode non-revertive
                                    value 15000
                                }
                            }
                        }
                    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
```

The **non-revertive** mode is configured for vES "vESI-23_2", but not for vES "vESI-23_1".

On PE-3, the ES configuration is modified as follows:

```
# on PE-3:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "vESI-23_1" {
                        df-election {
                            service-carving-mode manual
                            manual {
                                preference {
                                    value 5000
                                }
                            }
                        }
                    }
                    ethernet-segment "vESI-23_2" {
                        df-election {
                            service-carving-mode manual
                            manual {
                                evi 2 {
                                    end 2
                                }
                                preference {
                                    mode non-revertive
                                    value 20000
                                }
                            }
                        }
                    }
                }
```

For the single-active multi-homing vESs on PE-4 and PE-5, the same preferences are configured manually. The ES configuration on PE-4 is as follows:

```
# on PE-4:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "vESI-45_1" {
                        admin-state enable
                        type virtual
                        esi 01:00:00:00:00:45:01:00:00:01
                        multi-homing-mode single-active
                        df-election {
                            es-activation-timer 3
                            service-carving-mode manual
                            manual {
                                preference {
                                }
                            }
                        }
                        association {
                            sdp 46 {
                                virtual-ranges {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                                    vc-id 1 {
                                        end 1
                                    }
                                    vc-id 500 {
                                        end 501
                                    }
                                }
                            }
                        }
                    }
                    ethernet-segment "vESI-45_2" {
                        admin-state enable
                        type virtual
                        esi 01:00:00:00:00:45:02:00:00:01
                        multi-homing-mode single-active
                        df-election {
                            es-activation-timer 3
                            service-carving-mode manual
                            manual {
                                evi 2 {
                                    end 2
                                }
                                preference {
                                    mode non-revertive
                                    value 15000
                                }
                            }
                        }
                        association {
                            sdp 46 {
                                virtual-ranges {
                                    vc-id 2 {
                                        end 2
                                    }
                                }
                            }
                        }
                    }
                }
```

The ES configuration on PE-5 is as follows:

```
# on PE-5:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "vESI-45_1" {
                        admin-state enable
                        type virtual
                        esi 01:00:00:00:00:45:01:00:00:01
                        multi-homing-mode single-active
                        df-election {
                            es-activation-timer 3
                            service-carving-mode manual
                            manual {
                                preference {
                                    value 5000
                                }
                            }
                        }
                        association {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                            sdp 56 {
                                virtual-ranges {
                                    vc-id 1 {
                                        end 1
                                    }
                                    vc-id 500 {
                                        end 501
                                    }
                                }
                            }
                        }
                        ethernet-segment "vESI-45_2" {
                            admin-state enable
                            type virtual
                            esi 01:00:00:00:00:45:02:00:00:01
                            multi-homing-mode single-active
                            df-election {
                                es-activation-timer 3
                                service-carving-mode manual
                                manual {
                                    evi 2 {
                                        end 2
                                    }
                                    preference {
                                        mode non-revertive
                                        value 20000
                                    }
                                }
                            }
                            association {
                                sdp 56 {
                                    virtual-ranges {
                                        vc-id 2 {
                                            end 2
                                        }
                                    }
                                }
                            }
                        }
                    }
```

The preference configuration must be consistent across the PEs in the ES (manual or auto), otherwise the system reverts to the modulo-based DF election.

With preference-based DF election configured with default preference value 32767 and revertive, PE-4 sends the following BGP-EVPN update to the RR PE-2. The **df-election** extended community shows the DP=0 (revertive) and DF preference 32767.

```
# on PE-4:
60 2021/04/26 09:21:40.829 CEST MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 71
    Flag: 0x90 Type: 14 Len: 34 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.4
        Type: EVPN-ETH-SEG Len: 23 RD: 192.0.2.4:0
              ESI: 01:00:00:00:00:45:01:00:00:01, IP-Len: 4 Orig-IP-Addr: 192.0.2.4
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        df-election::DF-Type:Preference/DP:0/DF-Preference:32767/AC:1
        target:00:00:00:00:45:01
"
```

The following command shows the information in the preceding BGP-EVPN Ethernet-segment route for "vESI-45_1" sent by PE-4 to the RR PE-2:

```
[/]
A:admin@PE-4# show router bgp routes evpn eth-seg hunt
===============================================================================
 BGP Router ID:192.0.2.4        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Eth-Seg Routes
===============================================================================
---snip---
-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
Network        : n/a
Nexthop        : 192.0.2.4
To             : 192.0.2.2
Res. Nexthop   : n/a
Local Pref.    : 100                    Interface Name : NotAvailable
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : n/a
Connector      : None
Community      :
                 df-election::DF-Type:Preference/DP:0/DF-Preference:32767/AC:1
                 target:00:00:00:00:45:01
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.2
Origin         : IGP
AS-Path        : No As-Path
EVPN type      : ETH-SEG
ESI            : 01:00:00:00:00:45:01:00:00:01
Originator IP  : 192.0.2.4
Route Dist.    : 192.0.2.4:0
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
---snip---
```

The following command shows the DF preference election information for ES "vESI-45_1" with the preference mode revertive, the configured preference value on PE-4 (default 32767), and the operational preference value. No EVI ranges or ISID ranges are configured in this ES.

```
[/]
A:admin@PE-4# show service system bgp-evpn ethernet-segment name "vESI-45_1"

===============================================================================
Service Ethernet Segment
===============================================================================
Name                    : vESI-45_1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Eth Seg Type            : Virtual
Admin State             : Enabled           Oper State         : Up
ESI                     : 01:00:00:00:00:45:01:00:00:01
Multi-homing            : singleActive      Oper Multi-homing  : singleActive
ES SHG Label            : 524276
Source BMAC LSB         : <none>
Sdp Id                  : 46
ES Activation Timer     : 3 secs
Oper Group              : (Not Specified)
Svc Carving             : manual            Oper Svc Carving   : manual
Cfg Range Type          : lowest-pref

-------------------------------------------------------------------------------
DF Pref Election Information
-------------------------------------------------------------------------------
Preference     Preference     Last Admin Change        Oper Pref      Do No
Mode           Value                                   Value          Preempt
-------------------------------------------------------------------------------
revertive      32767          04/26/2021 09:21:41      32767          Disabled
-------------------------------------------------------------------------------
EVI Ranges: <none>
ISID Ranges: <none>
===============================================================================
```

The following command shows the DF preference election information for ES "vESI-45_2" with the preference mode non-revertive, the configured preference value on PE-4 (15000), and the operational preference value. The only configured EVI range is from 2 to 2. No ISID ranges are configured. For the configured EVI or ISID values, the lowest preference wins, as shown by the **Cfg Range Type : lowest-pref** parameter.

```
[/]
A:admin@PE-4# show service system bgp-evpn ethernet-segment name "vESI-45_2"

===============================================================================
Service Ethernet Segment
===============================================================================
Name                    : vESI-45_2
Eth Seg Type            : Virtual
Admin State             : Enabled           Oper State         : Up
ESI                     : 01:00:00:00:00:45:02:00:00:01
Multi-homing            : singleActive      Oper Multi-homing  : singleActive
ES SHG Label            : 524275
Source BMAC LSB         : <none>
Sdp Id                  : 46
ES Activation Timer     : 3 secs
Oper Group              : (Not Specified)
Svc Carving             : manual            Oper Svc Carving   : manual
Cfg Range Type          : lowest-pref

-------------------------------------------------------------------------------
DF Pref Election Information
-------------------------------------------------------------------------------
Preference     Preference     Last Admin Change        Oper Pref      Do No
Mode           Value                                   Value          Preempt
-------------------------------------------------------------------------------
non-revertive  15000          04/26/2021 09:21:41      15000          Enabled
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
EVI Ranges
-------------------------------------------------------------------------------
From                                To
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
2                                   2
-------------------------------------------------------------------------------
ISID Ranges: <none>
===============================================================================
```

It is important to note that a router will prune a remote PE from the DF candidate list for an ES if it does not receive the corresponding Auto Discovery (AD) per-EVI and AD per-ES routes for that PE. A remote PE will not be shown in the DF Candidate list if its AD per-ES route is withdrawn. This is only true for EVPN. In PBB-EVPN, there are no AD routes, therefore the DF Candidate list is built out of the ES routes only.

DF election: higher preference prevails for non-configured EVI ranges

The PEs run the DF election per PE per EVI, and the elected DF for a service will activate the SAP/Spoke-SDP when the es-activation-timer expires. PE-4 is the DF in "vESI-45_1" used in VPLS 1, as follows. The EVI is not configured in ES "vESI-45_1", so the higher preference prevails. The ES "vESI-45_1" has (default) preference 32767 on PE-4 (DF) and preference 5000 on PE-5 (Non-Designated Forwarder (NDF)).

```
[/]
A:admin@PE-4# show service id 1 ethernet-segment
No sap entries

===============================================================================
SDP Ethernet-Segment Information
===============================================================================
SDP                  Eth-Seg                         Status
-------------------------------------------------------------------------------
46:1                 vESI-45_1                       DF
===============================================================================
No vxlan instance entries
```

```
[/]
A:admin@PE-5# show service id 1 ethernet-segment
No sap entries

===============================================================================
SDP Ethernet-Segment Information
===============================================================================
SDP                  Eth-Seg                         Status
-------------------------------------------------------------------------------
56:1                 vESI-45_1                       NDF
===============================================================================
No vxlan instance entries
```

The preference value can be modified on the fly on an active ES. This allows the user to force a new DF for the ES for maintenance operations on the former DF or other reasons.

DF election: lowest preference prevails for configured EVI ranges

ES "vESI-45_2" is configured with EVI 2, so the lowest preference prevails. The admin preference value is 15000 on PE-4 and 20000 on PE-5. Both PE-4 and PE-5 are DF candidates, but PE-4 has the lowest preference, so it will be the DF, as follows:

```
[/]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
A:admin@PE-4# show service system bgp-evpn ethernet-segment name "vESI-45_2" all
                                                    | match "EVI Range" post-lines 28
EVI Ranges
-------------------------------------------------------------------------------
From                                    To
-------------------------------------------------------------------------------
2                                       2
-------------------------------------------------------------------------------
ISID Ranges: <none>
===============================================================================
===============================================================================
EVI Information
===============================================================================
EVI                 SvcId               Actv Timer Rem      DF
-------------------------------------------------------------------------------
2                   2                   0                   yes
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
-------------------------------------------------------------------------------
DF Candidate list
-------------------------------------------------------------------------------
EVI                                     DF Address
-------------------------------------------------------------------------------
2                                       192.0.2.4
2                                       192.0.2.5
-------------------------------------------------------------------------------
Number of entries: 2
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
===============================================================================
```

DF election: DP prevails when preferences are equal

The service carving in the ES is configured with default preference and non-revertive option, as follows:

```
# on PE-5:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "vESI-45_1" {
                        df-election {
                            es-activation-timer 3
                            service-carving-mode manual
                            manual {
                                preference {
                                    mode non-revertive
                                    delete value    # default value: 32767
                                }
                            }
                        }
```

The ES configuration on PE-4 remains unchanged, so the behavior is revertive. PE-4 and PE-5 have the same preference (default 32767), but PE-5 is non-revertive and becomes the DF, as follows:

```
[/]
A:admin@PE-5# show service id 1 ethernet-segment
No sap entries
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
SDP Ethernet-Segment Information
===============================================================================
SDP                  Eth-Seg                          Status
-------------------------------------------------------------------------------
56:1                 vESI-45_1                        DF
===============================================================================
No vxlan instance entries
```

DF election: lowest IP address prevails when preferences and DP are equal

The vES configuration on PE-4 is modified by enabling the non-revertive option, as follows:

```
# on PE-4:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "vESI-45_1" {
                        df-election {
                            manual {
                                preference {
                                    mode non-revertive
                                }
                            }
                        }
```

PE-4 and PE-5 have an equal preference (default value = 32767) and non-revertive behavior. The tiebreaker for the DF selection is the IP address. PE-4 has the lower IP address and becomes the DF, as follows:

```
[/]
A:admin@PE-4# show service id 1 ethernet-segment
No sap entries

===============================================================================
SDP Ethernet-Segment Information
===============================================================================
SDP                  Eth-Seg                          Status
-------------------------------------------------------------------------------
46:1                 vESI-45_1                        DF
===============================================================================
No vxlan instance entries
```

Service-carving configuration must be consistent

When the service carving on one of the PEs in the ES is configured in auto mode while one of the other PEs in the ES is configured in manual mode, the system reverts to modulo-based auto mode. The configuration of ES "vESI-45_1" remains unchanged on PE-4, but is modified on PE-5, as follows:

```
# on PE-5#
configure {
    service {
        system {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            bgp {
                evpn {
                    ethernet-segment "vESI-45_1" {
                        df-election {
                            service-carving-mode auto
                            delete manual
                        }
```

ES "vESI-45_1" will operate in auto mode on PE-4 and on PE-5. The following **show** command on PE-4 shows that the ES is configured in manual mode, but operates in auto mode:

```
[/]
A:admin@PE-4# show service system bgp-evpn ethernet-segment name "vESI-45_1"

===============================================================================
Service Ethernet Segment
===============================================================================
Name                    : vESI-45_1
Eth Seg Type            : Virtual
Admin State             : Enabled            Oper State         : Up
ESI                     : 01:00:00:00:00:45:01:00:00:01
Multi-homing            : singleActive       Oper Multi-homing  : singleActive
ES SHG Label            : 524273
Source BMAC LSB         : <none>
Sdp Id                  : 46
ES Activation Timer     : 3 secs
Oper Group              : (Not Specified)
Svc Carving             : manual             Oper Svc Carving   : auto
Cfg Range Type          : lowest-pref

-------------------------------------------------------------------------------
DF Pref Election Information
-------------------------------------------------------------------------------
Preference     Preference     Last Admin Change        Oper Pref      Do No
Mode           Value                                   Value          Preempt
-------------------------------------------------------------------------------
non-revertive  32767          04/26/2021 09:32:46      32767          Enabled
-------------------------------------------------------------------------------
EVI Ranges: <none>
ISID Ranges: <none>
===============================================================================
```

The following command on PE-5 shows that the ES is configured in auto mode and operates in auto mode:

```
[/]
A:admin@PE-5# show service system bgp-evpn ethernet-segment name "vESI-45_1"

===============================================================================
Service Ethernet Segment
===============================================================================
Name                    : vESI-45_1
Eth Seg Type            : Virtual
Admin State             : Enabled            Oper State         : Up
ESI                     : 01:00:00:00:00:45:01:00:00:01
Multi-homing            : singleActive       Oper Multi-homing  : singleActive
ES SHG Label            : 524276
Source BMAC LSB         : <none>
Sdp Id                  : 56
ES Activation Timer     : 3 secs
Oper Group              : (Not Specified)
Svc Carving             : auto               Oper Svc Carving   : auto
Cfg Range Type          : primary
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
```

For the remainder of the chapter, the vES configuration for "vESI-45_1" on PE-4 and PE-5 is restored to the initial settings. On PE-4, vESI-45_1" is configured with manual service-carving mode, reversive, and with default preference value (32767):

```
# on PE-4:
configure {
    service {
        system {
            bgp {
                evpn
                    ethernet-segment "vESI-45_1" {
                        df-election {
                            service-carving-mode manual
                            manual {
                                preference {
                                    delete mode     # default mode: revertive
                                }
                            }
```

On PE-5, "vESI-45_1" is configured with manual service-carving mode, reversive, and with preference value 5000:

```
# on PE-5:
configure {
    service {
        system {
            bgp {
                evpn
                    ethernet-segment "vESI-45_1" {
                        df-election {
                            service-carving-mode manual
                            manual {
                                preference {
                                    value 5000
                                }
                            }
```

When there are no failures, PE-4 is the DF, because it has a higher preference.

Revertive behavior

When SDP 64 fails on MTU-6, PE-4 becomes the NDF for ES "vESI-45_1" and PE-5 will be the DF instead, as follows. The failure is emulated by disabling the SDP on MTU-6.

```
# on MTU-6:
configure {
    service {
        sdp 64 {
            admin-state disable
```

When the PE is not a candidate DF because it cannot be used, the operational preference value equals 0, as follows:

```
[/]
A:admin@PE-4# show service system bgp-evpn ethernet-segment name "vESI-45_1"
                                            | match "DF Pref Election" post-lines 6
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
DF Pref Election Information
-------------------------------------------------------------------------------
Preference     Preference     Last Admin Change        Oper Pref      Do No
Mode           Value                                   Value          Preempt
-------------------------------------------------------------------------------
revertive      32767          04/26/2021 09:32:46      0              Disabled
-------------------------------------------------------------------------------
```

PE-5 is the only DF candidate in ES "vESI-45_1" for VPLS 1:

```
[/]
A:admin@PE-4# show service system bgp-evpn ethernet-segment name "vESI-45_1"
                                                                        evi evi-1 1

===============================================================================
EVI DF and Candidate List
===============================================================================
EVI           SvcId         Actv Timer Rem      DF  DF Last Change
-------------------------------------------------------------------------------
1             1             0                   no  04/26/2021 09:37:32
===============================================================================

===============================================================================
DF Candidates                           Time Added
-------------------------------------------------------------------------------
192.0.2.5                               04/26/2021 09:32:50
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
```

PE-5 is the DF in "vESI-45_1" for VPLS 1:

```
[/]
A:admin@PE-5# show service id 1 ethernet-segment
No sap entries

===============================================================================
SDP Ethernet-Segment Information
===============================================================================
SDP                  Eth-Seg                         Status
-------------------------------------------------------------------------------
56:1                 vESI-45_1                       DF
===============================================================================
No vxlan instance entries
```

The preference mode for this vES is revertive and the DF preference for PE-5 is 5000, as follows:

```
[/]
A:admin@PE-5# show service system bgp-evpn ethernet-segment name "vESI-45_1"
                                          | match "DF Pref Election" post-lines 6
DF Pref Election Information
-------------------------------------------------------------------------------
Preference     Preference     Last Admin Change        Oper Pref      Do No
Mode           Value                                   Value          Preempt
-------------------------------------------------------------------------------
revertive      5000           04/26/2021 09:37:02      5000           Disabled
-------------------------------------------------------------------------------
```

When the failure is restored, the system reverts and PE-4 will again be the DF for "vESI-45_1" in VPLS 1.

```
# on MTU-6:
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
configure {
    service {
        sdp 64 {
            admin-state enable
```

```
[/]
A:admin@PE-4# show service id 1 ethernet-segment
No sap entries

===============================================================================
SDP Ethernet-Segment Information
===============================================================================
SDP                  Eth-Seg                         Status
-------------------------------------------------------------------------------
46:1                 vESI-45_1                       DF
===============================================================================
No vxlan instance entries
```

Non-revertive behavior

When no failures have occurred, PE-4 is the DF for "vESI-45_2" because the lowest preference prevails for the configured EVI 2. The preference of PE-4 is 15000, which is lower than PE-5's preference of 20000.

```
[/]
A:admin@PE-4# show service id 2 ethernet-segment

No sap entries
===============================================================================
SDP Ethernet-Segment Information
===============================================================================
SDP                  Eth-Seg                         Status
-------------------------------------------------------------------------------
46:2                 vESI-45_2                       DF
===============================================================================
No vxlan instance entries
```

A failure is simulated as follows:

```
# on MTU-6:
configure {
    service {
        sdp 64 {
            admin-state disable
```

When SDP 64 on MTU-6 goes down, SDP 46 on PE-4 goes down which brings the vESs down on PE-4. PE-4 is no longer the DF for "vESI-45_2" and not even a DF candidate anymore. The operational preference value is 0.

```
[/]
A:admin@PE-4# show service system bgp-evpn ethernet-segment name "vESI-45_2"
                                          | match "DF Pref Election" post-lines 6
DF Pref Election Information
-------------------------------------------------------------------------------
Preference     Preference     Last Admin Change        Oper Pref      Do No
Mode           Value                                   Value          Preempt
-------------------------------------------------------------------------------
non-revertive  15000          04/26/2021 09:21:41      0              Disabled
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
```

PE-5 becomes the DF for "vESI-45_2" in VPLS 2, as follows:

```
[/]
A:admin@PE-5# show service id 2 ethernet-segment
No sap entries

===============================================================================
SDP Ethernet-Segment Information
===============================================================================
SDP                  Eth-Seg                         Status
-------------------------------------------------------------------------------
56:2                 vESI-45_2                       DF
===============================================================================
No vxlan instance entries
```

```
[/]
A:admin@PE-5# show service system bgp-evpn ethernet-segment name "vESI-45_2"
                                          | match "DF Pref Election" post-lines 6
DF Pref Election Information
-------------------------------------------------------------------------------
Preference     Preference     Last Admin Change        Oper Pref      Do No
Mode           Value                                   Value          Preempt
-------------------------------------------------------------------------------
non-revertive  20000          04/26/2021 09:21:50      20000          Enabled
-------------------------------------------------------------------------------
```

When the SDP is restored, the DF does not revert even though the list of DF candidates contains both PE-4 and PE-5. The preference mode is non-revertive; therefore, the DP bit has been set. PE-4 will not become the DF, as follows:

```
# on MTU-6:
configure {
    service {
        sdp 64 {
            admin-state enable
```

```
[/]
A:admin@PE-4# show service system bgp-evpn ethernet-segment name "vESI-45_2"
                                                                      evi evi-1 2

===============================================================================
EVI DF and Candidate List
===============================================================================
EVI           SvcId         Actv Timer Rem      DF  DF Last Change
-------------------------------------------------------------------------------
2             2             0                   no  04/26/2021 09:41:43
===============================================================================

===============================================================================
DF Candidates                           Time Added
-------------------------------------------------------------------------------
192.0.2.4                               04/26/2021 09:43:17
192.0.2.5                               04/26/2021 09:40:43
-------------------------------------------------------------------------------
Number of entries: 2
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The operational preference value on NDF PE-4 equals the preference value on DF PE-5, as follows. In this example, EVI 2 is included in the configured EVI range, so the lowest preference wins. To avoid the system reverting to the lower preference of 15000, the operational preference is raised to the value of 20000, which equals the preference of the current DF PE-5.

```
[/]
A:admin@PE-4# show service system bgp-evpn ethernet-segment name "vESI-45_2"
                                            | match "DF Pref Election" post-lines 6
DF Pref Election Information
-------------------------------------------------------------------------------
Preference     Preference     Last Admin Change        Oper Pref      Do No
Mode           Value                                   Value          Preempt
-------------------------------------------------------------------------------
non-revertive  15000          04/26/2021 09:21:41      20000          Disabled
-------------------------------------------------------------------------------
```

PE-4 checks its own administrative preference and compares it with the one of the Highest-PE and Lowest-PE that have DP=1 in their ES routes.

- The Highest-PE is the PE with higher preference, using the DP bit (with DP=1 being better) and, after that, the lower PE-IP address as tie-breakers.
- The Lowest-PE is the PE with lower preference, using the DP bit (with DP=1 being better) and, after that, the lower PE-IP address as tie-breakers.

Depending on this comparison, PE-4 will send the ES route with a preference and DP that may be different from its administrative values.

- If PE-4's preference value is higher than the Highest-PE's, PE-4 will send the ES route with an 'in-use' operational preference equal to the Highest-PE's and DP=0.
- If PE-4's preference value is lower than the Lowest-PE's, PE-4 will send the ES route with an 'in-use' operational preference equal to the Lowest-PE's and DP=0.
- If PE-4's preference value is neither higher nor lower than the Highest-PE's or the Lowest-PE's respectively, PE-4 will send the ES route with its administrative [preference,DP]=[15000,1].

In this example, NDF PE-4 sends operational preference 20000 and DP=0, because its admin preference value was lower than the Lowest-PE's (PE-5), as follows:

```
# on PE-4:
148 2021/04/26 09:43:17.492 CEST MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 71
    Flag: 0x90 Type: 14 Len: 34 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.4
        Type: EVPN-ETH-SEG Len: 23 RD: 192.0.2.4:0
              ESI: 01:00:00:00:00:45:02:00:00:01, IP-Len: 4 Orig-IP-Addr: 192.0.2.4
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        df-election::DF-Type:Preference/DP:0/DF-Preference:20000/AC:1
        target:00:00:00:00:45:02
"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

With equal operational preference, the current DF PE-5 sends DP=1, which is preferred over DP=0. The following output shows the BGP extended community of the ES routes for "vESI-45_2" in the RIB-In (received ES route from PE-5) and RIB-Out (sent ES route) on PE-4:

```
[/]
A:admin@PE-4# show router bgp routes evpn eth-seg hunt
                                    | match "target:00:00:00:00:45:02" pre-lines 2
Community      :                                  # in RIB-In
                  df-election::DF-Type:Preference/DP:1/DF-Preference:20000/AC:1
                  target:00:00:00:00:45:02
Community      :                                  # in RIB-Out
                  df-election::DF-Type:Preference/DP:0/DF-Preference:20000/AC:1
                  target:00:00:00:00:45:02
```

Either of the following events cause PE-4 to re-advertise its admin preference 15000 and DP=1:

- DF PE-5 withdraws its ES route.
- The admin preference for ES "vESI-45_2" on DF PE-5 is modified by configuration to a value preferred over PE-4's admin preference; in this case, to a value lower than 15000.

The admin preference value can be modified on ES "vESI-45_2" on DF PE-5, as follows:

```
# on PE-5:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "vESI-45_2" {
                        df-election {
                            manual {
                                preference {
                                    mode non-revertive
                                    value 10000
                                }
                            }
```

The preference value 10000 is lower than 15000 and, therefore, preferred when the lowest preference wins. PE-5 remains DF, but now there is no need to modify the preference of PE-4, because the system does not need to revert. Therefore, PE-4 can send the admin preference 15000 and configured DP=1, as follows:

```
# on PE-4:
151 2021/04/26 09:47:45.433 CEST MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 71
    Flag: 0x90 Type: 14 Len: 34 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.4
        Type: EVPN-ETH-SEG Len: 23 RD: 192.0.2.4:0
               ESI: 01:00:00:00:00:45:02:00:00:01, IP-Len: 4 Orig-IP-Addr: 192.0.2.4
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        df-election::DF-Type:Preference/DP:1/DF-Preference:15000/AC:1
        target:00:00:00:00:45:02
"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Conclusion

Preference-based DF election offers more control over the DF Election and applies to regular ESs and vESs, either in single-active or in all-active multi-homing mode, in VPLS, I-VPLS, or Epipe services. The DF election is by default revertive, but when preference mode is chosen, it can be configured as non-revertive to reduce service impact.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Proxy-ARP/ND MAC List for Dynamic Entries

This chapter provides information about Proxy-ARP/ND MAC List for Dynamic Entries.

Topics in this chapter include:

Applicability

This chapter was initially written based on SR OS Release 15.0.R4, but the MD-CLI in the current edition is based on SR OS Release 21.2.R2. Proxy-Address Resolution Protocol/Neighbor Discovery (proxy-ARP/ND) MAC list for dynamic entries is supported in SR OS Release 15.0.R1, and later.

Overview

In some EVPN networks, the use of static proxy-ARP/ND entries is preferred to dynamically learned entries. For example, this is the case with some Internet eXchange Points (IXPs) that use EVPN and proxy-ARP/ND technologies. The MAC address in the static entry can be a MAC address from a list of n preregistered MAC addresses. The advantage is that—in case of a router or card failure—the hardware can be replaced, and no reconfiguration is required if the new MAC address is within a list of allowed MAC addresses.

In SR OS, these allow lists are called MAC lists. The associated proxy-ARP/ND entries will not be added upon configuration, but dynamically through a resolve procedure. This follows *draft-ietf-bess-evpn-proxy-arp-nd*.

- When the dynamic proxy-ARP/ND IP address with its associated MAC list is configured, the system sends a resolve message to all its non-EVPN peers.
- The resolve message is an ARP request for IPv4, or a Neighbor Solicitation (NS) message for IPv6.
- The resolve message is sent at a configurable interval between 1 and 60 minutes; the default is 5 minutes.
- The system keeps sending resolve messages until a dynamic entry is created for the proxy-ARP/ND IP address. This entry is only created when two conditions are met:
  - An ARP/Gratuitous Address Resolution Protocol (GARP) or Neighbor Advertisement (NA) is received for the configured IP address.
  - The associated MAC address belongs to the MAC list configured for the IP address. If the MAC list is empty or not configured, the system will never create an entry for the IP address.

When the dynamic proxy-ARP/ND IP entry is created, the system advertises an EVPN-MAC update to its EVPN peers. The sticky bit will be set depending on how the corresponding MAC address is learned. If the MAC address is learned on a SAP/SDP-binding with Auto-Learn MAC Protect (ALMP) enabled, the EVPN-MAC route will be advertised as static.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 265: IXP with proxy-ARP/ND MAC list for dynamic entries shows an example of an IXP network that uses proxy-ARP/ND and a MAC list.

Figure 265: IXP with proxy-ARP/ND MAC list for dynamic entries

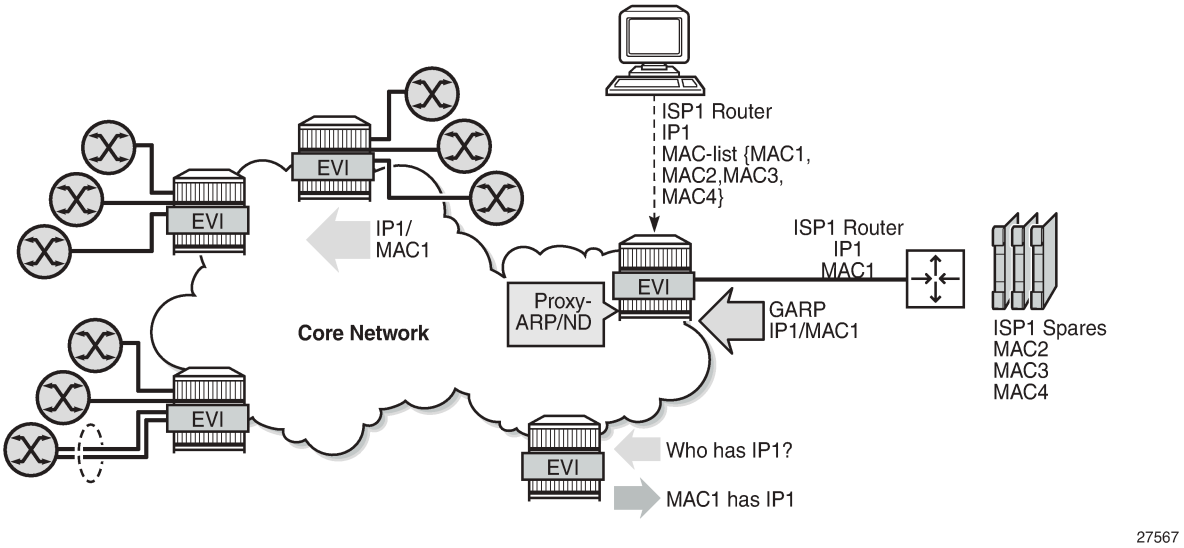

The ISP1 router with IP1 and MAC1 is connected to a PE in the core network that has proxy-ARP/ND enabled and a list of allowed MAC addresses. This MAC list contains four MAC addresses: MAC1 (for the hardware that is currently in use) and three MAC addresses for spares: MAC2, MAC3, and MAC4. The proxy-ARP/ND table will be populated as follows:

- The PE floods a resolve message for the configured IP address for proxy-ARP/ND to its non-EVPN peers.
- The ISP1 router that is connected to the network sends a GARP or ARP Reply message with IP1 and MAC1 that will be snooped by the PE.
- The PE checks whether IP1 is configured as a dynamic proxy-ARP/ND entry and MAC1 is in the MAC list assigned to proxy-ARP/ND entry IP1.
  - If true, the IP1/MAC1 entry is created in the proxy-ARP/ND table and advertised in EVPN.
  - If the GARP message contains MAC5, which is not in the MAC allow list, no proxy-ARP/ND entry is created, and IP/MAC is not advertised. If **proxy-arp>evpn>flood>gratuitous-arp false** is configured, the GARP containing MAC5 will be discarded.

If after the proxy-ARP/ND creation, the corresponding MAC address is flushed from the Forwarding Database (FDB), the entry goes inactive. After the age-time, the inactive entry will age out and the resolve process will restart.

MAC lists are configured with the following command:

```
[ex:/configure service proxy-arp-nd mac-list]
A:admin@PE-2# list "ISP1" ?

 list

 apply-groups          - Apply a configuration group at this level
 apply-groups-exclude  - Exclude a configuration group at this level
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
 mac                   - Add a list entry for mac
```

The MAC list contains the allowed MAC addresses and can be associated in one or more services with a proxy-ARP/ND IP address. A MAC list is associated with dynamic proxy-ARP IP 1.1.1.1 with the following command:

```
[ex:/configure service vpls "EVI-1" proxy-arp dynamic-arp]
A:admin@PE-2# ip-address 1.1.1.1 ?

 ip-address

 apply-groups          - Apply a configuration group at this level
 apply-groups-exclude  - Exclude a configuration group at this level
 mac-list              - MAC list for the dynamic entry
 resolve-retry-time    - Frequency at which the resolve messages are sent
```

The configuration for proxy-ND is similar:

```
[ex:/configure service vpls "EVI-1" proxy-nd dynamic-neighbor]
A:admin@PE-2# ip-address 2001:db8::99 ?

 ip-address

 apply-groups          - Apply a configuration group at this level
 apply-groups-exclude  - Exclude a configuration group at this level
 mac-list              - MAC list for the dynamic entry
 resolve-retry-time    - Frequency at which the resolve messages are sent
```

- The MAC list can be associated with multiple configured dynamic IP addresses:
  - In different services
  - In the same service, for proxy-ARP and proxy-ND
- An empty MAC list can be configured and applied, but no proxy-ARP/ND entries will be created when the PE receives a GARP message containing a MAC address that is not in the allow list.
- MAC lists can be modified at any time: MAC addresses can be added or removed even when the MAC lists are associated with configured dynamic IP addresses. If the MAC list changes, all the IP addresses associated with that MAC list will delete the proxy entries and restart the resolve process.

An existing dynamic proxy-ARP/ND entry IP1/MAC1 can be overridden when the system receives a GARP/ARP/NA for IP1 with another MAC address from the MAC list (IP1/MAC2). The system will first send a confirm message to check whether the old IP1/MAC1 is still reachable. Only when there is no answer, the entry IP1/MAC1 is replaced by IP1/MAC2. The existing duplicate-detect and confirm procedures are only applied for MAC address changes within the MAC list.

An existing dynamic proxy-ARP/ND entry IP1/MAC1 will be deleted when the system receives a GARP/ARP/NA IP1/MAC5 with a MAC address that is not contained in the MAC list. The GARP/ARP/NA message will be discarded and the resolve procedure is restarted.

Configuration

Figure 266: Example topology shows the example topology with three PEs. ISP router 1 is connected to PE-2. MAC1 is used; MAC2, MAC3, and MAC4 correspond to spares.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 266: Example topology

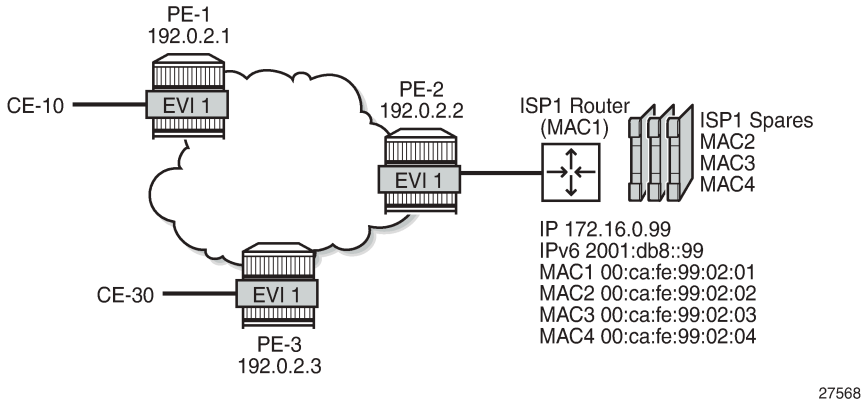

The initial configuration includes:

- Cards, MDAs, ports
- Router interfaces
- IS-IS between the PEs (alternatively, OSPF can be used)
- LDP between the PEs

BGP is enabled between the PEs for address family EVPN. The BGP configuration on PE-2 is as follows:

```
# on PE-2:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            rapid-withdrawal true
            split-horizon true
            rapid-update {
                evpn true
            }
            group "internal" {
                peer-as 64500
                family {
                    evpn true
                }
            }
            neighbor "192.0.2.1" {
                group "internal"
            }
            neighbor "192.0.2.3" {
                group "internal"
            }
        }
```

VPLS 1 is configured on PE-2 as follows. The configuration on the other PEs is similar.

```
# on PE-2:
configure {
    service {
        vpls "EVI-1" {
            admin-state enable
            service-id 1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            customer "1"
            bgp 1 {
            }
            bgp-evpn {
                evi 1
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            sap 1/2/1:1 {
            }
            sap 1/2/1:3 {
            }
        }
```

MAC list

The following MAC lists are configured on PE-2: ISP1 is an empty list; ISP2 is a MAC list containing four MAC addresses.

```
# on PE-2:
configure {
    service {
        proxy-arp-nd {
            mac-list {
                list "ISP1" {
                }
                list "ISP2" {
                    mac 00:ca:fe:99:02:01 { }
                    mac 00:ca:fe:99:02:02 { }
                    mac 00:ca:fe:99:02:03 { }
                    mac 00:ca:fe:99:02:04 { }
                }
            }
```

The following command shows the configured MAC lists on PE-2, with the number of MAC addresses and the number of associations. None of the MAC lists has been associated with a proxy-ARP/ND IP entry, so the number of associations is zero.

```
[/]
A:admin@PE-2# show service proxy-arp-nd mac-list

===============================================================================
MAC List Information
===============================================================================
MAC List Name                       Last Change         Num Macs    Num Assocs
-------------------------------------------------------------------------------
ISP1                                05/11/2021 13:58:23  0           0
ISP2                                05/11/2021 14:03:41  4           0
-------------------------------------------------------------------------------
Number of Entries: 2
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The following command shows the MAC addresses that are configured in MAC list ISP2. The timestamps show that all four MAC addresses were configured simultaneously, but MAC lists can be modified at any time.

```
[/]
A:admin@PE-2# show service proxy-arp-nd mac-list name "ISP2"

===============================================================================
MAC List MAC Addr Information
===============================================================================
MAC Addr                                Last Change
-------------------------------------------------------------------------------
00:ca:fe:99:02:01                       05/11/2021 14:03:41
00:ca:fe:99:02:02                       05/11/2021 14:03:41
00:ca:fe:99:02:03                       05/11/2021 14:03:41
00:ca:fe:99:02:04                       05/11/2021 14:03:41
-------------------------------------------------------------------------------
Number of Entries: 4
===============================================================================
```

MAC list associated with proxy-ARP/ND in VPLS

MAC lists can be associated with one or more services. An empty MAC list—such as ISP1—can be associated, but it is impossible to associate a non-existing MAC list with a service. The following error is raised when attempting to associate the non-existing MAC list ISP3 with proxy-ARP IP 1.1.1.1 in VPLS 1 on PE-2:

```
*[ex:/configure service vpls "EVI-1" proxy-arp dynamic-arp ip-address 1.1.1.1]
A:admin@PE-2# mac-list "ISP3"

*[ex:/configure service vpls "EVI-1" proxy-arp dynamic-arp ip-address 1.1.1.1]
A:admin@PE-2# commit
MINOR: MGMT_CORE #224: configure service vpls "EVI-1" proxy-arp dynamic-arp ip-address 1.1.1.1
 mac-list - Entry does not exist - configure service proxy-arp-nd mac-list list "ISP3"
```

MAC list ISP2 is associated with proxy-ARP IP 172.16.0.99 and with proxy-ND IP 2001:db8::99 in VPLS 1 on PE-2, as follows:

```
# on PE-2:
configure {
    service {
        vpls "EVI-1" {
            proxy-arp {
                admin-state enable
                dynamic-populate true
                dynamic-arp {
                    ip-address 172.16.0.99 {
                        mac-list "ISP2"
                        resolve-retry-time 1
                    }
                }
            }
            proxy-nd {
                admin-state enable
                dynamic-populate true
                dynamic-neighbor {
                    ip-address 2001:db8::99 {
                        mac-list "ISP2"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    }
                }
            }
```

For proxy-ARP IP 172.16.0.99, the resolve interval is 1 minute, which is the minimum; for proxy-ND IP 2001:db::99, the resolve interval is the default of 5 minutes. In scaled environments, Nokia recommends using the default interval, or even configuring a longer interval. The proxy-ARP and proxy-ND tables can be populated with dynamic entries (**dynamic-populate true**).

The following command shows all associations for MAC list ISP2: two associations are defined in VPLS 1: one for IP address 172.16.0.99 and another for IP address 2001:db8::99.

```
[/]
A:admin@PE-2# show service proxy-arp-nd mac-list name "ISP2" associations

===============================================================================
MAC List Associations
===============================================================================
Service Id                  IP Addr
-------------------------------------------------------------------------------
1                           172.16.0.99
1                           2001:db8::99
-------------------------------------------------------------------------------
Number of Entries: 2
===============================================================================
```

Different dynamic proxy-ARP/ND entries

A distinction is made between regular dynamic entries and configured dynamic entries:

- No IP address needs to be configured for regular dynamic proxy-ARP/ND entries. What only needs to be configured, is the option **dynamic-populate true**.
- IP address and MAC list need to be defined for configured proxy-ARP/ND entries.

Configured dynamic entries can override static and regular dynamic entries.

Regular dynamic proxy-ARP/ND entries can override configured dynamic entries.

EVPN entries cannot override configured dynamic entries, even though they can override regular dynamic entries.

Likewise, static entries can override regular dynamic entries, but they cannot override dynamic configured entries. The following error is raised when attempting to configure a static proxy-ARP entry for IP 172.16.0.99, which has already been configured as dynamic and associated with a MAC list.

```
*[ex:/configure service vpls "EVI-1" proxy-arp static-arp ip-address 172.16.0.99]A:admin@PE-2#
 commit
MINOR: MGMT_CORE #258: configure service vpls "EVI-1" proxy-arp dynamic-arp ip-address
 172.16.0.99 - Unique values required - configure service vpls "EVI-1" proxy-arp static-arp ip-
address 172.16.0.99
```

Debugging

Debugging for both proxy-ARP/ND IP entries is enabled—in classic CLI—on PE-2 as follows:

```
# on PE-2:
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
debug
    service
        id 1
            proxy-arp ip 172.16.0.99
            proxy-nd ip 2001:db8::99
        exit
    exit
exit
```

When the dynamic proxy-ARP IP 172.16.0.99 is configured with MAC list "ISP2", PE-2 floods a resolve message—in this case, an ARP request—to all its EVPN peers. Router ISP1 replies. PE-2 advertises an EVPN-MAC update to its EVPN peers PE-1 and PE-3. PE-2 adds a dynamic proxy-ARP entry for 172.16.0.99 with MAC address 00:ca:fe:99:02:01. Router ISP1 sends a GARP message. The following messages are logged:

```
29 2021/05/11 14:11:39.920 CEST MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.99 flood resolve"

31 2021/05/11 14:11:39.922 CEST MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.99 mac: 00:ca:fe:99:02:01 evpn advertise"

32 2021/05/11 14:11:39.922 CEST MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.99 type: Dyn mac: 00:ca:fe:99:02:01 Added"

37 2021/05/11 14:11:40.020 CEST MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.99 type: Dyn mac: 00:ca:fe:99:02:01 Gratuitous Update"
```

For proxy-ND, the following messages are logged:

```
30 2021/05/11 14:11:39.920 CEST MINOR: DEBUG #2001 Base proxy nd
"proxy nd:
svc: 1 ip: 2001:db8::99 flood resolve"

33 2021/05/11 14:11:39.922 CEST MINOR: DEBUG #2001 Base proxy nd
"proxy nd:
svc: 1 ip: 2001:db8::99 mac: 00:ca:fe:99:02:01 evpn advertise"

34 2021/05/11 14:11:39.922 CEST MINOR: DEBUG #2001 Base proxy nd
"proxy nd:
svc: 1 ip: 2001:db8::99 type: Dyn mac: 00:ca:fe:99:02:01 Added"

38 2021/05/11 14:11:40.020 CEST MINOR: DEBUG #2001 Base proxy nd
"proxy nd:
svc: 1 ip: 2001:db8::99 type: Dyn mac: 00:ca:fe:99:02:01 Gratuitous Update"
```

The following command shows the proxy-ARP details for VPLS 1 on PE-2. The only proxy-ARP entry is for IP address 172.16.0.99 with MAC address 00:ca:fe:99:02:01.

```
[/]
A:admin@PE-2# show service id 1 proxy-arp detail
-------------------------------------------------------------------------------
Proxy Arp
-------------------------------------------------------------------------------
Admin State       : enabled
Dyn Populate      : enabled
Age Time          : disabled            Send Refresh      : disabled
Table Size        : 250                 Total             : 1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Static Count      : 0                   EVPN Count        : 0
Dynamic Count     : 1                   Duplicate Count   : 0

Dup Detect
-------------------------------------------------------------------------------
Detect Window     : 3 mins              Num Moves         : 5
Hold down         : 9 mins
Anti Spoof MAC    : None

EVPN
-------------------------------------------------------------------------------
Garp Flood        : enabled             Req Flood         : enabled
Static Black Hole : disabled
EVPN Route Tag    : 0
-------------------------------------------------------------------------------

===============================================================================
VPLS Proxy Arp Entries
===============================================================================
IP Address          Mac Address         Type      Status    Last Update
-------------------------------------------------------------------------------
172.16.0.99         00:ca:fe:99:02:01   dyn       active    05/11/2021 14:11:40
-------------------------------------------------------------------------------
Number of entries : 1
===============================================================================
```

The following command shows the proxy-ND details for VPLS 1 on PE-2. The only proxy-ND entry if for IP address 2001:db8::99 with MAC address 00:ca:fe:99:02:01.

```
[/]
A:admin@PE-2# show service id 1 proxy-nd detail
-------------------------------------------------------------------------------
Proxy ND
-------------------------------------------------------------------------------
Admin State       : enabled
Dyn Populate      : enabled
Age Time          : disabled            Send Refresh      : disabled
Table Size        : 250                 Total             : 1
Static Count      : 0                   EVPN Count        : 0
Dynamic Count     : 1                   Duplicate Count   : 0

Dup Detect
-------------------------------------------------------------------------------
Detect Window     : 3 mins              Num Moves         : 5
Hold down         : 9 mins
Anti Spoof MAC    : None

EVPN
-------------------------------------------------------------------------------
Unknown NS Flood  : enabled             ND Advertise      : Router
Rtr Unsol NA Flood: enabled             Host Unsol NA Fld : enabled
EVPN Route Tag    : 0
-------------------------------------------------------------------------------

===============================================================================
VPLS Proxy ND Entries
===============================================================================
IP Address               Mac Address       Type Status Rtr/ Last Update
                                                       Host
-------------------------------------------------------------------------------
2001:db8::99             00:ca:fe:99:02:01 dyn  active Rtr  05/11/2021 14:11:40
-------------------------------------------------------------------------------
Number of entries : 1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
```

The proxy-ARP in VPLS 1 contains the following dynamic entry.

```
[/]
A:admin@PE-2# show service id 1 proxy-arp dynamic

===============================================================================
Proxy ARP Dyn Cfg Summary
===============================================================================
IP Addr                                 Mac List
-------------------------------------------------------------------------------
172.16.0.99                             ISP2
-------------------------------------------------------------------------------
Number of Entries: 1
===============================================================================
```

The following command shows the association for dynamic proxy-ARP IP address 172.16.0.99, with the configured resolve time in minutes and the remaining resolve time in seconds.

```
[/]
A:admin@PE-2# show service id 1 proxy-arp dynamic ip-address 172.16.0.99

===============================================================================
Proxy ARP Dyn Cfg Detail
===============================================================================
IP Addr        Mac List                         Resolve Time   Remaining
                                                (mins)          Resolve Time
                                                               (secs)
-------------------------------------------------------------------------------
172.16.0.99    ISP2                             1              0
-------------------------------------------------------------------------------
Number of Entries: 1
===============================================================================
```

The remaining resolve time is zero seconds because a dynamic proxy-ARP entry has been created and that suspends the resolve mechanism.

The proxy-ND in VPLS 1 contains the following dynamic entry.

```
[/]
A:admin@PE-2# show service id 1 proxy-nd dynamic

===============================================================================
Proxy ND Dyn Cfg Summary
===============================================================================
IP Addr                                 Mac List
-------------------------------------------------------------------------------
2001:db8::99                            ISP2
-------------------------------------------------------------------------------
Number of Entries: 1
===============================================================================
```

The following command shows the association for dynamic proxy-ND IP 2001:db8::99.

```
[/]
A:admin@PE-2# show service id 1 proxy-nd dynamic ipv6-address 2001:db8::99

===============================================================================
Proxy ND Dyn Cfg Detail
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
IP Addr                                 Mac List
 Resolve Time(mins)                      Remaining Resolve Time(secs)
-------------------------------------------------------------------------------
2001:db8::99                            ISP2
5                                       0
-------------------------------------------------------------------------------
Number of Entries: 1
===============================================================================
```

Tools command to trigger resolve procedure

The following tools command can be used to force the system to send a resolve message to its non-EVPN peers. The **force** option will trigger the resolve process even for existing entries in the proxy-ARP/ND table.

```
[/]
A:admin@PE-2# tools perform service id 1 proxy-arp dynamic-resolve ?

 dynamic-resolve all [force]
 dynamic-resolve <IP address> [force]

 [ip-address] (<ipv4-address> | <ipv6-address>)
 <ipv4-address>  - <d.d.d.d>
 <ipv6-address>  - (<x:x:x:x:x:x:x:x>|<x:x:x:x:x:x:d.d.d.d>)

 [ip-address]          - ipv4 address '<d.d.d.d>' or ipv6 address
 '(<x:x:x:x:x:x:x:x>|<x:x:x:x:x:x:d.d.d.d>)'
 all                   - <keyword>
 force                 - <keyword>
```

```
[/]
A:admin@PE-2# tools perform service id 1 proxy-nd dynamic-resolve ?

 dynamic-resolve all [force]
 dynamic-resolve <ipv6 address> [force]

 [ipv6-address] <ipv6-address>
 x:x:x:x:x:x:x:x   (eight 16-bit pieces)
 x:x:x:x:x:x:d.d.d.d
 x - [0..FFFF]H
 d - [0..255]D


    Attribute ipv6-address for dynamic-resolve

 [ipv6-address]        - Attribute ipv6-address for dynamic-resolve
 all                   - <keyword>
 force                 - <keyword>
```

Some examples:

```
[/]
A:admin@PE-2# tools perform service id 1 proxy-arp dynamic-resolve 172.16.0.99

[/]
A:admin@PE-2# tools perform service id 1 proxy-arp dynamic-resolve 172.16.0.99
                                                                        force

[/]
A:admin@PE-2# tools perform service id 1 proxy-arp dynamic-resolve all
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
[/]
A:admin@PE-2# tools perform service id 1 proxy-arp dynamic-resolve all force

[/]
A:admin@PE-2# tools perform service id 1 proxy-nd dynamic-resolve 2001:db8::99

[/]
A:admin@PE-2# tools perform service id 1 proxy-nd dynamic-resolve 2001:db8::99
                                                                          force

[/]
A:admin@PE-2# tools perform service id 1 proxy-nd dynamic-resolve all

[/]
A:admin@PE-2# tools perform service id 1 proxy-nd dynamic-resolve all force
```

Inactive proxy-ARP/ND entries

When the MAC address is flushed from the FDB, the proxy-ARP/ND entries become inactive.

```
[/]
A:admin@PE-2# clear service id 1 fdb mac 00:ca:fe:99:02:01
```

```
[/]
A:admin@PE-2# show service id 1 proxy-arp detail | match 172.16.0.99 pre-lines 6
                                                                    post-lines 3
-------------------------------------------------------------------------------
===============================================================================
VPLS Proxy Arp Entries
===============================================================================
IP Address          Mac Address         Type      Status    Last Update
-------------------------------------------------------------------------------
172.16.0.99         00:ca:fe:99:02:01   dyn       inActv    05/11/2021 14:16:37
-------------------------------------------------------------------------------
Number of entries : 1
===============================================================================
```

```
[/]
A:admin@PE-2# show service id 1 proxy-nd detail | match 2001:db8::99 pre-lines 7
                                                                    post-lines 3
-------------------------------------------------------------------------------
===============================================================================
VPLS Proxy ND Entries
===============================================================================
IP Address              Mac Address       Type Status Rtr/ Last Update
                                                      Host
-------------------------------------------------------------------------------
2001:db8::99            00:ca:fe:99:02:01 dyn  inActv Rtr  05/11/2021 14:16:37
-------------------------------------------------------------------------------
Number of entries : 1
===============================================================================
```

By default, aging is disabled, and the entries remain in the inactive status until the MAC address is learned again. However, if aging is enabled, the inactive proxy-ARP/ND entry will age out. After the entry is

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

deleted, the system sends a resolve message. When the ISP1 router replies, the entry is created again in the proxy-ARP/ND table. The age time is configured in seconds with the following command:

```
[ex:/configure service vpls "EVI-1" proxy-arp]
A:admin@PE-2# age-time ?

 age-time (<number> | <keyword>)
 <number>   - <60..86400>  - seconds
 <keyword>  - never        - seconds
 Default    - never

    Aging timer for proxy entries, where entries are flushed upon timer expiry
```

```
# on PE-2:
configure {
    service {
        vpls "EVI-1" {
            proxy-arp {
                age-time 60
            }
            proxy-nd {
                age-time 60
            }
```

The following debug messages for proxy ARP IP 172.16.0.99 show that an EVPN-MAC withdraw message is sent (when the MAC address is flushed from the FDB) and—after time-out—the proxy-ARP entry is deleted. PE-2 sends a resolve message to all its non-EVPN peers. Router ISP1 replies and the proxy-ARP entry is created again; an EVPN-MAC update is sent to the EVPN peers. Similar debug messages occur for proxy-ND.

```
57 2021/05/11 14:16:48.589 CEST MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.99 mac: 00:ca:fe:99:02:01 evpn withdraw"

62 2021/05/11 14:18:33.620 CEST MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.99 type: Dyn mac: 00:ca:fe:99:02:01 Deleted"

64 2021/05/11 14:18:33.720 CEST MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.99 flood resolve"

65 2021/05/11 14:18:33.722 CEST MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.99 mac: 00:ca:fe:99:02:01 evpn advertise"

66 2021/05/11 14:18:33.722 CEST MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.99 type: Dyn mac: 00:ca:fe:99:02:01 Added"

71 2021/05/11 14:18:33.820 CEST MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.99 type: Dyn mac: 00:ca:fe:99:02:01 Gratuitous Update"
```

The following command shows that the entry is created again with active status.

```
[/]
A:admin@PE-2# show service id 1 proxy-arp detail | match 172.16.0.99 pre-lines 6
                                                                    post-lines 3
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
VPLS Proxy Arp Entries
===============================================================================
IP Address          Mac Address         Type      Status    Last Update
-------------------------------------------------------------------------------
172.16.0.99         00:ca:fe:99:02:01   dyn       active    05/11/2021 14:19:34
-------------------------------------------------------------------------------
Number of entries : 1
===============================================================================
```

MAC address replacement

When the system receives a GARP/ARP/NA for the same IP address, but with another MAC address from the MAC list, it will first send a confirm message to ensure that the old MAC address is not used anymore for the IP address. If the existing proxy-ARP/ND entry is IP1/MAC1 and a GARP/ARP/NA message is received for IP1/MAC4, the system sends an EVPN-MAC withdraw message for MAC1 and changes MAC1 to MAC4 for proxy-ARP/ND IP1, but the status is pending (pendng), as follows:

```
[/]A:admin@PE-2# show service id 1 proxy-arp detail | match 172.16.0.99 pre-lines 6
                                                                      post-lines 3
-------------------------------------------------------------------------------
===============================================================================
VPLS Proxy Arp Entries
===============================================================================
IP Address          Mac Address         Type      Status    Last Update
-------------------------------------------------------------------------------
172.16.0.99         00:ca:fe:99:02:04   dyn       pendng    05/11/2021 14:23:32
-------------------------------------------------------------------------------
Number of entries : 1
===============================================================================
```

```
[/]
A:admin@PE-2# show service id 1 proxy-nd detail | match 2001:db8::99 pre-lines 7
                                                                      post-lines 3
-------------------------------------------------------------------------------
===============================================================================
VPLS Proxy ND Entries
===============================================================================
IP Address               Mac Address       Type Status Rtr/ Last Update
                                                       Host
-------------------------------------------------------------------------------
2001:db8::99             00:ca:fe:99:02:04 dyn  pendng Rtr  05/11/2021 14:23:31
-------------------------------------------------------------------------------
Number of entries : 1
===============================================================================
```

The system sends a confirm message (unicast ARP request) for the old entry IP1/MAC1 to ensure that there is no duplication. When there is no reply from MAC1, there is no duplication. An EVPN-MAC route is advertised for MAC4. The status of the proxy-ARP entry IP1/MAC4 changes to active. The following debug messages are logged for proxy-ARP 172.16.0.99:

```
151 2021/05/11 14:23:29.394 CEST MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.99 mac: 00:ca:fe:99:02:01 evpn withdraw"

152 2021/05/11 14:23:29.394 CEST MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.99 Mac Change: 00:ca:fe:99:02:01->00:ca:fe:99:02:04 "
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
157 2021/05/11 14:23:29.520 CEST MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.99 mac: 00:ca:fe:99:02:01 confirm"

160 2021/05/11 14:23:59.520 CEST MINOR: DEBUG #2001 Base proxy arp
"proxy arp:
svc: 1 ip: 172.16.0.99 mac: 00:ca:fe:99:02:04 evpn advertise"
```

The final status of the proxy-ARP IP 172.16.0.99 is active, as follows:

```
[/]
A:admin@PE-2# show service id 1 proxy-arp detail | match 172.16.0.99 pre-lines 6
                                                                    post-lines 3
-------------------------------------------------------------------------------
===============================================================================
VPLS Proxy Arp Entries
===============================================================================
IP Address          Mac Address         Type      Status    Last Update
-------------------------------------------------------------------------------
172.16.0.99         00:ca:fe:99:02:04   dyn       active    05/11/2021 14:24:34
-------------------------------------------------------------------------------
Number of entries : 1
===============================================================================
```

The mechanism is similar for proxy-ND.

The behavior is different when the system receives a GARP/ARP/NA for the IP address with a MAC address that is not contained in the MAC list. The GARP/ARP/NA message is discarded and the proxy-ARP/ND entry deleted. The resolve procedure gets restarted.

Modified MAC list

MAC lists can be modified at any time, as follows:

```
# on PE-2:
configure {
    service {
        proxy-arp-nd {
            mac-list {
                list "ISP2" {
                    mac 00:ca:fe:99:02:05 { }
```

```
[/]A:admin@PE-2# show service proxy-arp-nd mac-list name "ISP2"

===============================================================================
MAC List MAC Addr Information
===============================================================================
MAC Addr                                Last Change
-------------------------------------------------------------------------------
00:ca:fe:99:02:01                       05/11/2021 14:03:41
00:ca:fe:99:02:02                       05/11/2021 14:03:41
00:ca:fe:99:02:03                       05/11/2021 14:03:41
00:ca:fe:99:02:04                       05/11/2021 14:03:41
00:ca:fe:99:02:05                       05/11/2021 14:25:23
-------------------------------------------------------------------------------
Number of Entries: 5
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The timestamps show when the different MAC addresses were added to the MAC list.

When the MAC list ISP2 is modified, proxy-ARP entry 172.16.0.99 and proxy-ND entry 2001:db8::99 will be deleted, an EVPN-MAC withdraw message will be sent, and the resolve procedure will be restarted. The following log messages occur for proxy-ND 2001:db8::99.

```
182 2021/05/11 14:25:23.153 CEST MINOR: DEBUG #2001 Base proxy nd
"proxy nd:
svc: 1 ip: 2001:db8::99 mac: 00:ca:fe:99:02:04 evpn withdraw"

183 2021/05/11 14:25:23.153 CEST MINOR: DEBUG #2001 Base proxy nd
"proxy nd:
svc: 1 ip: 2001:db8::99 type: Dyn mac: 00:ca:fe:99:02:04 Deleted"

187 2021/05/11 14:25:23.320 CEST MINOR: DEBUG #2001 Base proxy nd
"proxy nd:
svc: 1 ip: 2001:db8::99 flood resolve"

190 2021/05/11 14:25:23.322 CEST MINOR: DEBUG #2001 Base proxy nd
"proxy nd:
svc: 1 ip: 2001:db8::99 mac: 00:ca:fe:99:02:04 evpn advertise"

191 2021/05/11 14:25:23.322 CEST MINOR: DEBUG #2001 Base proxy nd
"proxy nd:
svc: 1 ip: 2001:db8::99 type: Dyn mac: 00:ca:fe:99:02:04 Added"

195 2021/05/11 14:25:23.420 CEST MINOR: DEBUG #2001 Base proxy nd
"proxy nd:
svc: 1 ip: 2001:db8::99 type: Dyn mac: 00:ca:fe:99:02:04 Gratuitous Update"
```

Conclusion

MAC lists can be associated with configured dynamic proxy-ARP/ND IP addresses. The actual proxy entries will only be created after a GARP/ARP/NA message is received for the IP address and one of the MAC addresses from the MAC list.

This tool complements the SR OS EVPN proxy-ARP/ND solution for providers present at IXPs.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Static VXLAN Termination in Epipe Services

This chapter provides information about Static VXLAN Termination in Epipe Services.

Topics in this chapter include:

Applicability

This chapter was initially written for SR OS Release 15.0.R6, but the MD-CLI in the current edition is based on SR OS Release 21.5.R1. Static VXLAN termination for Epipe services is supported in SR OS Release 15.0.R1, and later.

Overview

Static Virtual eXtensible Local Area Network (VXLAN) termination on non-system IP addresses of the PEs is supported in VPLS services, as described in chapter VXLAN Forwarding Path Extension, and in Epipe services, as described in this chapter. Whereas VPLSs using VXLAN require BGP-EVPN control plane in the current release, Epipe services using VXLAN do not. This implies that only the configured values are used because no auto-discovery of the remote Termination Endpoints (TEPs) can be done without BGP-EVPN.

This chapter describes the configuration and use of static VXLAN as an access tunneling mechanism to a PBB-EVPN network. This is a design deployed in some service provider networks where the aggregation network is a non-MPLS IP network.

Static VXLAN termination for Epipe services can be applied on system IP addresses or non-system IP addresses.

Static VXLAN termination on system IP addresses

Figure 267: Static VXLAN termination on system IP addresses shows an example topology with three PEs and two CEs. Epipe 1 is configured on PE-1 and PE-2. PE-2 and PE-4 are part of a PBB-EVPN network. On PE-2, a port cross-connect (PXC) is configured to connect the SAP in Epipe 1 and the SAP in I-VPLS 101. CE-10 and CE-40 can send traffic to each other.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 267: Static VXLAN termination on system IP addresses

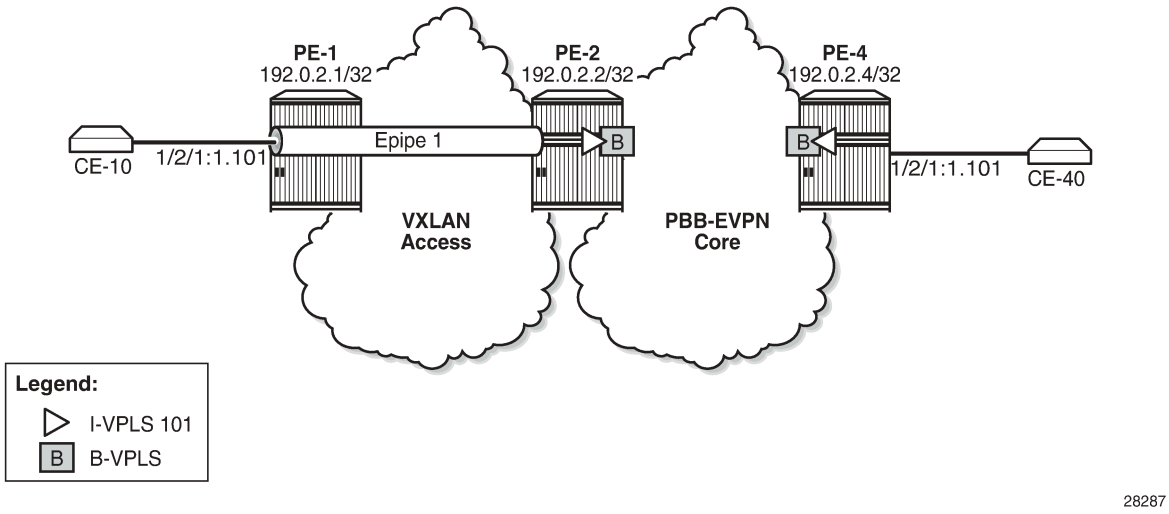

On PE-1, Epipe 1 is configured with egress VXLAN VNI 1, egress VXLAN Termination Endpoint (VTEP) 192.0.2.2, oper-group op-grp-1, and a SAP toward CE-10, as follows:

```
# on PE-1:
configure {
    service {
        oper-group "op-grp-1" {
        }
        epipe "Epipe 1" {
            admin-state enable
            service-id 1
            customer "1"
            sap 1/2/1:1.* {
            }
            vxlan {
                instance 1 {
                    vni 1
                    egress-vtep {
                        ip-address 192.0.2.2
                        oper-group "op-grp-1"
                    }
                }
            }
        }
```

where:

- The configured VXLAN Virtual Network Identifier (VNI) is used by the system as follows:
  - As the egress VNI when sending VXLAN packets for the Epipe service
  - As the source VNI that identifies the VXLAN packet to be part of the Epipe
  - Unique in the system, so it can only be configured in one service, either VPLS or Epipe

  The configuration of the VXLAN VNI in an Epipe is similar to the configuration of the VXLAN VNI in a VPLS, except that in a VPLS, the VNI is only used as the source VNI, because the egress VNI is learned from BGP-EVPN. However, in Epipe services with static VXLAN, the egress VNI is also the configured VNI.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- The egress VTEP is the system IP address of the remote PE. The system will add the configured egress VTEP IP address as the remote VTEP when encapsulating the frames into VXLAN packets. Only the egress VTEP is configured, not the source VTEP. The PE receiving VXLAN packets will not check the source VTEP.
- The egress VTEP IP address must be in the Routing Table Manager (RTM). An oper-group is associated with the egress VTEP IP address, so that when the egress VTEP disappears from the base route table, the oper-group is brought operationally down, which propagates the failure to other objects that have this oper-group associated. The status of the oper-group and the service will be as follows:
  - When the egress VTEP disappears from the RTM, the VXLAN binding goes operationally down and the oper-group associated with the egress VTEP goes operationally down.
  - When the Epipe SAP goes down, the service goes down too.
  - When the VXLAN binding goes down, the service remains up as long as the access SAP is up.
  - When the service is disabled, the VXLAN binding and the oper-group associated with the egress VTEP are both brought operationally down.
- Only SAPs can be associated with the Epipe; no spoke-SDPs are supported in SR OS Release 21.5.R1, as follows. Regular SAPs and PXC SAPs are supported.

```
*[ex:/configure service epipe "Epipe 1" spoke-sdp 11:1]
A:admin@PE-1# commit
MINOR: SVCMGR #12: configure service epipe "Epipe 1" vxlan instance 1 egress-vtep ip-address -
 Inconsistent Value error - vxlan-egr-vtep not supported with spoke-sdps in service - configure
 service epipe "Epipe 1" spoke-sdp 11:1
```

Frame encapsulation and forwarding

Incoming traffic in the PEs is treated as follows:

- For frames received from the SAPs, a SAP lookup identifies all frames matching the configured SAP (on PE-1, SAP 1/2/1:1.*). The matching frames will be encapsulated into VXLAN IPv4 packets with the following fields:
  - Source VTEP = system IP address
  - Destination VTEP = configured address in **egress-vtep**
  - VNI = configured VXLAN VNI
  - Source and destination UDP ports will be populated as per the existing VXLAN implementation VPLS services, with the source UDP port populated with the result of a hash on the ingress packets.
- For VXLAN frames received from the VXLAN network, a VNI lookup is done for packets with IP DA = system IP address. Frames with the configured VNI 1 are assigned to Epipe 1. The VXLAN encapsulation is removed and the frames are forwarded to the SAP.

Per-service hashing is not supported in Epipe-VXLAN services; only regular hashing and spraying in LAG/ECMP is supported as in any Epipe.

Static VXLAN termination on IPv6 or non-system IPv4 addresses

The non-system IPv4 or IPv6 VXLAN termination on Epipe services is configured in the same way as for VPLS services and described in the VXLAN Forwarding Path Extension chapter, using the FPE function for additional processing. The following steps are required for configuring the FPE for VXLAN termination:

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

1. Create FPE.
2. Associate FPE with VXLAN termination.
3. Configure the loopback router interface subnet for VXLAN termination and its advertisement into the routing protocol. The subnet can be IPv4 or IPv6.
4. Configure the loopback address for VXLAN termination.
5. Add the service configuration.

Configuration

In this section, static VXLAN termination for Epipe services is configured for the following cases:

- VXLAN termination on system IP addresses
- VXLAN termination on non-system IPv4 addresses
- VXLAN termination on IPv6 addresses
- Static VXLAN used as access network for PBB-EVPN core: all-active multi-homing

Static VXLAN termination on system IP addresses

Figure 268: Example topology for static VXLAN termination on system IP addresses shows the example topology for static VXLAN termination on system IP addresses. The initial configuration of the PEs includes the cards, MDAs, ports, router interfaces, and IGP. BGP is not required on PE-1; on PE-2 and PE-4, BGP is configured for address family EVPN.

Figure 268: Example topology for static VXLAN termination on system IP addresses

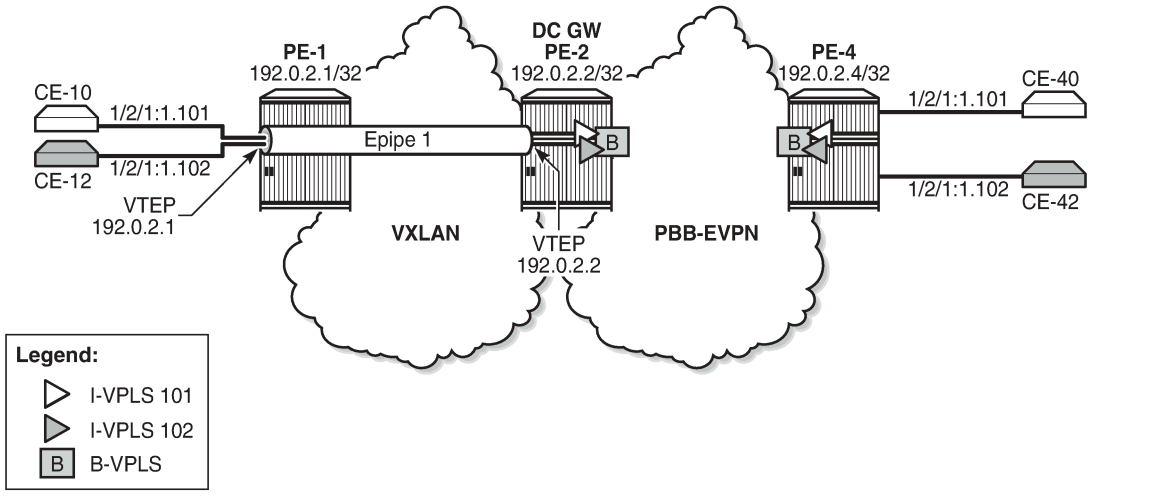

On PE-1, Epipe 1 is configured with egress VXLAN VNI 1, egress VTEP 192.0.2.2, oper-group op-grp-1, and a SAP toward CE-10, as follows. This configuration was explained in the text under Figure 267: Static VXLAN termination on system IP addresses.

```
# on PE-1:
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
configure {
    service {
        oper-group "op-grp-1" {
        }
        epipe "Epipe 1" {
            admin-state enable
            service-id 1
            customer "1"
            sap 1/2/1:1.* {
            }
            vxlan {
                instance 1 {
                    vni 1
                    egress-vtep {
                        ip-address 192.0.2.2
                        oper-group "op-grp-1"
                    }
                }
            }
        }
```

On PE-2, BGP is configured for address family EVPN, as follows:

```
# on PE-2:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            rapid-withdrawal true
            split-horizon true
            rapid-update {
                evpn true
            }
            group "internal" {
                peer-as 64500
                family {
                    evpn true
                }
            }
            neighbor "192.0.2.4" {
                group "internal"
            }
        }
```

There is a PXC configured on port 1/2/1 that will connect SAP pxc-21.a:1.* in Epipe 1, SAP pxc-21.b:1.101 in I-VPLS 101, and SAP pxc-21.b:1.102 in I-VPLS 102. The PXC is configured on PE-2 as follows. See the "Port Cross-Connect (PCX)" chapter in the Interface Configuration volume in the *7450 ESS, 7750 SR, and 7950 XRS MD-CLI Advanced Configuration Guide - Part I* for more information.

```
# on PE-2:
configure {
    port-xc {
        pxc 21 {
            admin-state enable
            port-id 1/2/1
        }
    }
    port pxc-21.a {
        admin-state enable
        ethernet {
            encap-type qinq
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    }
    port pxc-21.b {
        admin-state enable
        ethernet {
            encap-type qinq
        }
    }
    port 1/2/1 {
        admin-state enable
        ethernet {
            mode hybrid
            dot1x {
                tunneling true
            }
        }
    }
```

The service configuration on PE-2 includes Epipe 1, B-VPLS 100, and I-VPLSs 101-102, as follows:

```
# on PE-2:
configure {
    service {
        oper-group "op-grp-1" {
        }
        epipe "Epipe 1" {
            admin-state enable
            service-id 1
            customer "1"
            sap pxc-21.a:1.* {
            }
            vxlan {
                instance 1 {
                    vni 1
                    egress-vtep {
                        ip-address 192.0.2.1
                        oper-group "op-grp-1"
                    }
                }
            }
        }
        vpls "B-VPLS 100" {
            admin-state enable
            service-id 100
            customer "1"
            service-mtu 2000
            pbb-type b-vpls
            pbb {
                source-bmac {
                    address 00:00:00:00:00:02
                }
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 100
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
        vpls "I-VPLS 101" {
            admin-state enable
            service-id 101
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS 100" {
                    isid 101
                }
            }
            sap pxc-21.b:1.101 {
            }
        }
        vpls "I-VPLS 102" {
            admin-state enable
            service-id 102
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS 100" {
                    isid 102
                }
            }
            sap pxc-21.b:1.102 {
            }
        }
```

The service configuration on PE-4 is similar for the B-VPLS and the I-VPLSs, but Epipe 1 is not configured on PE-4.

The following command shows the VXLAN information for Epipe 1 on PE-1. By default, the source VTEP is the system IP address 192.0.2.1.

```
[/]
A:admin@PE-1# show service id 1 vxlan
===============================================================================
Vxlan Src Vtep IP: N/A
===============================================================================

===============================================================================
Vxlan Instance
===============================================================================
VXLAN Instance                VNI                     Oper-flags
-------------------------------------------------------------------------------
1                             1                       none
-------------------------------------------------------------------------------
Number of Entries : 1
-------------------------------------------------------------------------------
===============================================================================
```

```
[/]
A:admin@PE-1# show service id 1 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address                        Egress VNI          Oper      Vxlan
                                                        State     Type
-------------------------------------------------------------------------------
192.0.2.2                           1                   Up        static
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Number of Egress VTEP, VNI : 1
-------------------------------------------------------------------------------
===============================================================================
---snip---
```

The following command shows the oper-group information on PE-1 with the list of egress VTEP members.

```
[/]
A:admin@PE-1# show service oper-group "op-grp-1" detail

===============================================================================
Service Oper Group Information
===============================================================================
Oper Group       : op-grp-1
Creation Origin  : manual                          Oper Status: up
Hold DownTime    : 0 secs                          Hold UpTime: 4 secs
Members          : 1                               Monitoring : 0
===============================================================================

===============================================================================
Member Egr-Vtep for OperGroup: op-grp-1
===============================================================================
Svc Id              VNI                 VTEP Address
-------------------------------------------------------------------------------
1                   1                   192.0.2.2
-------------------------------------------------------------------------------
Egr-Vtep Entries found: 1
===============================================================================
```

The oper-group with member egress VTEP 192.0.2.2 cannot be monitored on a SAP in the same Epipe. The following error is raised when attempting to configure the same oper-group for the SAP in Epipe 1 on PE-1:

```
[ex:/configure service epipe "Epipe 1" sap 1/2/1:1.*]
A:admin@PE-1# oper-group "op-grp-1"

*[ex:/configure service epipe "Epipe 1" sap 1/2/1:1.*]
A:admin@PE-1# commit
MINOR: MGMT_CORE #5001: configure service epipe "Epipe 1" sap 1/2/1:1.* - Oper-group has an Egr
Vtep member, no other members allowed
```

The following ports on PE-2 are disabled to make the destination VTEP unreachable from PE-1:

```
# on PE-2:
configure {
    port 1/1/1 {
        admin-state disable
    }
    port 1/1/2 {
        admin-state disable
    }
```

When the destination VTEP disappears from the RTM, the oper-group op-grp-1 goes down and the VXLAN binding in Epipe 1 goes down, while the Epipe service remains up, as follows:

```
[/]
A:admin@PE-1# show service oper-group "op-grp-1"

===============================================================================
Service Oper Group Information
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Oper Group       : op-grp-1
Creation Origin  : manual                          Oper Status: down
Hold DownTime    : 0 secs                          Hold UpTime: 4 secs
Members          : 1                               Monitoring : 0
===============================================================================
```

```
[/]
A:admin@PE-1# show service id 1 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address                          Egress VNI          Oper      Vxlan
                                                          State     Type
-------------------------------------------------------------------------------
192.0.2.2                             1                   Down      static
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 1
-------------------------------------------------------------------------------
===============================================================================
---snip---
```

```
[/]
A:admin@PE-1# show service id 1 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 1                 Vpn Id            : 0
Service Type      : Epipe
---snip---

Admin State       : Up                Oper State        : Up
---snip---

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                            Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/2/1:1.*                         qinq         1578    1578    Up   Up
===============================================================================
```

The output is similar on PE-2. The ports are re-enabled on PE-2, which will cause the VXLAN binding and the oper-group to be operationally up again:

```
# on PE-2:
configure {
    port 1/1/1 {
        admin-state enable
    }
    port 1/1/2 {
        admin-state enable
    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The preceding example proved that the Epipe service remains up when the VXLAN binding goes down. The following example shows that the Epipe service goes down when the SAP goes down. On PE-1, port 1/2/1 is disabled, as follows:

```
# on PE-1:
configure {
    port 1/2/1
        admin-state disable
```

The following command shows that SAP 1/2/1:1.* and Epipe 1 are down on PE-1:

```
[/]
A:admin@PE-1# show service id 1 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 1                   Vpn Id            : 0
Service Type      : Epipe
---snip---

Admin State       : Up                  Oper State        : Down
---snip---

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                               Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/2/1:1.*                            qinq         1578    1578    Up   Down
===============================================================================
```

The port is re-enabled and SAP 1/2/1:1 and service Epipe 1 will be up again.

```
# on PE-1:
configure {
    port 1/2/1 {
        admin-state enable
```

When the service is disabled (**admin-state disable**), the SAP goes down, the VXLAN binding goes down, and the oper-group goes down, as follows:

```
# on PE-1:
configure {
    service {
        epipe "Epipe 1" {
            admin-state disable
```

```
[/]
A:admin@PE-1# show service id 1 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 1                   Vpn Id            : 0
Service Type      : Epipe
---snip---

Admin State       : Down                Oper State        : Down
---snip---
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                              Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/2/1:1.*                           qinq         1578    1578    Up   Down
===============================================================================
```

```
[/]
A:admin@PE-1# show service id 1 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address                            Egress VNI          Oper      Vxlan
                                                            State     Type
-------------------------------------------------------------------------------
192.0.2.2                               1                   Down      static
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 1
-------------------------------------------------------------------------------
===============================================================================
```

```
[/]
A:admin@PE-1# show service oper-group

===============================================================================
Service Oper Group Information
===============================================================================
Name                            Oper   Creation Hold   Hold   Members Monitor
                                Status Origin   UpTime DnTime
                                                (secs) (secs)
-------------------------------------------------------------------------------
op-grp-1                        down   manual   4      0      1       0
-------------------------------------------------------------------------------
Entries found: 1
===============================================================================
```

Static VXLAN termination on non-system IPv4 addresses

Non-system IP VXLAN termination is provisioned as follows:

1. Create FPE
2. Associate FPE with VXLAN termination
3. Configure router loopback interface
4. Configure non-system VXLAN termination VTEP addresses
5. Add the service configuration

Figure 269: Example topology for static VXLAN termination on non-system IPv4 addresses shows the example topology with PE-1 and PE-2 in a VXLAN network. The non-system loopback address on PE-2 will be used for VXLAN termination, whereas the system IP address will be used on PE-1.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 269: Example topology for static VXLAN termination on non-system IPv4 addresses

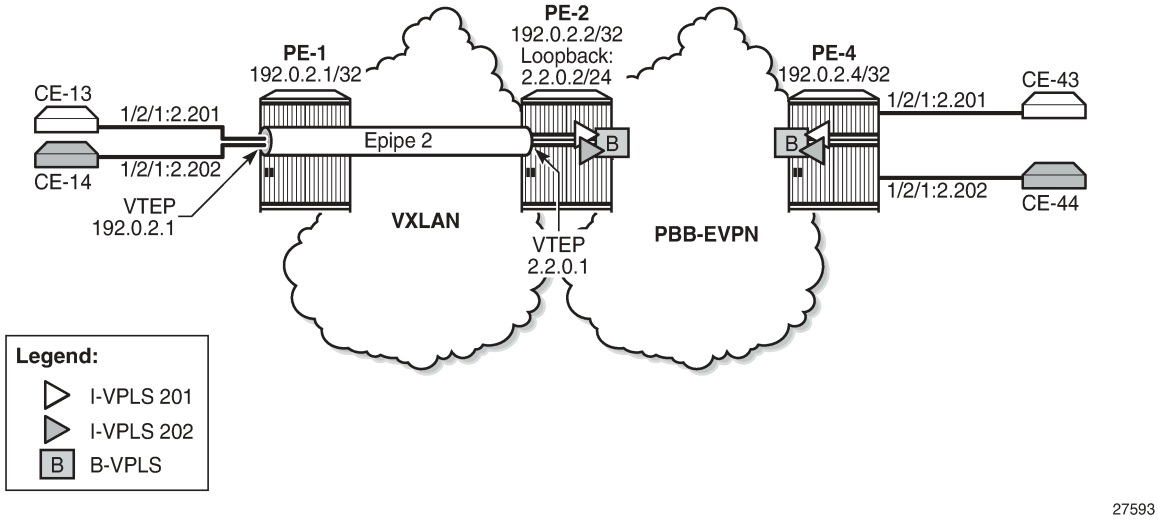

Create FPE

FPE uses the back-to-back PXC, either a PXC port or a LAG-based PXC. PXC 1 is created on PE-2:

```
# on PE-2:
configure {
    port-xc {
        pxc 1 {
            admin-state enable
            port-id 1/2/5
        }
    }
    port pxc-1.a {
        admin-state enable
    }
    port pxc-1.b {
        admin-state enable
    }
    port 1/2/5 {
        admin-state enable
        ethernet {
            mode hybrid
            dot1x {
                tunneling true
            }
        }
    }
```

```
[/]
A:admin@PE-2# show port pxc 1

===============================================================================
Ports on Port Cross Connect 1
===============================================================================
Port          Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id            State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
pxc-1.a       Up    Yes  Up      1574 1574    - hybr dotq xgige
pxc-1.b       Up    Yes  Up      1574 1574    - hybr dotq xgige
===============================================================================
```

The following FPE uses the PXC:

```
# on PE-2:
configure {
    fwd-path-ext {
        fpe 1 {
            path {
                pxc 1
            }
```

The following shows that FPE 1 uses PXC 1 and has no VXLAN termination associated:

```
[/]
A:admin@PE-2# show fwd-path-ext fpe 1

===============================================================================
FPE Id: 1
===============================================================================
Description        : (Not Specified)
Path               : pxc 1
Pw Port            : Disabled                          Oper     : down
Sub Mgmt Extension : Disabled                          Oper     : N/A
Vxlan Termination  : Disabled                          Oper     : down
Segment-Routing V6 : Disabled
===============================================================================
```

Associate FPE with VXLAN termination

The following command associates FPE 1 with VXLAN termination:

```
# on PE-2:
configure {
    fwd-path-ext {
        sdp-id-range {
            start 10000
            end 10127
        }
        fpe 1 {
            path {
                pxc 1
            }
            application {
                vxlan-termination {
                }
            }
        }
```

When attempting to associate the FPE with VXLAN termination without configuring a range of SDP IDs for FPE, the following error is raised:

```
*[ex:/configure fwd-path-ext fpe 1 application vxlan-termination]
A:admin@PE-2# commit
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
MINOR: FPE #1021: configure fwd-path-ext fpe 1 - sdp-id-range is not configured - configure
 fwd-path-ext sdp-id-range
```

The following shows the range of SDP IDs for FPE and the list of configured FPEs; see the VXLAN Forwarding Path Extension chapter for more information about the use of SDP IDs. The application for FPE 1 is VXLAN termination.

```
[/]
A:admin@PE-2# show fwd-path-ext

===============================================================================
FPE Info
===============================================================================
FPE Id                  Path                          Application
                        pxc/xc-a, xc-b
-------------------------------------------------------------------------------
1                       pxc 1                         vxlan-term
-------------------------------------------------------------------------------
Number of entries : 1
-------------------------------------------------------------------------------
SDP-Id Range: 10000 - 10127
===============================================================================
```

After the FPEs are associated with VXLAN termination, the system creates two internal router interfaces per FPE, one per PXC sub-port, as follows:

```
[/]
A:admin@PE-2# show router interface

===============================================================================
Interface Table (Router: Base)
===============================================================================
Interface-Name                   Adm       Opr(v4/v6)  Mode    Port/SapId
   IP-Address                                                   PfxState
-------------------------------------------------------------------------------
_tmnx_fpe_1.a                    Up        Up/Up       Network pxc-1.a:1
   fe80::100/64                                                 PREFERRED
_tmnx_fpe_1.b                    Up        Up/Up       Network pxc-1.b:1
   fe80::101/64                                                 PREFERRED
---snip---
```

Configure router loopback interface

The following loopback interface is configured in PE-2 and added to the IS-IS context. The IPv6 address is not required yet.

```
# on PE-2:
configure {
    router "Base" {
        interface "loopback1" {
            loopback
            ipv4 {
                primary {
                    address 2.2.0.2
                    prefix-length 24
                }
            }
            ipv6 {
                address 220::2 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    prefix-length 120
                }
            }
        }
        isis 0 {
            interface "loopback1" {
            }
        }
```

A subnet must be assigned to the loopback interface, but not a /32 or /128 subnet mask, because the system cannot terminate VXLAN on a local interface address. In the preceding example, all addresses in the subnet 2.2.0.0/24 can be used for VXLAN tunnel termination, except for 2.2.0.2. The subnet will be advertised by the IGP. The subnet can be as small as /31 or /127.

Configure non-system VTEP addresses

On PE-2, non-system IP address 2.2.0.1 in the subnet of the loopback address 2.2.0.2/24 is configured as VTEP, as follows. Up to three non-system VTEP addresses can be configured to terminate VXLAN tunnels and their corresponding FPEs.

```
# on PE-2:
configure {
    service {
        system {
            vxlan {
                tunnel-termination 2.2.0.1 {
                    fpe-id 1
                }
            }
```

No non-system VTEP addresses need to be configured on PE-1.

When the non-system VTEP address is configured, an internal loopback interface _tmnx_vli_vxlan_1_131075 with VTEP address 2.2.0.1/32 is auto-created that can respond to ICMP requests.

```
[/]
A:admin@PE-2# show router interface

===============================================================================
Interface Table (Router: Base)
===============================================================================
Interface-Name                   Adm       Opr(v4/v6)  Mode    Port/SapId
   IP-Address                                                  PfxState
-------------------------------------------------------------------------------
_tmnx_fpe_1.a                    Up        Up/Up       Network pxc-1.a:1
   fe80::100/64                                                PREFERRED
_tmnx_fpe_1.b                    Up        Up/Up       Network pxc-1.b:1
   fe80::101/64                                                PREFERRED
_tmnx_vli_vxlan_1_131075         Up        Up/Up       Network loopback
   2.2.0.1/32                                                  n/a
   fe80::13:ffff:fe00:0/64                                     PREFERRED
---snip---
```

The system does not verify if there is a local base router loopback interface with a subnet corresponding to the VTEP address. If a tunnel termination address is configured and the FPE is up, the system will start terminating VXLAN traffic and responding ICMP for that address, regardless of the presence of a loopback

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

in the base router. It is also possible that a non-loopback interface has an IP address in the configured subnet.

Configure the services

Epipe 2 is configured on PE-1 as follows. By default, the system IP address will be used as source VTEP of the VXLAN-encapsulated frames. The non-system IP address 2.2.0.1 is used as egress VTEP.

```
# on PE-1:
configure {
    service {
        epipe "Epipe 2" {
            admin-state enable
            service-id 2
            customer "1"
            sap 1/2/1:2.* {
            }
            vxlan {
                instance 1 {
                    vni 2
                    egress-vtep {
                        ip-address 2.2.0.1
                    }
                }
            }
        }
```

The configuration of Epipe 2 on PE-2 defines the non-system IP address 2.2.0.1 as source VTEP, as follows. The egress VTEP is 192.0.2.1, the system IP address of PE-1. The configuration of the B-VPLS is the same as in the preceding example; the configuration of the I-VPLSs 201 and 202 is similar to the configuration of I-VPLS 101 in the preceding example.

```
# on PE-2:
configure {
    service {
        epipe "Epipe 2" {
            admin-state enable
            service-id 2
            customer "1"
            sap pxc-21.a:2.* {
            }
            vxlan {
                source-vtep 2.2.0.1
                instance 1 {
                    vni 2
                    egress-vtep {
                        ip-address 192.0.2.1
                    }
                }
            }
        }
```

The following **show** command on PE-1 shows that no VXLAN source VTEP IP address is configured:

```
[/]
A:admin@PE-1# show service id 2 vxlan
===============================================================================
Vxlan Src Vtep IP: N/A
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Vxlan Instance
===============================================================================
VXLAN Instance               VNI                     Oper-flags
-------------------------------------------------------------------------------
1                            2                       none
-------------------------------------------------------------------------------
Number of Entries : 1
-------------------------------------------------------------------------------
===============================================================================
```

The following shows that the egress VTEP is 2.2.0.1, which is a non-system VTEP on PE-2. The VXLAN tunnel is operationally up.

```
[/]
A:admin@PE-1# show service id 2 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address                          Egress VNI          Oper      Vxlan
                                                          State     Type
-------------------------------------------------------------------------------
2.2.0.1                               2                   Up        static
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 1
-------------------------------------------------------------------------------
===============================================================================
---snip---
```

The same commands on PE-2 show that source VTEP IP address 2.2.0.1 is configured and the egress VTEP is 192.0.2.1, which is the system IP address of PE-1, as follows:

```
[/]
A:admin@PE-2# show service id 2 vxlan
===============================================================================
Vxlan Src Vtep IP: 2.2.0.1
===============================================================================

===============================================================================
Vxlan Instance
===============================================================================
VXLAN Instance               VNI                     Oper-flags
-------------------------------------------------------------------------------
1                            2                       none
-------------------------------------------------------------------------------
Number of Entries : 1
-------------------------------------------------------------------------------
===============================================================================
```

```
[/]
A:admin@PE-2# show service id 2 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address                          Egress VNI          Oper      Vxlan
                                                          State     Type
-------------------------------------------------------------------------------
192.0.2.1                             2                   Up        static
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 1
-------------------------------------------------------------------------------
===============================================================================
---snip---
```

Static VXLAN termination on IPv6 addresses

IPv6 VXLAN termination is provisioned as follows:

1. Create FPE
2. Associate FPE with VXLAN termination
3. Configure router loopback interface
4. Configure non-system VXLAN termination VTEP addresses
5. Add the service configuration

Figure 270: Example topology for static VXLAN termination on IPv6 addresses shows the example topology with PE-1 and PE-2 in a VXLAN network. The loopback addresses on PE-1 and PE-2 will be used for IPv6 VXLAN termination. The existing PXC 1 on PE-2 is reused for FPE; only an IPv6 VTEP address needs to be added.

Figure 270: Example topology for static VXLAN termination on IPv6 addresses

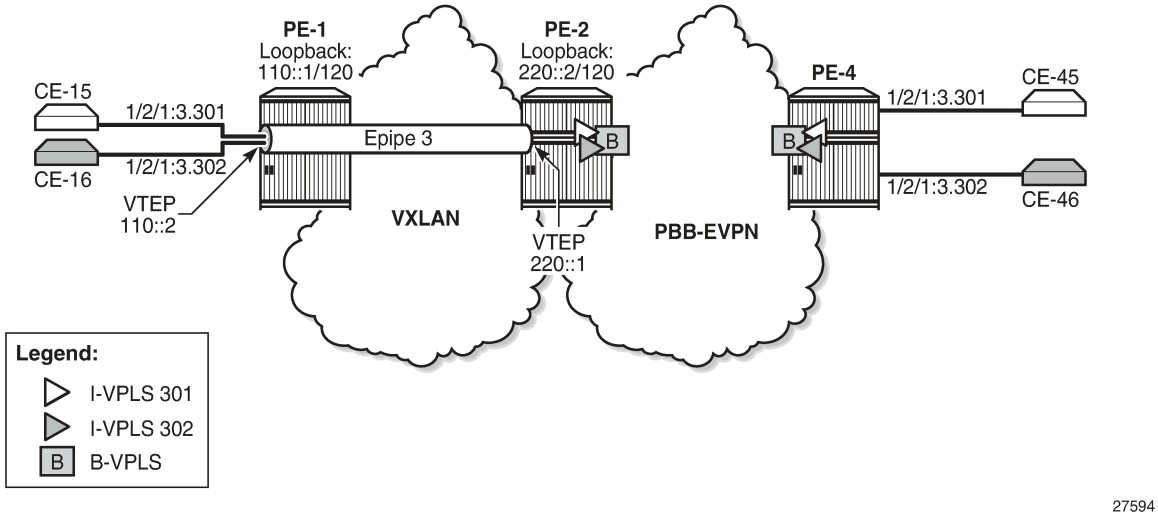

For IPv6 routing, the following option is configured for IS-IS on all nodes:

```
# on all PEsL
configure {
    router "Base" {
        isis 0 {
            ipv6-routing native
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Create FPE

The following PXC is created on PE-1; PXC 1 will be used for FPE:

```
# on PE-1:
configure {
    port-xc {
        pxc 1 {
            admin-state enable
            port-id 1/2/5
        }
    }
    port pxc-1.a {
        admin-state enable
    }
    port pxc-1.b {
        admin-state enable
    }
    port 1/2/5 {
        admin-state enable
        ethernet {
            mode hybrid
            dot1x {
                tunneling true
            }
        }
    }
```

```
[/]
A:admin@PE-1# show port pxc 1

===============================================================================
Ports on Port Cross Connect 1
===============================================================================
Port          Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id            State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
pxc-1.a       Up    Yes  Up      1574 1574    - hybr dotq xgige
pxc-1.b       Up    Yes  Up      1574 1574    - hybr dotq xgige
===============================================================================
```

Configure FPE with VXLAN termination

The following command associates FPE 1 with VXLAN termination:

```
# on PE-1:
configure {
    fwd-path-ext {
        sdp-id-range {
            start 10000
            end 10127
        }
        fpe 1 {
            path {
                pxc 1
            }
            application {
                vxlan-termination {
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
```

The following shows the range of SDP IDs for FPE and the list of configured FPEs. The application for FPE 1 is VXLAN termination.

```
[/]
A:admin@PE-1# show fwd-path-ext

===============================================================================
FPE Info
===============================================================================
FPE Id                  Path                          Application
                        pxc/xc-a, xc-b
-------------------------------------------------------------------------------
1                       pxc 1                         vxlan-term
-------------------------------------------------------------------------------
Number of entries : 1
-------------------------------------------------------------------------------
SDP-Id Range: 10000 - 10127
===============================================================================
```

The following shows that FPE 1 has a VXLAN termination that is oper up:

```
[/]
A:admin@PE-1# show fwd-path-ext fpe 1

===============================================================================
FPE Id: 1
===============================================================================
Description        : (Not Specified)
Path               : pxc 1
Pw Port            : Disabled                        Oper     : down
Sub Mgmt Extension : Disabled                        Oper     : N/A
Vxlan Termination  : Router: Base                    Oper     : up
Segment-Routing V6 : Disabled
===============================================================================
```

After the FPEs are associated with VXLAN termination, the system creates two internal router interfaces per FPE, one per PXC sub-port, as follows:

```
[/]
A:admin@PE-1# show router interface

===============================================================================
Interface Table (Router: Base)
===============================================================================
Interface-Name                   Adm       Opr(v4/v6)  Mode    Port/SapId
   IP-Address                                                   PfxState
-------------------------------------------------------------------------------
_tmnx_fpe_1.a                    Up        Up/Up       Network pxc-1.a:1
   fe80::100/64                                                 PREFERRED
_tmnx_fpe_1.b                    Up        Up/Up       Network pxc-1.b:1
   fe80::101/64                                                 PREFERRED
---snip---
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Configure router loopback interface

The following loopback interface is configured in PE-1 and added to the IS-IS context:

```
# on PE-1:
configure {
    router "Base" {
        interface "loopback1" {
            loopback
            ipv4 {
                primary {
                    address 1.1.0.1
                    prefix-length 24
                }
            }
            ipv6 {
                address 110::1 {
                    prefix-length 120
                }
            }
        }
        isis 0 {
            interface "loopback1" {
            }
        }
```

All IPv6 addresses in the 110::/120 subnet can be used for VXLAN tunnel termination, except for 110::1.

Configure non-system VTEP addresses

On PE-1, IPv6 address 110::2 in the subnet of the loopback address 110::1/120 is configured as VTEP, as follows:

```
# on PE-1:
configure {
    service {
        system {
            vxlan {
                tunnel-termination 110::2 {
                    fpe-id 1
                }
            }
```

On PE-2, IPv6 address 220::1 in the subnet of the loopback address 220::2/120 is configured as VTEP, as follows:

```
# on PE-2:
configure {
    service {
        system {
            vxlan {
                tunnel-termination 220::1 {
                    fpe-id 1
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

When the IPv6 VTEP address is configured on PE-1, an internal loopback interface _tmnx_vli_vxlan_1_131075 is created, as follows.

```
[/]
A:admin@PE-1# show router interface

===============================================================================
Interface Table (Router: Base)
===============================================================================
Interface-Name                   Adm       Opr(v4/v6)  Mode    Port/SapId
   IP-Address                                                   PfxState
-------------------------------------------------------------------------------
_tmnx_fpe_1.a                    Up        Up/Up       Network pxc-1.a:1
   fe80::100/64                                                 PREFERRED
_tmnx_fpe_1.b                    Up        Up/Up       Network pxc-1.b:1
   fe80::101/64                                                 PREFERRED
_tmnx_vli_vxlan_1_131075         Up        Down/Up     Network loopback
   110::2/128                                                   PREFERRED
   fe80::f:ffff:fe00:0/64                                       PREFERRED
---snip---
```

The following IPv6 route table on PE-1 contains an internal static route for source VTEP 110::2/128 using the FPE internal interface _tmnx_fpe_1.a:

```
[/]
A:admin@PE-1# show router route-table ipv6

===============================================================================
IPv6 Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
110::/120                                     Local   Local     01h34m32s  0
       loopback1                                                    0
110::2/128                                    Remote  Static    00h33m20s  5
       fe80::101-"_tmnx_fpe_1.a"                                    1
220::/120                                     Remote  ISIS      00h15m03s  15
       fe80::616:1ff:fe01:2-"int-PE-1-PE-2"                         10
---snip---
```

The following IPv6 route table on PE-2 shows that an internal static route is configured for the source VTEP 220::1/128 using the FPE internal interface _tmnx_fpe_1.a:

```
[/]
A:admin@PE-2# show router route-table ipv6

===============================================================================
IPv6 Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
110::/120                                     Remote  ISIS      00h00m46s  15
       fe80::10:1ff:fe01:1-"int-PE-2-PE-1"                          10
220::/120                                     Local   Local     00h05m08s  0
       loopback1                                                    0
220::1/128                                    Remote  Static    00h00m24s  5
       fe80::101-"_tmnx_fpe_1.a"                                    1
---snip---
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Configure the services

Epipe 3 is configured on PE-1 with **source-vtep** 110::2, which is the VTEP address configured in the preceding step (VXLAN tunnel termination). The egress VTEP is 220::1, which is the VXLAN termination configured on PE-2.

```
# on PE-1:
configure {
    service {
        epipe "Epipe 3" {
            admin-state enable
            service-id 3
            customer "1"
            sap 1/2/1:3.* {
            }
            vxlan {
                source-vtep 110::2
                instance 1 {
                    vni 3
                    egress-vtep {
                        ip-address 220::1
                    }
                }
            }
        }
```

Epipe 3 on PE-2 has VXLAN source VTEP 220::1 and egress VTEP 110::2.

```
# on PE-2:
configure {
    service {
        epipe "Epipe 3" {
            admin-state enable
            service-id 3
            customer "1"
            sap pxc-21.a:3.* {
            }
            vxlan {
                source-vtep 220::1
                instance 1 {
                    vni 3
                    egress-vtep {
                        ip-address 110::2
                    }
                }
            }
        }
```

The configuration of the B-VPLS is the same as in the preceding example. The configuration of I-VPLS 302 is similar.

```
# on PE-2:
configure {
    service {
        vpls "I-VPLS 301" {
            admin-state enable
            service-id 301
            customer "1"
            pbb-type i-vpls
            pbb {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                backbone-vpls "B-VPLS 100" {
                    isid 301
                }
            }
            sap pxc-21.b:3.301 {
            }
        }
```

The following **show** commands on PE-1 show that the VXLAN source VTEP IP address is 110::2 and the egress VTEP is 220::1. The VXLAN tunnel is operationally up.

```
[/]
A:admin@PE-1# show service id 3 vxlan
===============================================================================
Vxlan Src Vtep IP: 110::2
===============================================================================

===============================================================================
Vxlan Instance
===============================================================================
VXLAN Instance                VNI                     Oper-flags
-------------------------------------------------------------------------------
1                             3                       none
-------------------------------------------------------------------------------
Number of Entries : 1
-------------------------------------------------------------------------------
===============================================================================
```

```
[/]
A:admin@PE-1# show service id 3 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address                          Egress VNI          Oper      Vxlan
                                                          State     Type
-------------------------------------------------------------------------------
220::1                                3                   Up        static
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 1
-------------------------------------------------------------------------------
===============================================================================
---snip---
```

The same commands on PE-2 show VXLAN source VTEP 220::1 and egress VTEP 110::2, as follows:

```
[/]
A:admin@PE-2# show service id 3 vxlan
===============================================================================
Vxlan Src Vtep IP: 220::1
===============================================================================

===============================================================================
Vxlan Instance
===============================================================================
VXLAN Instance                VNI                     Oper-flags
-------------------------------------------------------------------------------
1                             3                       none
-------------------------------------------------------------------------------
Number of Entries : 1
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
```

```
[/]
A:admin@PE-2# show service id 3 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address                            Egress VNI          Oper      Vxlan
                                                            State     Type
-------------------------------------------------------------------------------
110::2                                  3                   Up        static
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 1
-------------------------------------------------------------------------------
===============================================================================
---snip---
```

Static VXLAN used as access network for PBB-EVPN core: all-active multi-homing and anycast VTEPs

Figure 271: Example topology for static VXLAN termination using anycast shows the example topology with PE-1, PE-2, and PE-3 in the VXLAN access network. Epipe 4 is configured on PE-1, PE-2, and PE-3. On PE-1, the system IP address 192.0.2.1 is used as source VTEP, while (anycast) IP address 23.23.23.4 is used as source VTEP on PE-2 and PE-3.

In the PBB-EVPN core network, all-active multi-homing virtual Ethernet segments vES-23_401 and vES-23_402 are configured on PE-2 and PE-3.

Figure 271: Example topology for static VXLAN termination using anycast

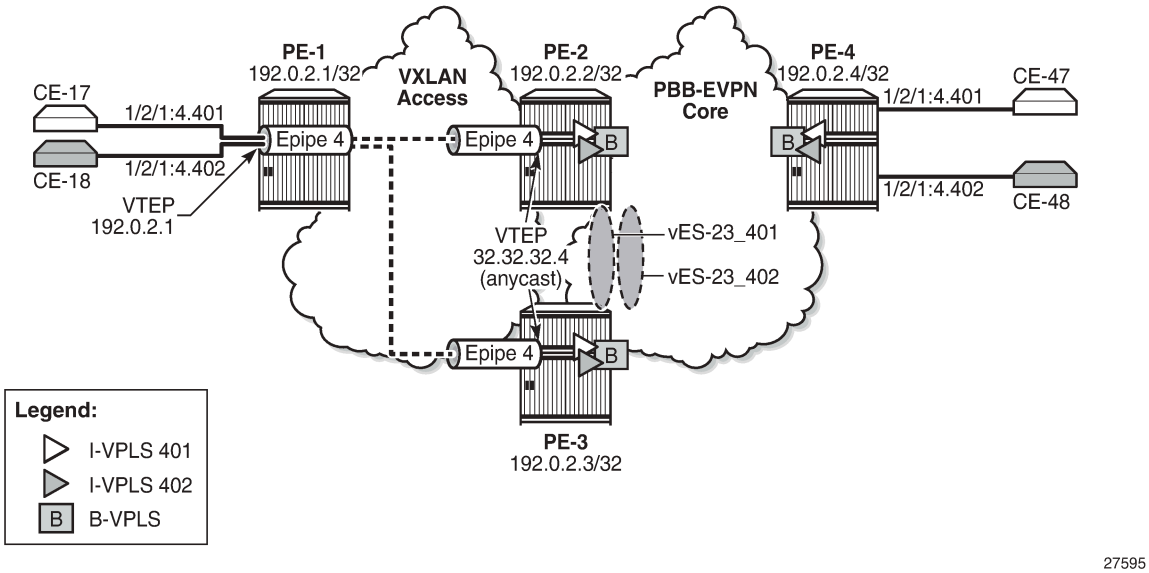

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

VXLAN access network

On PE-2 and PE-3, PXC ports are configured: PXC 2 will be used as FPE, whereas PXC-3 and PXC-4 will be used to make a LAG for the PXC between Epipe and I-VPLS services. The PXC sub-ports for FPE have dot1q encapsulation whereas the PXC sub-ports for port cross-connect have qinq encapsulation. The configuration of the PXC ports and sub-ports is as follows:

```
# on PE-2, PE-3:
configure {
    port 1/2/6 {
        admin-state enable
        ethernet {
            mode hybrid
            dot1x {
                tunneling true
            }
        }
    }
    port 1/2/7 {
        admin-state enable
        ethernet {
            mode hybrid
            dot1x {
                tunneling true
            }
        }
    }
    port 1/2/8 {
        admin-state enable
        ethernet {
            mode hybrid
            dot1x {
                tunneling true
            }
        }
    }
    port pxc-2.a {
        admin-state enable
    }
    port pxc-2.b {
        admin-state enable
    }
    port pxc-3.a {
        admin-state enable
        ethernet {
            encap-type qinq
        }
    }
    port pxc-3.b {
        admin-state enable
        ethernet {
            encap-type qinq
        }
    }
    port pxc-4.a {
        admin-state enable
        ethernet {
            encap-type qinq
        }
    }
    port pxc-4.b {
        admin-state enable
        ethernet {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            encap-type qinq
        }
    }
    port-xc {
        pxc 2 {
            admin-state enable
            port-id 1/2/6
        }
        pxc 3 {
            admin-state enable
            port-id 1/2/7
        }
        pxc 4 {
            admin-state enable
            port-id 1/2/8
        }
    }
```

On PE-2 and PE-3, FPE 2 is configured. FPE 2 is associated with VXLAN termination and two internal interfaces will be auto-created: _tmnx_fpe_2.a and _tmnx_fpe_2.b.

```
# on PE-2, PE-3:
configure {
    fwd-path-ext {
        sdp-id-range {
            start 10000
            end 10127
        }
        fpe 2 {
            path {
                pxc 2
            }
            application {
                vxlan-termination {
                }
            }
        }
```

```
[/]
A:admin@PE-2# show router interface

===============================================================================
Interface Table (Router: Base)
===============================================================================
Interface-Name                   Adm       Opr(v4/v6)  Mode    Port/SapId
   IP-Address                                                  PfxState
-------------------------------------------------------------------------------
_tmnx_fpe_1.a                    Up        Up/Up       Network pxc-1.a:1
   fe80::100/64                                                PREFERRED
_tmnx_fpe_1.b                    Up        Up/Up       Network pxc-1.b:1
   fe80::101/64                                                PREFERRED
_tmnx_fpe_2.a                    Up        Up/Up       Network pxc-2.a:1
   fe80::200/64                                                PREFERRED
_tmnx_fpe_2.b                    Up        Up/Up       Network pxc-2.b:1
   fe80::201/64                                                PREFERRED
---snip---
```

A router loopback interface with IP address 23.23.23.2/24 is created on PE-2, and on PE-3 with IP address 23.23.23.3/24:

```
# on PE-2:
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
configure {
    router "Base" {
        interface "loopback2" {
            loopback
            ipv4 {
                primary {
                    address 23.23.23.2
                    prefix-length 24
                }
            }
        }
        isis 0 {
            interface "loopback2" {
            }
        }
```

On PE-2 and PE-3, the VTEP 23.23.23.4 is configured for FPE 2, as follows:

```
# on PE-2, PE-3:
configure {
    service {
        system {
            vxlan {
                tunnel-termination 23.23.23.4 {
                    fpe-id 2
                }
            }
        }
```

The following command shows an additional VTEP 23.23.23.4 to the existing router interface _tmnx_vli_vxlan_1_131075 on PE-2:

```
[/]
A:admin@PE-2# show router interface "_tmnx_vli_vxlan_1_131075"

===============================================================================
Interface Table (Router: Base)
===============================================================================
Interface-Name                   Adm       Opr(v4/v6)  Mode    Port/SapId
   IP-Address                                                  PfxState
-------------------------------------------------------------------------------
_tmnx_vli_vxlan_1_131075         Up        Up/Up       Network loopback
   2.2.0.1/32                                                  n/a
   220::1/128                                                  PREFERRED
   23.23.23.4/32                                               n/a
   fe80::13:ffff:fe00:0/64                                     PREFERRED
-------------------------------------------------------------------------------
Interfaces : 1
===============================================================================
```

On PE-2 and PE-3, the VXLAN Epipe 4 uses LAG 4 (composed of pxc-3.b and pxc-4.b) to extend the VXLAN toward the I-VPLSs 401 and 402. The I-VPLS SAPs use LAG 3 (composed of pxc-3.a and pxc-4.a). The PXC LAGs provide higher bandwidth and better resiliency. The LAGs are configured as follows on both PE-2 and PE-3:

```
# on PE-2, PE-3:
configure {
    lag "lag-3" {
        admin-state enable
        encap-type qinq
        mode hybrid
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        max-ports 64
        port pxc-3.a {
        }
        port pxc-4.a {
        }
    }
    lag "lag-4" {
        admin-state enable
        encap-type qinq
        mode hybrid
        max-ports 64
        port pxc-3.b {
        }
        port pxc-4.b {
        }
    }
```

Epipe 4 is configured on PE-1, PE-2, and PE-3. On PE-1, no FPE is required because the system IP address is used as VTEP. Epipe 4 is configured on PE-1 with egress VTEP 23.23.23.4, as follows:

```
# on PE-1:
configure {
    service {
        epipe "Epipe 4" {
            admin-state enable
            service-id 4
            customer "1"
            sap 1/2/1:4.* {
            }
            vxlan {
                instance 1 {
                    vni 4
                    egress-vtep {
                        ip-address 23.23.23.4
                    }
                }
            }
        }
```

Epipe 4 is configured on PE-2 and PE-3 with source VTEP 23.23.23.4 and egress VTEP 192.0.2.1, as follows. The SAP uses LAG 4, which is composed of PXC sub-ports pxc-3.b and pxc-4.b.

```
# on PE-2, PE-3:
configure {
    service {
        epipe "Epipe 4" {
            admin-state enable
            service-id 4
            customer "1"
            sap lag-4:4.* {
            }
            vxlan {
                source-vtep 23.23.23.4
                instance 1 {
                    vni 4
                    egress-vtep {
                        ip-address 192.0.2.1
                    }
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The following command on PE-1 shows that the egress VTEP in Epipe 4 equals 23.23.23.4.

```
[/]
A:admin@PE-1# show service id 4 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address                          Egress VNI          Oper      Vxlan
                                                          State     Type
-------------------------------------------------------------------------------
23.23.23.4                            4                   Up        static
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 1
-------------------------------------------------------------------------------
===============================================================================
---snip---
```

The following commands for Epipe 4 on PE-2 show a source VTEP equal to 23.23.23.4 and an egress VTEP equal to the system address of PE-1 (192.0.2.1), as follows:

```
[/]
A:admin@PE-2# show service id 4 vxlan
===============================================================================
Vxlan Src Vtep IP: 23.23.23.4
===============================================================================

===============================================================================
Vxlan Instance
===============================================================================
VXLAN Instance               VNI                     Oper-flags
-------------------------------------------------------------------------------
1                            4                       none
-------------------------------------------------------------------------------
Number of Entries : 1
-------------------------------------------------------------------------------
===============================================================================
```

```
[/]
A:admin@PE-2# show service id 4 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
VTEP Address                          Egress VNI          Oper      Vxlan
                                                          State     Type
-------------------------------------------------------------------------------
192.0.2.1                             4                   Up        static
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 1
-------------------------------------------------------------------------------
===============================================================================
---snip---
```

The output on PE-3 is identical: source VTEP 23.23.23.4 and egress VTEP 192.0.2.1.

The following route table on PE-1 shows that the best route toward 23.23.23.4 is via PE-2:

```
[/]
A:admin@PE-1# show router route-table 23.23.23.4
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
23.23.23.0/24                                 Remote  ISIS      00h04m13s  15
       192.168.12.2                                                 10
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

PBB-EVPN core network

Two all-active multi-homing virtual ESs are configured on PE-2 and PE-3. The preference for the DF election is configured manually, with opposite preference values for the vESs so that DF load balancing is achieved. While vES-23_401 has preference 5000 on PE-2 and preference 10000 on PE-3, vES-23_402 has preference 10000 on PE-2 and preference 5000 on PE-3. When no event has occurred that caused a DF switchover, PE-2 is DF for vES-23_402 and PE-3 is DF for vES-23_401. Both vESs use LAG 3, which is composed of pxc-3.a and pxc-4.a. For vES-23_401, the qinq encapsulation must match S-tag 4 and C-tag 401; for vES-23_402, the S-tag must be 4 and the C-tag 402. On PE-2, the vESs are configured as follows.

```
# on PE-2:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "vES-23_401" {
                        admin-state enable
                        type virtual
                        esi 0x01000000230401000001
                        multi-homing-mode all-active
                        df-election {
                            service-carving-mode manual
                            manual {
                                preference {
                                    mode non-revertive
                                    value 5000
                                }
                            }
                        }
                        association {
                            lag "lag-3" {
                                virtual-ranges {
                                    qinq {
                                        s-tag-c-tag 4 c-tag-start 401 {
                                            c-tag-end 401
                                        }
                                    }
                                }
                            }
                        }
                        pbb {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        source-bmac-lsb 0x2341
                    }
                }
                ethernet-segment "vES-23_402" {
                    admin-state enable
                    type virtual
                    esi 0x01000000230402000001
                    multi-homing-mode all-active
                    df-election {
                        service-carving-mode manual
                        manual {
                            preference {
                                mode non-revertive
                                value 10000
                            }
                        }
                    }
                    association {
                        lag "lag-3" {
                            virtual-ranges {
                                qinq {
                                    s-tag-c-tag 4 c-tag-start 402 {
                                        c-tag-end 402
                                    }
                                }
                            }
                        }
                    }
                    pbb {
                        source-bmac-lsb 0x2342
                    }
                }
            }
        }
    }
```

The B-VPLS 100 is configured to use the ES-BMAC. On PE-2, the B-VPLS is configured as follows.

```
# on PE-2:
configure {
    service {
        vpls "B-VPLS 100" {
            admin-state enable
            service-id 100
            customer "1"
            service-mtu 2000
            pbb-type b-vpls
            pbb {
                source-bmac {
                    address 00:00:00:00:00:02
                    use-es-bmac-lsb true
                }
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 100
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    auto-bind-tunnel {
                        resolution any
                    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
            }
        }
```

On PE-4, the following configuration sets ECMP to a value of 2 in the **bgp-evpn mpls** context of the B-VPLS, so that aliasing is possible.

```
# on PE-4:
configure {
    service {
        vpls "B-VPLS 100" {
            bgp-evpn {
                mpls 1 {
                    ecmp 2
                }
            }
```

On PE-2 and PE-3, the I-VPLSs are configured with SAP LAG 3, which is composed of pxc-3.a and pxc-4.a, as follows. The qinq encapsulation 4.401 in I-VPLS 401 matches the condition in vES-23_401, whereas qinq 4.402 in I-VPLS 402 matches vES-23_402.

```
# on PE-2, PE-3:
configure {
    service {
        vpls "I-VPLS 401" {
            admin-state enable
            service-id 401
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS 100" {
                    isid 401
                }
            }
            sap lag-3:4.401 {
            }
        }
        vpls "I-VPLS 402" {
            admin-state enable
            service-id 402
            customer "1"
            pbb-type i-vpls
            pbb {
                backbone-vpls "B-VPLS 100" {
                    isid 402
                }
            }
            sap lag-3:4.402 {
            }
        }
```

With the preceding configuration, PBB-EVPN all-active multi-homing and the anycast VTEP at the access VXLAN network can be combined for an efficient and fully redundant network. PE-4 can alias the known unicast traffic to PE-2 and PE-3 on a per-flow basis, whereas if ECMP (and shared queuing) is enabled on PE-1, traffic can also be load-balanced to PE-2 and PE-3. BUM traffic sent from PE-4 will be forwarded by the corresponding DF for the ES.

See chapter EVPN for PBB over MPLS (PBB-EVPN) for more information about PBB-EVPN and all-active multi-homing.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Verification

The following command shows that PE-2 is NDF in vES-23_401 in I-VPLS 401:

```
[/]
A:admin@PE-2# show service id 401 ethernet-segment

===============================================================================
SAP Ethernet-Segment Information
===============================================================================
SAP                   Eth-Seg                         Status
-------------------------------------------------------------------------------
lag-3:4.401           vES-23_401                      NDF
===============================================================================
No sdp entries
No vxlan instance entries
```

For I-VPLS 402, PE-2 is DF, as follows:

```
[/]
A:admin@PE-2# show service id 402 ethernet-segment

===============================================================================
SAP Ethernet-Segment Information
===============================================================================
SAP                   Eth-Seg                         Status
-------------------------------------------------------------------------------
lag-3:4.402           vES-23_402                      DF
===============================================================================
No sdp entries
No vxlan instance entries
```

For PE-3, the reverse is true: PE-3 is DF in vES-23_401 for I-VPLS 401 and NDF in vES-23_402 for I-VPLS 402.

Within B-VPLS 100, the BMAC addresses are advertised via BGP-EVPN. On PE-2, the following FDB for B-VPLS 100 contains the BMAC addresses of PE-3 and PE-4, which are advertised via BGP-EVPN:

```
[/]
A:admin@PE-2# show service id 100 fdb detail

===============================================================================
Forwarding Database, Service 100
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
100        00:00:00:00:00:03 mpls:                   EvpnS:P  06/08/21 15:01:37
                             192.0.2.3:524279
           ldp:65540
100        00:00:00:00:00:04 mpls:                   EvpnS:P  06/08/21 15:01:37
                             192.0.2.4:524283
           ldp:65538
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Likewise, the following FDB for B-VPLS 100 on PE-3 contains the BMAC addresses of PE-2 and PE-4:

```
[/]
A:admin@PE-3# show service id 100 fdb detail

===============================================================================
Forwarding Database, Service 100
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
100        00:00:00:00:00:02 mpls:                   EvpnS:P  06/08/21 15:16:08
                             192.0.2.2:524283
           ldp:65537
100        00:00:00:00:00:04 mpls:                   EvpnS:P  06/08/21 15:16:08
                             192.0.2.4:524283
           ldp:65539
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

The following FDB for B-VPLS 100 on PE-4 contains the BMAC addresses of PE-2 and PE-3, but also the BMAC addresses of vES-23_401 and vES-23_402:

```
[/]
A:admin@PE-4# show service id 100 fdb detail

===============================================================================
Forwarding Database, Service 100
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
100        00:00:00:00:00:02 mpls:                   EvpnS:P  06/08/21 14:09:43
                             192.0.2.2:524283
           ldp:65538
100        00:00:00:00:00:03 mpls:                   EvpnS:P  06/08/21 14:50:11
                             192.0.2.3:524279
           ldp:65540
100        00:00:00:00:23:41 eES:                    EvpnS:P  06/08/21 14:50:02
                             MAX-ESI
100        00:00:00:00:23:42 eES:                    EvpnS:P  06/08/21 14:50:02
                             MAX-ESI
-------------------------------------------------------------------------------
No. of MAC Entries: 4
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

On PE-4, the following list of BGP EVPN routes for ES-BMAC 00:00:00:00:23:41 of vES-23_401 shows that PE-4 learned the ES-BMAC address via two PEs: PE-2 and PE-3.

```
[/]
A:admin@PE-4# show router bgp routes evpn mac mac-address 00:00:00:00:23:41
===============================================================================
 BGP Router ID:192.0.2.4        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN MAC Routes
===============================================================================
Flag  Route Dist.         MacAddr           ESI
      Tag                 Mac Mobility      Label1
                          Ip Address
                          NextHop
-------------------------------------------------------------------------------
u*>i  192.0.2.2:100       00:00:00:00:23:41 ESI-MAX
      0                   Static            LABEL 524283
                          n/a
                          192.0.2.2

u*>i  192.0.2.3:100       00:00:00:00:23:41 ESI-MAX
      0                   Static            LABEL 524279
                          n/a
                          192.0.2.3

-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

PE-4 also learned ES-BMAC 00:00:00:00:23:42 via PE-2 and PE-3, as follows:

```
[/]
A:admin@PE-4# show router bgp routes evpn mac mac-address 00:00:00:00:23:42
===============================================================================
 BGP Router ID:192.0.2.4        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN MAC Routes
===============================================================================
Flag  Route Dist.         MacAddr           ESI
      Tag                 Mac Mobility      Label1
                          Ip Address
                          NextHop
-------------------------------------------------------------------------------
u*>i  192.0.2.2:100       00:00:00:00:23:42 ESI-MAX
      0                   Static            LABEL 524283
                          n/a
                          192.0.2.2

u*>i  192.0.2.3:100       00:00:00:00:23:42 ESI-MAX
      0                   Static            LABEL 524279
                          n/a
                          192.0.2.3

-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

When a ping is initiated from CE-17 to CE-47, the ICMP packets are forwarded from PE-1 to PE-2, because the best route to 23.23.23.4 is via PE-2. PE-2 learns MAC address ca:fe:01:17:17:17 of CE-17 on the local I-VPLS SAP. PE-2 forwards the ICMP packets through I-VPLS 401 and B-VPLS 100 toward PE-4.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

PE-4 learns MAC ca:fe:01:17:17:17 of CE-17 via the ES-BMAC. When the reply is sent, PE-4 learns MAC address ca:fe:04:47:47:47 of CE-47 on the local SAP.

The FDB for I-VPLS 401 on PE-2 shows that MAC ca:fe:04:47:47:47 is learned on the local SAP and MAC ca:fe:04:47:47:47 can be reached via the B-VPLS to PE-4.

```
[/]
A:admin@PE-2# show service id 401 fdb detail

===============================================================================
Forwarding Database, Service 401
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
401        ca:fe:01:17:17:17 sap:lag-3:4.401         L/0      06/08/21 15:19:19
401        ca:fe:04:47:47:47 b-mpls:                 L/0      06/08/21 15:19:19
                             192.0.2.4:524283
           ldp:65538
-------------------------------------------------------------------------------
No. of MAC Entries: 2
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

The following FDB for I-VPLS 401 on PE-3 shows that MAC ca:fe:04:47:47:47 is learned via BGP-EVPN from PE-4.

```
[/]
A:admin@PE-3# show service id 401 fdb detail
===============================================================================
Forwarding Database, Service 401
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
401        ca:fe:04:47:47:47 b-mpls:                 L/0      06/08/21 15:19:19
                             192.0.2.4:524283
           ldp:65539
-------------------------------------------------------------------------------
No. of MAC Entries: 1
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

The following FDB for I-VPLS 401 on PE-4 shows that MAC ca:fe:04:47:47:47 is learned on a local SAP, whereas MAC ca:fe:01:17:17:17 is learned via ES-BMAC 00:00:00:00:23:41 of vES-23_401.

```
[/]
A:admin@PE-4# show service id 401 fdb detail
===============================================================================
Forwarding Database, Service 401
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
           Transport:Tnl-Id                          Age
-------------------------------------------------------------------------------
401        ca:fe:01:17:17:17 eES-BMAC:               L/0      06/08/21 15:19:19
                             00:00:00:00:23:41
401        ca:fe:04:47:47:47 sap:1/2/1:4.401         L/0      06/08/21 15:19:19
-------------------------------------------------------------------------------
No. of MAC Entries: 2
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

Conclusion

VXLAN FPE is required to terminate non-system IPv4/IPv6 VXLAN tunnels. The examples in this chapter show how VXLAN FPE can be applied in Epipe services, to stitch static VXLAN to other services, such as I-VPLS services.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Three-byte EVI in EVPN Services

This chapter provides information about the three-byte EVI in EVPN services.

Topics in this chapter include:

Applicability

The information and configuration in this chapter are based on SR OS Release 22.10.R1. The three-byte EVI is supported in EVPN services in SR OS Release 21.10.R1 and later. Three-byte EVI values can be configured in VPLS, R-VPLS, B-VPLS, and Epipe services for MPLS, VXLAN, and SRv6 instances.

Overview

In SR OS implementations earlier than SR OS Release 21.10.R1, the EVPN instance (EVI) is defined as a two-byte integer value, providing up to 65535 unique identifiers. The EVI is a unique value per service that can be used for three purposes:

- service route target (RT) auto-derivation – autonomous system number (ASN):EVI; for example, 64496:10
- service route distinguisher (RD) auto-derivation – system IP address:EVI; for example, 192.0.2.1:10
- designated forwarder (DF) election, as described in the Preference-based and Non-revertive EVPN DF Election chapter

However, in large networks, more than 65535 EVI values are required if the EVI is desired to be unique network-wide. The three-byte EVI provides up to 16777215 values and is supported in SR OS Release 21.10.R1 and later.

All DF election procedures support the extended EVI range. The RD auto-derivation is only possible for the two-byte EVI; the RT auto-derivation for the three-byte EVI can be enabled with the **evi-three-byte-auto-rd** command.

Auto-derived RT

The figure Figure 272: Auto-derived RT in RFC 8365shows the RT auto-derivation for configured EVI values in the range up to 16777215.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 272: Auto-derived RT in RFC 8365

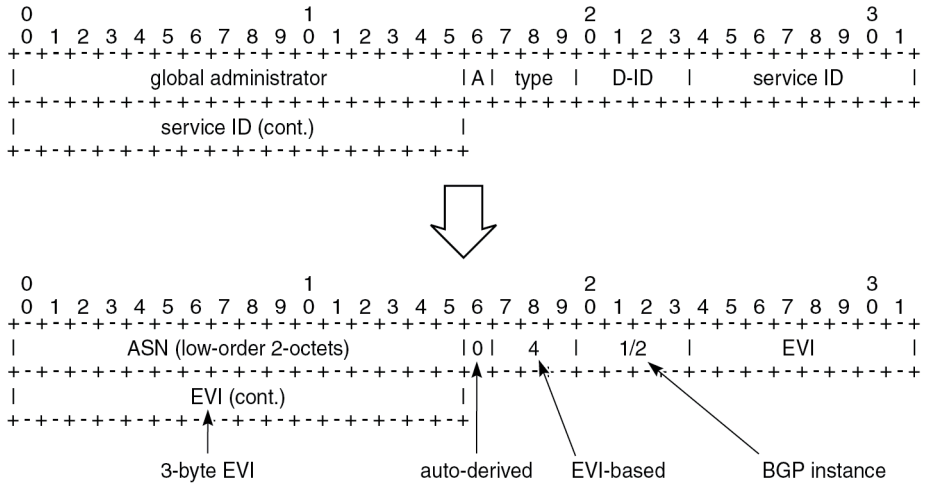

For three-octet EVI values, the fields in the RT format are:

- the global administrator field, which contains (the lower two octets of) the autonomous system number (ASN)
- the single-bit field A, which indicates if the RT is auto-derived: A=0 for auto-derivation
- the three-bit type field, which indicates the space in which the three-byte service ID is defined:
  - 0: VID (802.1Q VLAN ID)
  - 1: VXLAN
  - 2: NVGRE
  - 3: I-SID
  - **4: EVI**
  - 5: dual-VID (QinQ VLAN ID)
- the four-bit D-ID field, which encodes the domain ID. For type 4 (EVI), the D-ID corresponds to the BGP instance ID in the EVPN service.
- the three-octet service ID, which is set to the EVI (for type 4)

As an example, in a dual-instance EVPN-VPLS service with the following characteristics:

- ASN 64496
- EVI 100002 (0x186A2)
- BGP 1 for EVPN-VXLAN; BGP 2 for EVPN-MPLS
- **evi-three-byte-auto-rt** enabled

The two auto-derived RTs are:

- 64496:1090619042 (0x4**1**0186A2) for BGP 1
- 64496:1107396258 (0x4**2**0186A2) for BGP 2

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The RT can also be configured manually, for example, 64496:100002. A manually configured RT has precedence over an auto-derived RT.

Auto-derived RD

Each BGP instance in an EVPN service has an RD. Only for EVI values smaller than or equal to 65535, the RD for BGP instance 1 can be auto-derived out of the system IP address and the EVI, for example, 192.0.2.2:10. EVI values greater than 65535 do not generate RDs automatically.

The VPLS RD is selected based on the following precedence order:

- manually configured RD or auto-RD take precedence when configured,
- if there is no manual RD or auto-RD configuration, the RD is derived from the **bgp-ad>vpls-id**,
- if there is no manual RD, auto-RD, or VPLS ID configuration, the RD is derived from the EVI for EVI values up to 65535 and except for **bgp-mh** which does not support EVI-derived RD,

The Epipe RD is determined in a similar way, but there is no VPLS ID in Epipes.

The following error messages are raised when attempting to enable **bgp-evpn** with an EVI value greater than 65535 without having configured a manual RD, auto-RD, or BGP-AD VPLS ID:

```
*[ex:/configure service vpls "VPLS-99"]
A:admin@PE-1# commit
MINOR: MGMT_CORE #4001: configure service vpls "VPLS-99" bgp-evpn vxlan 1 admin-state - No
 route-distinguisher configured - configure service vpls "VPLS-99" bgp 1 route-distinguisher
MINOR: MGMT_CORE #4001: configure service vpls "VPLS-99" bgp-evpn vxlan 1 admin-state - No
 import route-target configured - configure service vpls "VPLS-99" bgp 1 route-target import
MINOR: MGMT_CORE #4001: configure service vpls "VPLS-99" bgp-evpn vxlan 1 admin-state - No
 export route-target configured - configure service vpls "VPLS-99" bgp 1 route-target export
MINOR: MGMT_CORE #4001: configure service vpls "VPLS-99" bgp-evpn vxlan 1 admin-state - No vsi-
import configured - configure service vpls "VPLS-99" bgp 1 vsi-import
MINOR: MGMT_CORE #4001: configure service vpls "VPLS-99" bgp-evpn vxlan 1 admin-state - No vsi-
export configured - configure service vpls "VPLS-99" bgp 1 vsi-import
MINOR: MGMT_CORE #4001: configure service vpls "VPLS-99" bgp-evpn vxlan 1 admin-state - No bgp-
ad vpls-id configured to derive RD or RT - configure service
MINOR: MGMT_CORE #4001: configure service vpls "VPLS-99" bgp-evpn vxlan 1 admin-state - No bgp-
evpn evi to derive RD or RT - configure service vpls "VPLS-99" bgp-evpn evi
```

The RD configuration can be changed dynamically. When the RD changes, the active routes for the service are withdrawn and readvertised with the new RD.

EVI RT set for AD per-ES routes

As described in the EVPN for MPLS Tunnels chapter, Auto-discovery per Ethernet Segment (AD per-ES) routes carry the ESI label and the multi-homing mode. When multiple EVIs are defined in an ES, the AD per-ES routes can be aggregated.

EVI RT set for AD per-ES routes with two-byte EVI

The following command enables the aggregation of AD per-ES routes for two-byte EVI values:

```
configure {
    service {
        system {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            bgp {
                evpn {
                    ad-per-es-route {
                        route-target-type evi-route-target-set
                        route-distinguisher-ip-address <ip-address>
                    }
```

The RD is specific for this EVI RT set feature. If enabled, a single AD per-ES route with the associated RD and a set of maximum 128 EVI RTs can be advertised. The EVI RTs are distributed in routes with the RD configured in the preceding command and one of the following *comm-val* values (the *comm-val* range is not configurable):

- EVIs from 1 to 128 – *comm-val* = 1
- EVIs from 129 to 256 – *comm-val* = 2
- ...
- EVIs from 65409 to 65535 – *comm-val* = 512

EVI RT set for AD per-ES routes with three-byte EVI

The command to enable AD per-ES route aggregation with extended EVI range is:

```
configure {
    service {
        system {
            bgp {
                evpn {
                    ad-per-es-route {
                        route-target-type evi-route-target-set
                        route-distinguisher-ip-address <ip-address>
                        extended-evi-range true
                    }
```

For three-byte EVIs, the *comm-val* range is extended from 512 to 65535 and the maximum number of AD per-ES routes that can be aggregated is increased from 128 to 257. The 257 RTs per route packing is done for any configured EVI, regardless of the value being greater than 65535 or not.

- EVIs from 1 to 257 – *comm-val* = 1
- EVIs from 258 to 514 – *comm-val* = 2
- ...
- EVIs from 16776961 to 16777215 – *comm-val* = 65281

Configuration

The figure Figure 273: Example topology with dual-instance VPLS shows the example topology with dual-instance VPLS 2: VXLAN is used between PE-1, PE-2, PE-3 and also between PE-4, PE-5, and PE-6. MPLS is used between the core PEs PE-2, PE-3, PE-4, and PE-5.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 273: Example topology with dual-instance VPLS

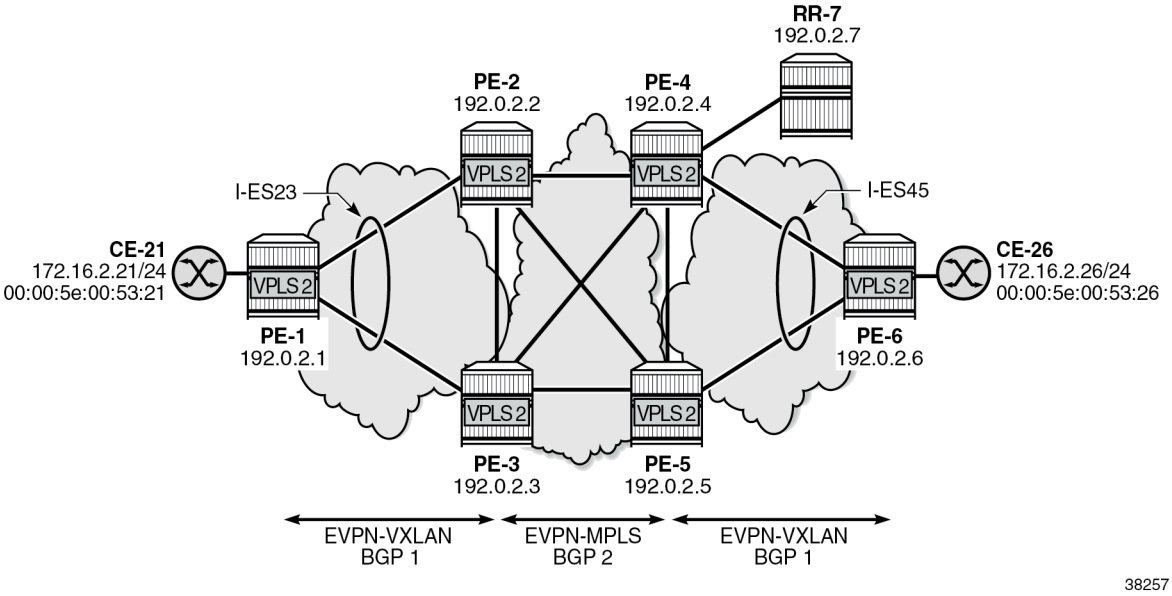

The initial configuration includes:

- cards, MDAs, ports
- router interfaces
- IS-IS level 2 between core PEs PE-2, PE-3, PE-4, PE-5, and RR-7
- IS-IS level 1 between PE-1, PE-2, and PE-3
- IS-IS level 1 between PE-6, PE-4, and PE-5
- SR-ISIS between core PEs PE-2, PE-3, PE-4, and PE-5

The BGP configuration and the used policies for dual-instance VPLSs in ESs are described in the EVPN Interconnect Ethernet Segmentschapter. Policies are required to prevent loops. RR-7 acts as route reflector for the core PEs PE-2, PE-3, PE-4, and PE-5. The policy and BGP configuration on PE-2 is as follows:

```
# on PE-2:
configure {
    policy-options {
        community "SOO-DCGW-23" {
            member "origin:64500:23" { }
        }
        community "vxlan" {
            member "bgp-tunnel-encap:VXLAN" { }
        }
        policy-statement "add SOO to vxlan routes" {
            entry 10 {
                from {
                    family [evpn]
                    community {
                        name "vxlan"
                    }
                }
                action {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    action-type accept
                    community {
                        add ["SOO-DCGW-23"]
                    }
                }
            }
            default-action {
                action-type accept
            }
        }
        policy-statement "allow only mpls" {
            entry 10 {
                from {
                    family [evpn]
                    community {
                        name "vxlan"
                    }
                }
                action {
                    action-type reject
                }
            }
        }
        policy-statement "allow only vxlan" {
            entry 10 {
                from {
                    family [evpn]
                    community {
                        name "vxlan"
                    }
                }
                action {
                    action-type accept
                }
            }
            default-action {
                action-type reject
            }
        }
        policy-statement "drop SOO-DCGW-23" {
            entry 10 {
                from {
                    family [evpn]
                    community {
                        name "SOO-DCGW-23"
                    }
                }
                action {
                    action-type reject
                }
            }
        }
    }
    router "Base" {
        autonomous-system 64496
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            split-horizon true
            rapid-update {
                evpn true
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            group "WAN" {
                peer-as 64496
                family {
                    evpn true
                }
                export {
                    policy ["allow only mpls"]
                }
            }
            group "access1" {
                peer-as 64496
                family {
                    evpn true
                }
                export {
                    policy ["allow only vxlan"]
                }
            }
            neighbor "192.0.2.1" {
                group "access1"
            }
            neighbor "192.0.2.3" {
                group "access1"
                import {
                    policy ["drop SOO-DCGW-23"]
                }
                export {
                    policy ["add SOO to vxlan routes"]
                }
            }
            neighbor "192.0.2.7" {
                group "WAN"
            }
```

The all-active interconnect ES "I-ES23" is configured on PE-2 and PE-3; the single-active interconnect ES "I-ES45" is configured on PE-4 and PE-5. VPLS 1 with EVI 1 (0x1) and VPLS 2 with EVI 100002 (0x186A2) are configured on all PEs. Both VPLSs have BGP 1 for VXLAN and BGP 2 for MPLS (SR-ISIS) in the core. For VPLS 1, no extended EVI range is required. The RD can be auto-derived for BGP instance 1, but not for BGP instance 2. For VPLS 2, the EVI is greater than 65535, so the RD must always be configured (manual configuration or auto-RD). The RT is auto-derived in VPLS 1 and VPLS 2. For VPLS 2, the **evi-three-byte-auto-rt** command is configured to enable auto-derivation of RTs for EVI values up to 16777215. On all core PEs, **evi-route-target-set** is enabled for the aggregation of AD per-ES routes. The service configuration on PE-2 is as follows:

```
# on PE-2:
configure exclusive
    service {
        system {
            bgp-auto-rd-range {
                ip-address 192.0.2.2
                community-value {
                    start 2000
                    end 2999
                }
            }
            bgp {
                evpn {
                    ethernet-segment "I-ES23" {
                        admin-state enable
                        type virtual
                        esi 00:00:00:00:00:23:23:00:00:01
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    multi-homing-mode all-active
                    df-election {
                        service-carving-mode manual
                        manual {
                            evi 1 {
                                end 200000
                            }
                            preference {
                                mode non-revertive
                                value 150
                            }
                        }
                    }
                    association {
                        network-interconnect-vxlan 1 {
                            virtual-ranges {
                                service-id 1 {
                                    end 2
                                }
                            }
                        }
                    }
                }
                ad-per-es-route {
                    route-target-type evi-route-target-set
                    route-distinguisher-ip-address 10.0.2.2
                    extended-evi-range true
                }
            }
        }
    }
    vpls "VPLS-1" {
        admin-state enable
        service-id 1
        customer "1"
        vxlan {
            instance 1 {
                vni 1
            }
        }
        bgp 1 {   # RD will be auto-derived from EVI
        }
        bgp 2 {
            route-distinguisher auto-rd
        }
        bgp-evpn {
            evi 1
            vxlan 1 {
                admin-state enable
                vxlan-instance 1
            }
            mpls 2 {
                admin-state enable
                ingress-replication-bum-label true
                ecmp 2
                auto-bind-tunnel {
                    resolution filter
                    resolution-filter {
                        sr-isis true
                    }
                }
            }
        }
    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        vpls "VPLS-2" {
            admin-state enable
            service-id 2
            customer "1"
            vxlan {
                instance 1 {
                    vni 2
                }
            }
            bgp 1 {
                route-distinguisher auto-rd    # RD cannot be auto-derived from EVI
            }
            bgp 2 {
                route-distinguisher auto-rd
            }
            bgp-evpn {
                evi 100002
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                    evi-three-byte-auto-rt true
                }
                mpls 2 {
                    admin-state enable
                    ingress-replication-bum-label true
                    ecmp 2
                    evi-three-byte-auto-rt true
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
        }
```

When configuring the **evi-route-target-set** command, the RD must be different from the RD in the auto-rd range. If not, the following error message is raised:

```
*[ex:/configure service system bgp evpn ad-per-es-route]
A:admin@PE-2# route-distinguisher-ip-address 192.0.2.2

*[ex:/configure service system bgp evpn ad-per-es-route]
A:admin@PE-2# commit
MINOR: SVCMGR #12: configure service system bgp-auto-rd-range ip-address - Inconsistent Value
 error - cannot be the same as ad-per-es-route/route-distinguisher-ip-address - configure
 service system bgp evpn ad-per-es-route route-distinguisher-ip-address
```

The following command shows the RD and RT values for both BGP instances in VPLS 1 on PE-2:

```
[/]
A:admin@PE-2# show service id 1 bgp

===============================================================================
BGP Information
===============================================================================
Bgp Instance         : 1
Vsi-Import           : None
Vsi-Export           : None
Route Dist           : None
Oper Route Dist      : 192.0.2.2:1
Oper RD Type         : derivedEvi
Rte-Target Import    : None                Rte-Target Export: None
Oper RT Imp Origin   : derivedEvi          Oper RT Import   : 64496:1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Oper RT Exp Origin   : derivedEvi           Oper RT Export   : 64496:1
ADV Service MTU      : -1

Bgp Instance         : 2
Vsi-Import           : None
Vsi-Export           : None
Route Dist           : auto-rd
Oper Route Dist      : 192.0.2.2:2000
Oper RD Type         : auto
Rte-Target Import    : None                 Rte-Target Export: None
Oper RT Imp Origin   : derivedEvi           Oper RT Import   : 64496:1
Oper RT Exp Origin   : derivedEvi           Oper RT Export   : 64496:1
ADV Service MTU      : -1

PW-Template Id       : None
-------------------------------------------------------------------------------
===============================================================================
```

RD 192.0.2.2:1 for BGP instance 1 is auto-derived whereas RD 192.0.2.2:2000 for BGP instance 2 is the result of auto-RD. RT 64496:1 is auto-derived based on the system IP address and the EVI. It is possible to configure **evi-three-byte-auto-rt true** in VPLS 1, even though the EVI value is smaller than 65535.

```
# on PE-2:
configure {
    service {
        vpls "VPLS-1" {
            bgp-evpn {
                vxlan 1 {
                    evi-three-byte-auto-rt true
                }
                mpls 2 {
                    evi-three-byte-auto-rt true
                }
            }
```

In this case, the auto-derivation is based on RFC 8365 and the value is 64496:1090519041 (0x41000001) for BGP 1 and 64496:1107296257 (0x42000001) for BGP 2.

```
[/]
A:admin@PE-2# show service id 1 bgp

===============================================================================
BGP Information
===============================================================================
Bgp Instance         : 1
Vsi-Import           : None
Vsi-Export           : None
Route Dist           : None
Oper Route Dist      : 192.0.2.2:1
Oper RD Type         : derivedEvi
Rte-Target Import    : None                 Rte-Target Export: None
Oper RT Imp Origin   : derivedEvi           Oper RT Import   : 64496:1090519041
Oper RT Exp Origin   : derivedEvi           Oper RT Export   : 64496:1090519041
ADV Service MTU      : -1

Bgp Instance         : 2
Vsi-Import           : None
Vsi-Export           : None
Route Dist           : auto-rd
Oper Route Dist      : 192.0.2.2:2000
Oper RD Type         : auto
Rte-Target Import    : None                 Rte-Target Export: None
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Oper RT Imp Origin   : derivedEvi           Oper RT Import   : 64496:1107296257
Oper RT Exp Origin   : derivedEvi           Oper RT Export   : 64496:1107296257
ADV Service MTU      : -1

PW-Template Id       : None
-------------------------------------------------------------------------------
===============================================================================
```

The auto-derived RTs on the other nodes are the same as on PE-2 when the BGP instance is the same. On PE-2, PE-3, PE-4, and PE-5, BGP 1 is for VXLAN and BGP 2 is for MPLS in the core. That way, EVPN messages can be exchanged in BGP instance 2 between the core PEs.

AD per-ES aggregation is enabled on the core nodes. The following command on PE-2 shows one AD per-ES with RD 10.0.2.4:390 (10.0.2.4 is the RD configured for **evi-route-target-set** and 390 is the *comm-val* value for the EVI range from 99974 to 100230) and one AD per-EVI with RD 192.0.2.4:2002 (auto-RD).

```
*A:PE-2# show router bgp routes evpn auto-disc detail
===============================================================================
 BGP Router ID:192.0.2.2        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Auto-Disc Routes
===============================================================================
---snip---
                           ## AD per-ES
Network        : n/a
Nexthop        : 192.0.2.4
Path Id        : None
From           : 192.0.2.7
Res. Nexthop   : 192.168.24.2
Local Pref.    : 100                    Interface Name : int-PE-2-PE-4
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : 10
Connector      : None
Community      : target:64496:1107396258    # auto-derived RT
                 esi-label:524271/Single-Active
Cluster        : 192.0.2.7
Originator Id  : 192.0.2.4              Peer Router Id : 192.0.2.7
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
EVPN type      : AUTO-DISC
ESI            : 00:00:00:00:00:45:45:00:00:01
Tag            : MAX-ET              # AD per-ES has MAX-ET
Route Dist.    : 10.0.2.4:390        # RD for evi-rt-set
MPLS Label     : LABEL 0
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
Add Paths Send : Default
Last Modified  : 00h02m16s
---snip---

                           ## AD per-EVI
Network        : n/a
Nexthop        : 192.0.2.4
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Path Id         : None
From            : 192.0.2.7
Res. Nexthop    : 192.168.24.2
Local Pref.     : 100                    Interface Name : int-PE-2-PE-4
Aggregator AS   : None                   Aggregator     : None
Atomic Aggr.    : Not Atomic             MED            : None
AIGP Metric     : None                   IGP Cost       : 10
Connector       : None
Community       : target:64496:1107396258 bgp-tunnel-encap:MPLS
Cluster         : 192.0.2.7
Originator Id   : 192.0.2.4              Peer Router Id : 192.0.2.7
Flags           : Used Valid Best IGP
Route Source    : Internal
AS-Path         : No As-Path
EVPN type       : AUTO-DISC
ESI             : 00:00:00:00:00:45:45:00:00:01
Tag             : 0                    # AD per-EVI has Ethernet tag 0
Route Dist.     : 192.0.2.4:2002       # RD for BGP 2 in VPLS 2
MPLS Label      : LABEL 524268
Route Tag       : 0
Neighbor-AS     : n/a
Orig Validation: N/A
Source Class    : 0                      Dest Class     : 0
Add Paths Send : Default
Last Modified   : 00h02m16s
---snip---
```

The figure Figure 274: Example topology with VPLS 4 and Epipe 5 shows an example topology with EVPN-MPLS in the core, all-active ES "vESI-23" on PE-2 and PE-3, and single-active ES "vESI-45" on PE-4 and PE-5. VPLS 4 with EVI 100004 is configured on all PEs with manually configured RD and RT.

Figure 274: Example topology with VPLS 4 and Epipe 5

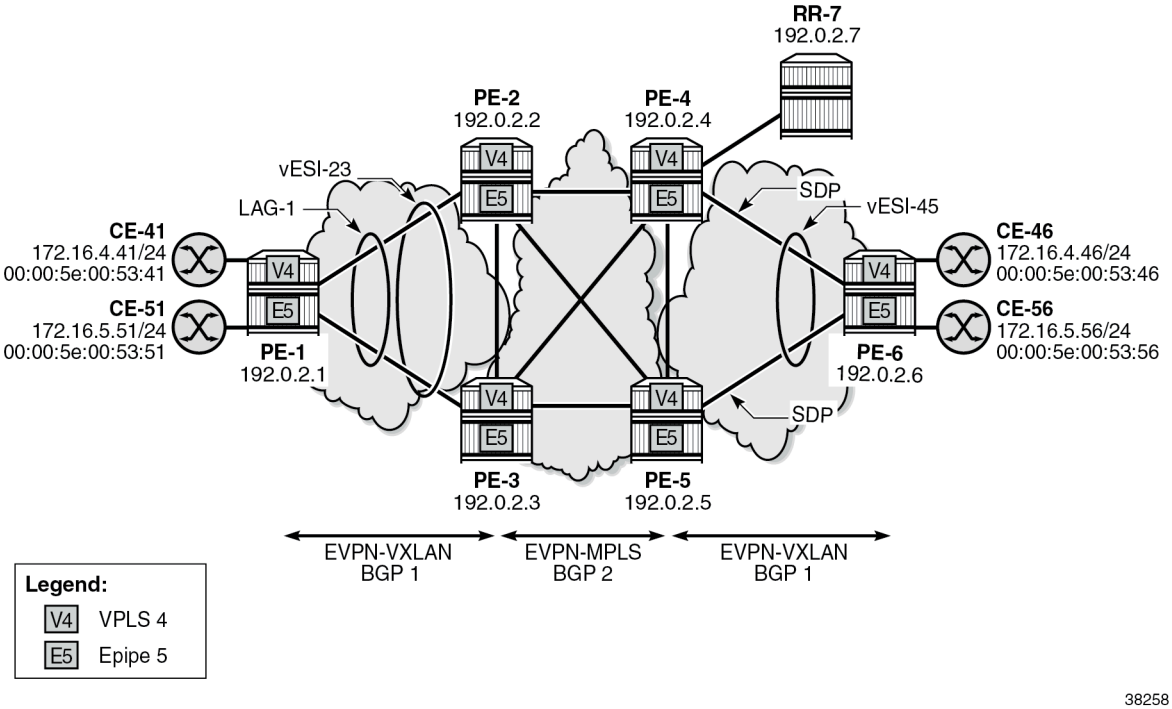

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The configuration on PE-2 is as follows:

```
# on PE-2:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "vESI-23" {
                        admin-state enable
                        type virtual
                        esi 01:00:00:00:00:23:03:09:00:01
                        multi-homing-mode all-active
                        df-election {
                            es-activation-timer 3
                        }
                        association {
                            lag "lag-1" {
                                virtual-ranges {
                                    dot1q {
                                        q-tag 3 {
                                            end 9
                                        }
                                    }
                                }
                            }
                        }
                    }
                }
            }
        }
        vpls "VPLS-4" {
            admin-state enable
            service-id 4
            customer "1"
            bgp 1 {
                route-distinguisher "192.0.2.2:4"
                route-target {
                    export "target:64496:100004"
                    import "target:64496:100004"
                }
            }
            bgp-evpn {
                evi 100004
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    ecmp 2
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            sap lag-1:4 {
            }
        }
```

With the configured RD 192.0.2.2:4 and RT 64496:100004, the following BGP information is retrieved on PE-2 for VPLS 4:

```
[/]
A:admin@PE-2# show service id 4 bgp
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
BGP Information
===============================================================================
Bgp Instance         : 1
Vsi-Import           : None
Vsi-Export           : None
Route Dist           : 192.0.2.2:4
Oper Route Dist      : 192.0.2.2:4
Oper RD Type         : configured
Rte-Target Import    : 64496:100004         Rte-Target Export: 64496:100004
Oper RT Imp Origin   : configured           Oper RT Import   : 64496:100004
Oper RT Exp Origin   : configured           Oper RT Export   : 64496:100004
ADV Service MTU      : -1

PW-Template Id       : None
-------------------------------------------------------------------------------
===============================================================================
```

Epipe 5 is configured on all PEs with EVI 100005 (0x186A5) and **evi-three-byte-auto-rt** enabled. The configuration on PE-2 is as follows:

```
# on PE-2:
configure {
    service {
        epipe "Epipe-5" {
            admin-state enable
            service-id 5
            customer "1"
            bgp 1 {
                route-distinguisher auto-rd
            }
            sap lag-1:5 {
            }
            bgp-evpn {
                evi 100005
                local-attachment-circuit "AC-ESI-23-PE-1" {
                    eth-tag 231
                }
                remote-attachment-circuit "AC-ESI-45-PE-6" {
                    eth-tag 456
                }
                mpls 1 {
                    admin-state enable
                    ecmp 2
                    evi-three-byte-auto-rt true
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
```

The auto-derived RT is 64496:1090619045 (0x410186A5):

```
[/]
A:admin@PE-2# show service id 5 bgp

===============================================================================
BGP Information
===============================================================================
Vsi-Import           : None
Vsi-Export           : None
Route Dist           : auto-rd
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Oper Route Dist      : 192.0.2.2:2003
Oper RD Type         : auto
Rte-Target Import    : None                Rte-Target Export: None
Oper RT Imp Origin   : derivedEvi          Oper RT Import   : 64496:1090619045
Oper RT Exp Origin   : derivedEvi          Oper RT Export   : 64496:1090619045
ADV Service MTU      : -1
PW-Template Id       : None
-------------------------------------------------------------------------------
===============================================================================
```

Instead of VXLAN or MPLS, SRv6 can be used too. As an example,VPLS 6 with EVI 100006 (0x186A6) is configured on PE-1, PE-2, PE-4, and PE-6. SRv6 is configured between PE-4 and PE-6, as described in the *Segment Routing over IPv6* chapter. The configuration on PE-2 is as follows:

```
# on PE-2:
configure {
    service {
        vpls "VPLS-6" {
            admin-state enable
            service-id 6
            customer "1"
            vxlan {
                instance 1 {
                    vni 6
                }
            }
            segment-routing-v6 1 {
                locator "PE-2_loc" {
                    function {
                        end-dt2u {
                        }
                        end-dt2m {
                        }
                    }
                }
            }
            bgp 1 {
                route-distinguisher auto-rd
            }
            bgp 2 {
                route-distinguisher auto-rd
            }
            bgp-evpn {
                evi 100006
                segment-routing-v6 2 {
                    admin-state enable
                    ecmp 2
                    evi-three-byte-auto-rt true
                    srv6 {
                        instance 1
                        default-locator "PE-2_loc"
                    }
                }
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                    evi-three-byte-auto-rt true
                }
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

On PE-2, RT 64496:1090619046 (0x410186A6) is auto-derived for BGP 1 and RT 64496:1107396262 (0x420186A6) for BGP 2:

```
[/]
A:admin@PE-2# show service id 6 bgp

===============================================================================
BGP Information
===============================================================================
Bgp Instance         : 1
Vsi-Import           : None
Vsi-Export           : None
Route Dist           : auto-rd
Oper Route Dist      : 192.0.2.2:2004
Oper RD Type         : auto
Rte-Target Import    : None                Rte-Target Export: None
Oper RT Imp Origin   : derivedEvi          Oper RT Import   : 64496:1090619046
Oper RT Exp Origin   : derivedEvi          Oper RT Export   : 64496:1090619046
ADV Service MTU      : -1

Bgp Instance         : 2
Vsi-Import           : None
Vsi-Export           : None
Route Dist           : auto-rd
Oper Route Dist      : 192.0.2.2:2005
Oper RD Type         : auto
Rte-Target Import    : None                Rte-Target Export: None
Oper RT Imp Origin   : derivedEvi          Oper RT Import   : 64496:1107396262
Oper RT Exp Origin   : derivedEvi          Oper RT Export   : 64496:1107396262
ADV Service MTU      : -1

PW-Template Id       : None
-------------------------------------------------------------------------------
===============================================================================
```

Conclusion

In large networks, a three-byte EVI can be required as a unique identifier for services. RTs can be auto-derived based on a three-byte EVI, but RDs cannot be auto-derived that way.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

VCCV BFD for Epipe Services

This chapter describes the VCCV BFD for Epipe services.

Topics in this chapter include:

Applicability

This chapter was initially written based on SR OS Release 15.0.R7. The MD-CLI in the current edition corresponds to SR OS Release 23.7.R1.

Overview

Virtual circuit connectivity verification (VCCV) is defined by RFC 5085. Bidirectional forwarding detection (BFD) is defined by RFC 5880.

VCCV is an end-to-end fault-detection tool for testing pseudowires (PWs), and typically requires an operator to take manual actions. The PWs can be used for virtual leased line (VLL), virtual private LAN service (VPLS), and Internet enhanced service (IES)/virtual private routed network (VPRN) services with Epipe or Ipipe spoke-SDPs.

SR OS supports RFC 5885 which specifies a method for carrying BFD messages in a PW-associated channel and is referred to as VCCV BFD in SR OS. Because the associated channel shares fate with the data plane, VCCV BFD monitors the PW between two terminating PEs (T-PEs), regardless of the number of provider routers or switching PEs (S-PEs) the PW may traverse; see Figure 275: PW reference model. When enabled, faults in individual PWs can be detected quickly, whether or not other provider routers or S-PEs also carry other PWs. VCCV BFD can monitor specific high-value services, where detecting forwarding failures (and potentially recovering from them) in a minimum amount of time is critical.

Figure 275: PW reference model

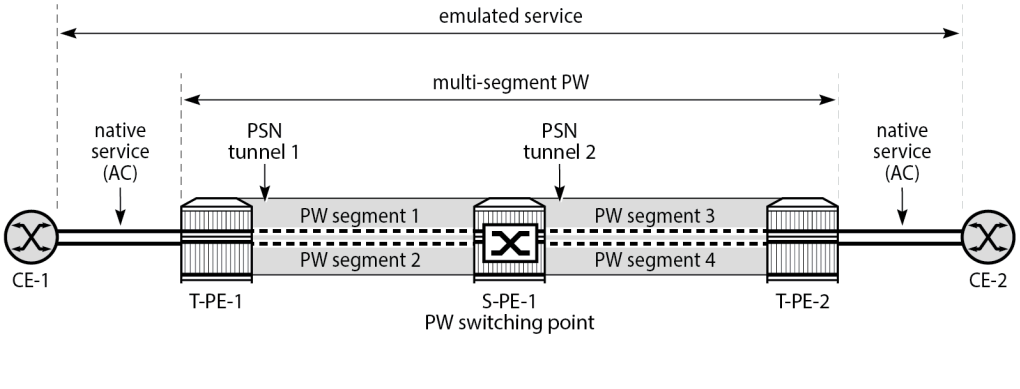

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

VCCV BFD avoids manual hop-by-hop troubleshooting of each element along the path of the PW, which minimizes the probability of not detecting silent failures on intermediate routers.

VCCV BFD sessions run end-to-end on a switched or single-hop PW, from T-PE to T-PE. They do not terminate on an intermediate S-PE; therefore, the TTL of the PW label on VCCV BFD packets is always set to 255, to ensure that the packets reach the far-end T-PE of a multi-segment PW.

BFD is only used for fault detection. While RFC 5885 provides a mode in which VCCV BFD can be used to signal PW status, this mode is only applicable for PWs that have no other status signaling mechanism in use. LDP status and static PW status signaling always take precedence over BFD-signaled PW status, and BFD-signaled PW status is not used on PWs that use LDP status or static PW status signaling mechanisms.

Configuration

Figure 276: Example topology shows the example topology with Epipes "Epipe-1" and "Epipe-2" using LSP "lsp-1" between PE-5 and PE-6 and Epipe "Epipe-3" using LSP "lsp-2" between PE-5 and S-PE PE-4 and LSP "lsp-3" between PE-4 and PE-6.

Figure 276: Example topology

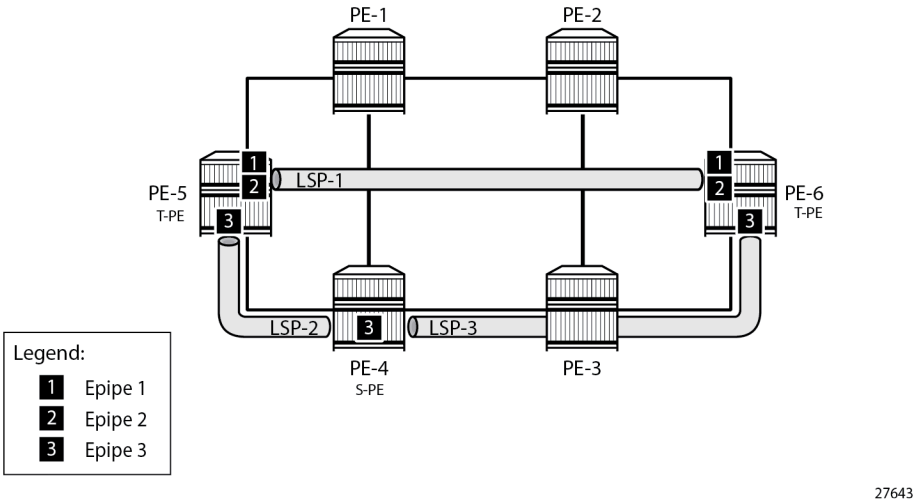

The initial configuration includes:

- Cards, MDAs, and ports
- Router interfaces
- IS-IS as IGP on all interfaces (alternatively, OSPF can be used), with traffic engineering enabled
- MPLS paths and LSPs:
  - LSP "lsp-1" configured on PE-5 with primary path "path-5-1-2-6" and on PE-6 with primary path "path-6-2-1-5"
  - LSP "lsp-2" configured on PE-5 with primary path "path-5-4" and on PE-4 with primary path "path-4-5"
  - LSP "lsp-3" configured on PE-4 with primary path "path-4-3-6" and on PE-6 with primary path "path-6-3-4"

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

VCCV BFD configuration

Three steps are needed when configuring VCCV BFD:

1. Configure the BFD template
2. Apply the BFD template
3. Enable BFD

Step 1: configure BFD template

The **bfd-template** command provides the control packet timer values for the BFD.

The general command to define a BFD template is as follows:

```
configure {
    bfd {
        bfd-template <name> {
            echo-receive <echo-interval>
            multiplier <multiplier>
            receive-interval <receive-interval>
            transmit-interval <transmit-interval>
            type {cpm-np}
```

However, network processor BFD (cpm-np) is not supported for VCCV, and the minimum supported receive or transmit timer interval is 100 ms. An error is generated if a user tries to apply a BFD template with the **type cpm-np** command or any unsupported transmit or receive interval value. An error is also generated when the user attempts to commit changes to a BFD template that is already bound to a spoke-SDP.

Steps 2 and 3: apply BFD template and enable BFD

To apply and enable the BFD template to a spoke-SDP where LDP is used as the SDP signaling protocol for a service, the following command can be used, depending on the service:

```
configure {
    service {
        [epipe|cpipe|ipipe|vpls|ies|vprn] <name> {
            spoke-sdp <sdp-binding-id> {
                bfd {
                    bfd-template <name>
                    bfd-liveness {
                    }
```

If BGP is used as the SDP signaling protocol, the following command is used:

```
configure {
    service {
        [epipe|vpls] <name> {
            bgp 1 {
                pw-template-binding <reference> {
                    bfd 1 {
                        bfd-template <name>
                        bfd-liveness true
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

In this example, the following BFD templates are configured on PE-5 and PE-6:

```
# on PE-5, PE-6:
configure {
    bfd {
        bfd-template "bfdt-1" {
            multiplier 5
            receive-interval 2000
            transmit-interval 2000
        }
        bfd-template "bfdt-2" {
            receive-interval 1000
            transmit-interval 1000
        }
```

These BFD templates are used in the Epipe services configured in the next section.

Service configuration

LDP VLL "Epipe-1"

The service "Epipe-1" is an LDP VLL running between PE-5 and PE-6, and uses manually configured SDP 56 on PE-5 and SDP 65 on PE-6, respectively, so the signaling is set to T-LDP. On PE-5, the spoke-SDP 56:1 has BFD template *bfdt-2* applied, and BFD is enabled. The configuration on PE-6 is similar.

```
# on PE-5:
configure {
    service {
        epipe "Epipe-1" {
            admin-state enable
            service-id 1
            customer "1"
            spoke-sdp 56:1 {
                bfd {
                    bfd-template "bfdt-2"
                    bfd-liveness {
                    }
                }
            }
            sap 1/1/c4/1:1 {
            }
        }
        sdp 56 {
            admin-state enable
            delivery-type mpls
            far-end {
                ip-address 192.0.2.6
            }
            lsp "lsp-1" { }
        }
```

Log 99 indicates the local discriminator value used for the VCCV BFD session, as follows:

```
95 2023/08/21 08:23:10.590 CEST MINOR: VRTR #2070 Base 127.0.0.1
"The vccv BFD session with Local Discriminator 1 on Svc 1 SdpBind 56:1 is up"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Configuring the spoke-SDP with a template with an invalid type (for example, type "cpm-np") or with invalid transmit or receive intervals leads to errors, as follows:

```
[ex:/configure service epipe "Epipe-1" spoke-sdp 56:1 bfd]
A:admin@PE-5# bfd-template "bfdt-cpm-np-50ms"

*[ex:/configure service epipe "Epipe-1" spoke-sdp 56:1 bfd]
A:admin@PE-5# commit
MINOR: MGMT_CORE #4001: configure service epipe "Epipe-1" spoke-sdp 56:1 bfd
bfd-template
- bfd-template transmit-interval must be minimum 100 for this application
- configure bfd bfd-template "bfdt-cpm-np-50ms" transmit-interval
MINOR: MGMT_CORE #4001: configure service epipe "Epipe-1" spoke-sdp 56:1 bfd
bfd-template
- bfd-template receive-interval must be minimum 100 for this application
- configure bfd bfd-template "bfdt-cpm-np-50ms" receive-interval
MINOR: MGMT_CORE #4001: configure service epipe "Epipe-1" spoke-sdp 56:1 bfd
bfd-template
- bfd-template type is not valid for this application
- configure bfd bfd-template "bfdt-cpm-np-50ms" type
```

BGP VPWS "Epipe-2"

The service "Epipe-2" is a BGP VPWS, also running between PE-5 and PE-6, using the manually configured SDPs 561 and 651 with the signaling set to BGP, as follows. Again, BFD template *bfdt-2* is used, but now the BFD template is referred to from the **pw-template-binding** context. See the BGP Virtual Private Wire Services chapter for more information.

```
# PE-5:
configure {
    service {
        pw-template "1" {
            provisioned-sdp prefer
        }
        epipe "Epipe-2" {
            admin-state enable
            service-id 2
            customer "1"
            bgp 1 {
                route-distinguisher "65545:2"
                route-target {
                    export "target:65545:2"
                    import "target:65545:2"
                }
                pw-template-binding "1" {
                    bfd-template "bfdt-2"
                    bfd-liveness true
                }
            }
            bgp-vpws {
                admin-state enable
                local-ve {
                    name "PE-5"
                    id 5
                }
                remote-ve "PE-6" {
                    id 6
                }
            }
            sap 1/1/c4/1:2 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
        }
        sdp 561 {
            admin-state enable
            delivery-type mpls
            signaling bgp
            far-end {
                ip-address 192.0.2.6
            }
            lsp "lsp-1" { }
        }
```

LDP VLL "Epipe-3" with switching node PE-4

The service "Epipe-3" is another LDP VLL running between PE-5 and PE-6, but switched at PE-4. It uses the manually configured SDPs 54 and 45 between PE-5 and PE-4, and SDPs 46 and 64 between PE-4 and PE-6. All these SDPs are using T-LDP for the signaling. On PE-5, the spoke-SDP 54:3 has BFD template "bfdt-1" applied, and **control-word** is active. This ensures that BFD packets get into the PW mapping to that spoke-SDP and that these packets are forwarded between the VC-switched spoke-SDPs at PE-4. The configuration on PE-6 is similar.

```
# on PE-5:
configure {
    bfd {
        bfd-template "bfdt-1" {
            multiplier 5
            receive-interval 2000
            transmit-interval 2000
        }
    }
    service {
        epipe "Epipe-3" {
            admin-state enable
            service-id 3
            customer "1"
            spoke-sdp 54:3 {
                control-word true
                bfd {
                    bfd-template "bfdt-1"
                    bfd-liveness {
                    }
                }
            }
            sap 1/1/c4/1:3 {
            }
        }
        sdp 54 {
            admin-state enable
            delivery-type mpls
            far-end {
                ip-address 192.0.2.4
            }
            lsp "lsp-2" { }
        }
```

For PE-4 to switch traffic from one VC to another, **vc-switching true** is configured in the Epipe, as follows:

```
# on PE-4:
configure {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    service {
        epipe "Epipe-3" {
            admin-state enable
            service-id 3
            customer "1"
            vc-switching true      # S-PE
            spoke-sdp 45:3 {
            }
            spoke-sdp 46:3 {
            }
        }
        sdp 45 {
            admin-state enable
            delivery-type mpls
            far-end {
                ip-address 192.0.2.5
            }
            lsp "lsp-2" { }
        }
        sdp 46 {
            admin-state enable
            delivery-type mpls
            far-end {
                ip-address 192.0.2.6
            }
            lsp "lsp-3" { }
        }
```

VCCV BFD verification

The following command shows that BFD template "bfdt-2" is applied to SDP 56:1 in the "Epipe-1" service and to SDP 561:4294967295 in the "Epipe-2" service:

```
[/]
A:admin@PE-5# show router bfd bfd-template "bfdt-2"

===============================================================================
BFD Template bfdt-2
===============================================================================
Template Name           : bfdt-2       Template Type             : auto
Transmit Timer          : 1000 msec    Receive Timer             : 1000 msec
Template Multiplier     : 3            Echo Receive Interval     : 100 msec

LSP-LDP  Association Count : 0
LSP-RSVP Association Count : 0
LSP-RSVP Template Association Count : 0
LSP-SR-TE Association Count : 0
LSP-SR-TE Template Association Count : 0
LSP-SR-TE Association Count : 0
LSP-SR-TE Template Association Count : 0
Static SR-Policy Association Count : 0
BGP SR-Policy Association Count : 0

Mpls-tp Association
None

-------------------------------------------------------------------------------
Service Associations
-------------------------------------------------------------------------------
SvcId               Sdp Bind                      BFD Enable     BFD Encap
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
1                   56:1                            yes            ipv4
2                   561:4294967295                  yes            ipv4
===============================================================================
```

The BFD configuration for SDP 56:1 on the "Epipe-1" service is listed in the detailed output for the SDP, as follows. The BFD template used is *bfdt-2*, BFD is enabled, and the BFD encapsulation used is IPv4. The peer VCCV CV bits indicate that the remote end supports LSP ping as well as BFD fault detection.

```
[/]
A:admin@PE-5# show service id 1 sdp 56:1 detail

===============================================================================
Service Destination Point (Sdp Id : 56:1) Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 56:1  -(192.0.2.6)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id             : 56:1                       Type              : Spoke
Spoke Descr     : (Not Specified)
VC Type            : Ether                      VC Tag            : n/a
Admin Path MTU     : 0                          Oper Path MTU     : 8974
Delivery           : MPLS
Far End            : 192.0.2.6                  Tunnel Far End    : n/a
Oper Tunnel Far End: 192.0.2.6
LSP Types          : RSVP
Hash Label         : Disabled                   Hash Lbl Sig Cap  : Disabled
Oper Hash Label    : Disabled
Entropy Label      : Disabled

Admin State        : Up                         Oper State        : Up
MinReqd SdpOperMTU : 1514
Adv Service MTU    : n/a
Acct. Pol          : None                       Collect Stats     : Disabled
Ingress Label      : 524284                     Egress Label      : 524284
Ingr Mac Fltr-Id   : n/a                        Egr Mac Fltr-Id   : n/a
Ingr IP Fltr-Id    : n/a                        Egr IP Fltr-Id    : n/a
Ingr IPv6 Fltr-Id  : n/a                        Egr IPv6 Fltr-Id  : n/a
Admin ControlWord  : Not Preferred              Oper ControlWord  : False
Admin BW(Kbps)     : 0                          Oper BW(Kbps)     : 0
BFD Template       : bfdt-2
BFD-Enabled        : yes                        BFD-Encap         : ipv4
BFD Fail Action    : none                       BFD Oper State    : connected
BFD WaitForUpTimer : 0 secs
BFD Time Remain    : 0 secs
Last Status Change : 08/21/2023 08:23:09        Signaling         : TLDP
Last Mgmt Change   : 08/21/2023 08:27:14
Endpoint           : N/A                        Precedence        : 4
ICB                : False
PW Status Sig      : Enabled
Force Vlan-Vc      : Disabled                   Force Qinq-Vc     : none
Class Fwding State : Down
Flags              : None
Local Pw Bits      : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : lspPing bfdFaultDet
Peer Vccv CC Bits  : mplsRouterAlertLabel
---snip---

-------------------------------------------------------------------------------
Control Channel Status
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
PW Status          : disabled                Refresh Timer     : <none>
Peer Status Expire : false
Request Timer      : <none>
Acknowledgement    : false

-------------------------------------------------------------------------------
---snip---

-------------------------------------------------------------------------------
RSVP/Static LSPs
-------------------------------------------------------------------------------
Associated LSP List :
Lsp Name           : lsp-1
Admin State        : Up                      Oper State        : Up
Time Since Last Tr*: 00h10m42s

-------------------------------------------------------------------------------
---snip---

-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

The full set of VCCV BFD sessions running with the currently used parameters can be shown as follows:

```
[/]
A:admin@PE-5# show service vccv-bfd

===============================================================================
BFD Session
===============================================================================
Svc-Id                                          State      Tx Pkts    Rx Pkts
  Sdp-Id:Vc-Id                                Multipl     Tx Intvl   Rx Intvl
  Protocols                                      Type     LAG Port     LAG ID
                                                                     LAG name
-------------------------------------------------------------------------------
1                                                  Up          88         90
  56:1                                              3        1000       1000
  vccv                                          central         N/A        N/A
  127.0.0.2
2                                                  Up         218        219
  561:4294967295                                    3        1000       1000
  vccv                                          central         N/A        N/A
  127.0.0.2
3                                                  Up          93         91
  54:3                                              5        2000       2000
  vccv                                          central         N/A        N/A
  127.0.0.2
-------------------------------------------------------------------------------
No. of System BFD sessions: 3
===============================================================================
```

The VCCV BFD sessions for a single service can be shown as follows:

```
[/]
A:admin@PE-5# show service id 3 vccv-bfd session

===============================================================================
BFD Session
===============================================================================
Svc-Id                                          State      Tx Pkts    Rx Pkts
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
  Sdp-Id:Vc-Id                                  Multipl     Tx Intvl   Rx Intvl
  Protocols                                        Type     LAG Port     LAG ID
                                                                       LAG name
-------------------------------------------------------------------------------
3                                                    Up          109        108
  54:3                                                5         2000       2000
  vccv                                           central          N/A        N/A
  127.0.0.2
-------------------------------------------------------------------------------
No. of BFD sessions: 1
===============================================================================
```

Similar output can be obtained on PE-6.

Disconnecting the link between PE-1 and PE-2 affects the traffic taking the upper path; the VCCV BFD sessions for the services "Epipe-1" and "Epipe-2" go down, and so do the SDPs and the services. This is reflected in log 99, as follows:

```
# on PE-5:
132 2023/08/21 08:29:40.782 CEST WARNING: MPLS #2012 Base VR 1:
"LSP path lsp-1::path-5-1-2-6 is operationally disabled ('shutdown') because
resvTear"

133 2023/08/21 08:29:40.782 CEST WARNING: MPLS #2010 Base VR 1:
"LSP lsp-1 is operationally disabled ('shutdown') because noPathIsOperational"

134 2023/08/21 08:29:40.783 CEST MINOR: SVCMGR #2303 Base
"Status of SDP 56 changed to admin=up oper=down"

135 2023/08/21 08:29:40.783 CEST MINOR: SVCMGR #2303 Base
"Status of SDP 561 changed to admin=up oper=down"

136 2023/08/21 08:29:40.783 CEST MINOR: SVCMGR #2326 Base
"Status of SDP Bind 56:1 in service 1 (customer 1) local PW status bits changed
to psnIngressFault psnEgressFault "

137 2023/08/21 08:29:40.784 CEST MAJOR: SVCMGR #2316 Base
"Processing of a SDP state change event is finished and the status of all affected
 SDP Bindings on SDP 561 has been updated."

138 2023/08/21 08:29:40.785 CEST MAJOR: SVCMGR #2316 Base
"Processing of a SDP state change event is finished and the status of all affected
 SDP Bindings on SDP 56 has been updated."

139 2023/08/21 08:29:43.777 CEST MINOR: VRTR #2069 Base 127.0.0.1
"The vccv BFD session with Local Discriminator 2 on Svc 2 SdpBind 561:4294967295
 is down due to noHeartBeat "

140 2023/08/21 08:29:43.787 CEST MINOR: VRTR #2069 Base 127.0.0.1
"The vccv BFD session with Local Discriminator 5 on Svc 1 SdpBind 56:1 is down
 due to noHeartBeat "

141 2023/08/21 08:29:47.318 CEST MINOR: SVCMGR #2313 Base
"Status of SDP Bind 56:1 in service 1 (customer 1) peer PW status bits changed
to psnIngressFault psnEgressFault "
```

This status of the VCCV BFD sessions then is as follows:

```
[/]
A:admin@PE-5# show service vccv-bfd

===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
BFD Session
===============================================================================
Svc-Id                                          State      Tx Pkts    Rx Pkts
  Sdp-Id:Vc-Id                                Multipl     Tx Intvl   Rx Intvl
  Protocols                                      Type     LAG Port     LAG ID
                                                                     LAG name
-------------------------------------------------------------------------------
1                                                Down          168        169
  56:1                                              3         1000       1000
  vccv                                         central          N/A        N/A
  127.0.0.2
2                                                Down          298        299
  561:4294967295                                    3         1000       1000
  vccv                                         central          N/A        N/A
  127.0.0.2
3                                                  Up          223        222
  54:3                                              5         2000       2000
  vccv                                         central          N/A        N/A
  127.0.0.2
-------------------------------------------------------------------------------
No. of System BFD sessions: 3
===============================================================================
```

Consequently, the "Epipe-1" and "Epipe-2" services are operationally down, as follows:

```
[/]
A:admin@PE-5# show service service-using epipe

===============================================================================
Services [epipe]
===============================================================================
ServiceId    Type      Adm  Opr  CustomerId Service Name
-------------------------------------------------------------------------------
1            Epipe     Up   Down 1          Epipe-1
2            Epipe     Up   Down 1          Epipe-2
3            Epipe     Up   Up   1          Epipe-3
-------------------------------------------------------------------------------
Matching Services : 3
-------------------------------------------------------------------------------
===============================================================================
```

Conclusion

VCCV BFD can monitor specific high-value services, where detecting forwarding failures (and potentially recovering from them) in the minimal amount of time is critical. VCCV BFD complements other on-demand tools such as VCCV ping and VCCV trace by providing proactive detection of faults. VCCV ping and VCCV trace can later be used to localize and diagnose the root cause of the fault.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Virtual Ethernet Segments

This chapter provides information about Virtual Ethernet Segments.

Topics in this chapter include:

- Applicability
- Overview
- Configuration
- Conclusion

Applicability

This chapter was initially written based on SR OS Release 15.0.R3, but the MD-CLI in the current edition is based on SR OS Release 21.2.R2. Virtual Ethernet segments are supported in SR OS Release 15.0.R1, and later.

Overview

RFC 7432 describes the use and procedures for Ethernet segments (ESs) that can be associated with physical Ethernet ports and LAGs. The SR OS implementation also allows an ES to be associated with SDPs. ESs meet the redundancy requirements of directly connected CEs. However, ESs will not work when an aggregation network exists between CEs and ES PEs, which requires different ESs to be defined for the port, LAG, or SDP. *Draft-ietf-bess-evpn-virtual-eth-segment* describes how virtual ESs (vESs) can be defined with an Attachment Circuit (AC) level granularity. Figure 277: vESs for PWs shows an example where vES definition at the pseudowire (PW) granularity level is required:

Figure 277: vESs for PWs

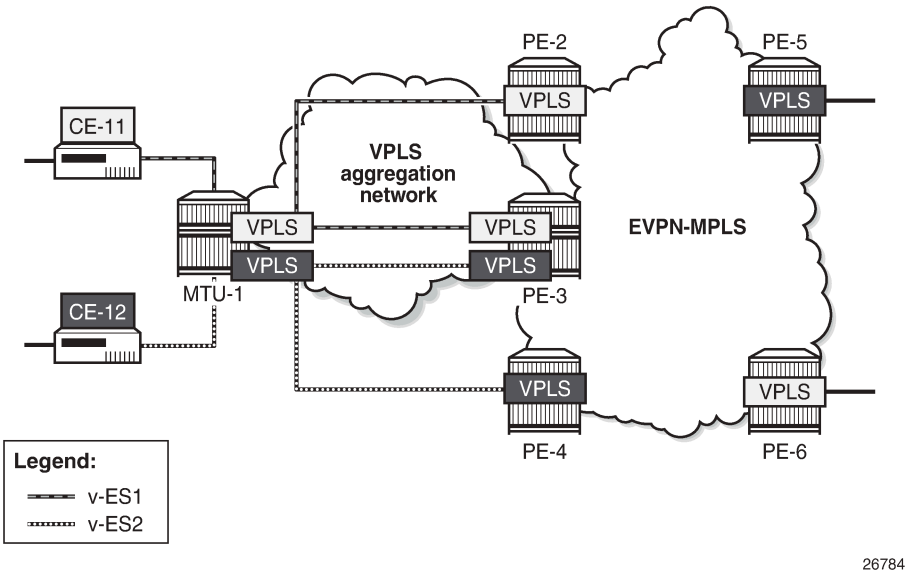

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

When a Layer 2 aggregation network is used to get access to EVPN, the association of ACs that belong to the same ES and physical ports or SDPs can be arbitrary. For example, the SDP between MTU-1 and PE-3 (Figure 277: vESs for PWs) cannot be associated with only one ES, because it is being used by two different CEs that require different ESs. The association must be at spoke-SDP level. The RFC 7432 port/lag-based ES definition is not sufficient, so vESs need to be defined. Virtual ESs can be configured with up to eight ranges of one or more:

- VC-IDs (spoke-SDPs)
- Q-tags (dot1q)
- S-tags (qinq)
- C-tags for a fixed S-tag (qinq)

Mesh-SDPs are not allowed for an SDP used by a vES.

Virtual ESs are configured as Ethernet segments of type virtual:

```
*[ex:/configure service system bgp evpn ethernet-segment "ESI-1"]
A:admin@PE-2# type ?

 type <keyword>
 <keyword>  - (none|virtual)
 Default    - none

 'type' is: immutable

    Type of the ethernet segment.

    Warning: Modifying this element recreates
    'configure service system bgp evpn ethernet-segment "ESI-1"' automatically for the
 new value to take effect.
```

Virtual ES "vESI-23_600" is associated with LAG 1 and one service-delimiting VLAN range is defined for the S-tag, as follows:

```
# on PE-2, PE-3:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "vESI-23_600" {
                        admin-state enable
                        type virtual
                        esi 01:00:00:00:00:23:06:00:00:01
                        multi-homing-mode all-active
                        df-election {
                            es-activation-timer 3
                            service-carving-mode manual
                            manual {
                                evi 2 {
                                    end 2
                                }
                            }
                        }
                        association {
                            lag "lag-1" {
                            virtual-ranges {
                                qinq {
                                    s-tag 600 {
                                        end 602
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                                        }
                                    }
                                }
                            }
                        }
                    }
                }
```

The configured ES will match all the SAPs for which the top (outer) service-delimiting tag is within the 600 to 602 range.

When the ES is created as virtual, a port, LAG, or SDP needs to be created before any VLAN or VC-ID can be associated.

- For VC-ID, only spoke-SDPs are allowed, no mesh-SDPs. Manual spoke-SDP VC-IDs and BGP-AD VC-IDs can be included in the range.
- For dot1q, only those SAPs that match the service-delimiting VLAN range will be associated with the vES
- For qinq, the following two commands can be configured, with a mutually exclusive S-tag:
  - **s-tag <qtag1> end <qtag1>** - associates all qinq SAPs with outer tag between the configured qtags.
  - **s-tag-c-tag <qtag1> c-tag-start <qtag2> c-tag-end <qtag2>** - associates all qinq SAPs with outer qtag1 and inner qtag between the configured qtag2 values to the vES

  A mutually exclusive S-tag means that a value for the S-tag can be configured in either of the two commands, but not in both.

Table 14: Supported examples for Q-tag values between 1 and 4094 shows the supported examples for qtag values between 1 and 4094; Table 15: Supported examples for Q-tag values 0, *, and null shows the supported examples for qtag values 0, *, and null:

Table 14: Supported examples for Q-tag values between 1 and 4094

vES configuration for port 1/1/1	SAP association
dot1q qtag 100	1/1/1:100
dot1q qtag-range 100 to 102	1/1/1:100, 1/1/1:101, 1/1/1:102
qinq s-tag 100 c-tag 200	1/1/1:100.200
qinq s-tag 100 c-tag-range 200 to 202	1/1/1:100.200, 1/1/1:100.201, 1/1/1:100.202
qinq s-tag 100	All SAPs 1/1/1:100.x (x being 1 to 4094, 0, or *)
qinq s-tag-range 100 to 102	All SAPs 1/1/1:100.x, 1/1/1:101.x, 1/1/1:102.x (x being 1 to 4094, 0, or *)

*Table 15: Supported examples for Q-tag values 0, *, and null*

vES configuration for port 1/1/1	SAP association
dot1q qtag 0	1/1/1:0

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

vES configuration for port 1/1/1	SAP association
dot1q qtag *	1/1/1:*
qinq s-tag 0 c-tag *	1/1/1:0.*
qinq s-tag * c-tag *	1/1/1:*.*
qinq s-tag * c-tag null	1/1/1:*.null

Considerations:

- The ranges can be modified on the fly for qtag, s-tag/c-tag, or vc-id.
- For port-based vESs, PXC sub-ports are supported. For more information about PXC, see the "Port Cross-Connect (PCX)" chapter in the Interface Configuration volume in the *7450 ESS, 7750 SR, and 7950 XRS MD-CLI Advanced Configuration Guide - Part I*.
- Virtual ESs are supported in EVPN-MPLS, PBB-EVPN, and EVPN-VPWS
- Virtual ESs are supported in single-active and all-active EVPN multi-homing
  - Two all-active vESs must use different ES-BMAC addresses, even if they are defined in the same LAG.
- Virtual ESs implement CMAC flush procedures described in RFC 7623. Optionally, ISID-based CMAC-flush can be used where the single-active vES does not use ES-BMAC allocation. See chapter PBB-EVPN ISID-based CMAC Flush.
- Connection-profile-vlan SAPs (CP-SAPs) cannot be associated with a vES and cannot be configured on ports where vESs are defined. For more information about CP-SAPs, see chapter VLAN Range SAPs for VPLS and Epipe Services.

Configuration

Figure 278: Example topology shows the example topology with four core PEs in an EVPN-MPLS network and two MTUs. VPLS 1 is configured in all the nodes. EVPN is configured on the core PEs, not on the MTUs. LAG 1 is configured on MTU-1, PE-2, and PE-3 and associated with an all-active vES "ESI-23_1" on PE-2 and PE-3. A single-active vES "ESI-45_1" is configured on PE-4 and PE-5, associated with SDPs.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 278: Example topology

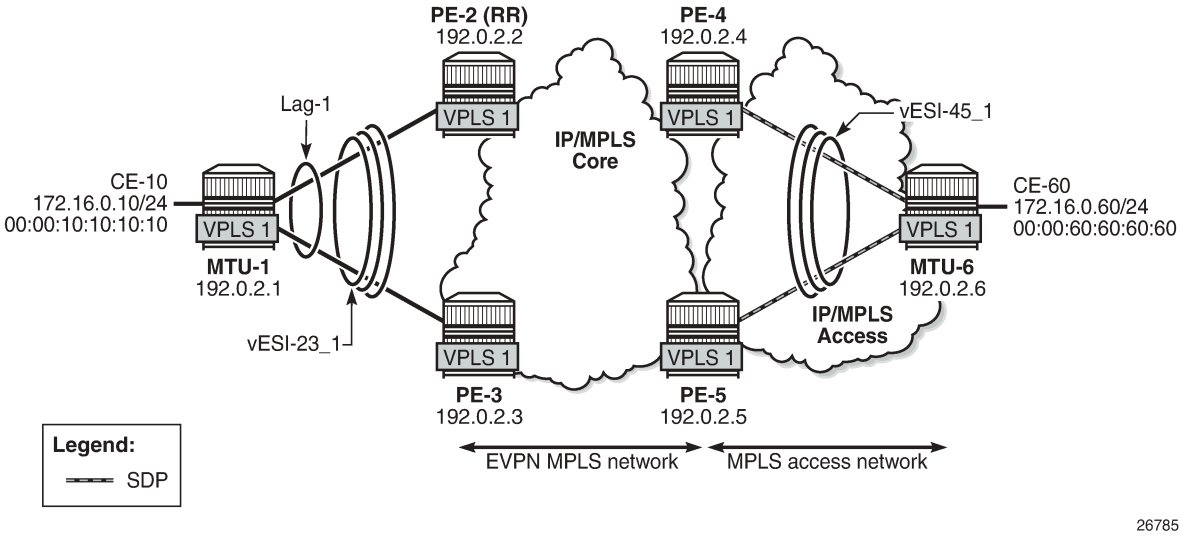

The configuration is similar to the one in chapter EVPN for MPLS Tunnels, where the parameters are described in detail.

The initial configuration on the nodes includes the following:

- Cards, MDAs, ports
- Router interfaces
- IS-IS (alternatively, OSPF can be configured)
- LDP in the IP/MPLS core and IP/MPLS access network

LAG 1 is configured with qinq encapsulation. The LAG configuration on MTU-1 is as follows:

```
# on MTU-1:
configure {
    lag "lag-1" {
        admin-state enable
        encap-type qinq
        mode access
        max-ports 64
        lacp {
            mode active
            administrative-key 32768
        }
        port 1/1/1 {
        }
        port 1/1/2 {
        }
    }
```

BGP is configured on all PEs for address family EVPN. PE-2 is the Route Reflector (RR) and is configured as follows.

```
# on RR PE-2:
configure {
    router "Base" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        autonomous-system 64500
        bgp {
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            peer-ip-tracking true
            split-horizon true
            rapid-update {
                evpn true
            }
            group "internal" {
                peer-as 64500
                family {
                    evpn true
                }
                cluster {
                    cluster-id 1.1.1.1
                }
            }
            neighbor "192.0.2.3" {
                group "internal"
            }
            neighbor "192.0.2.4" {
                group "internal"
            }
            neighbor "192.0.2.5" {
                group "internal"
            }
        }
```

VPLS 1 is configured on all nodes. On the PEs, BGP-EVPN is enabled for MPLS. The following is configured on PE-2:

```
# on PE-2:
configure {
    service {
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
            }
            bgp-evpn {
                evi 1
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    ecmp 2
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            sap lag-1:1.1 {
            }
        }
```

The configuration on the other PEs is similar, but on PE-4 and PE-5, a spoke-SDP is configured instead of a SAP. The service configuration on PE-4 is as follows:

```
# on PE-4:
configure {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    service {
        sdp 46 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.6
            }
        }
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
            }
            bgp-evpn {
                evi 1
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    ecmp 2
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            spoke-sdp 46:1 {
            }
        }
```

Virtual ESs must be configured with type **virtual**; if not, the following error is raised after an attempt to define virtual ranges:

```
*[ex:/configure service system bgp evpn ethernet-segment "ESI-3" association lag "lag-1"]
A:admin@PE-2# virtual-ranges {
MINOR: MGMT_CORE #2203: configure service system bgp evpn ethernet-segment "ESI-3" association
 lag "lag-1" virtual-ranges - Invalid element - virtual-ranges allowed only on virtual
 ethernet-segments
```

On PE-2 and PE-3, the two following two all-active multi-homing vESs are created, each with a unique ESI:

```
# on PE-2, PE-3:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "vESI-23_1" {
                        admin-state enable
                        type virtual
                        esi 01:00:00:00:00:23:01:00:00:01
                        multi-homing-mode all-active
                        df-election {
                            es-activation-timer 3
                        }
                        association {
                            lag "lag-1" {
                                virtual-ranges {
                                    qinq {
                                        s-tag-c-tag 495 c-tag-start 100 {
                                            c-tag-end 102
                                        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                                    s-tag 1 {
                                        end 1
                                    }
                                    s-tag 500 {
                                        end 501
                                    }
                                }
                            }
                        }
                    }
                ethernet-segment "vESI-23_600" {
                    admin-state enable
                    type virtual
                    esi 01:00:00:00:00:23:06:00:00:01
                    multi-homing-mode all-active
                    df-election {
                        es-activation-timer 3
                        service-carving-mode manual
                        manual {
                            evi 2 {
                                end 2
                            }
                        }
                    }
                    association {
                        lag "lag-1" {
                            virtual-ranges {
                                qinq {
                                    s-tag 600 {
                                        end 602
                                    }
                                }
                            }
                        }
                    }
                }
            }
```

When attempting to configure another vES with the ESI of an existing ES/vES, the following error is raised:

```
*[ex:/configure service system bgp evpn ethernet-segment "vESI-23_610"]
A:admin@PE-2# esi 01:00:00:00:00:23:06:00:00:01

*[ex:/configure service system bgp evpn ethernet-segment "vESI-23_610"]
A:admin@PE-2# commit
MINOR: SVCMGR #1003: configure service system bgp evpn ethernet-segment "vESI-23_610" -
 Inconsistent value - esi 01:00:00:00:00:23:06:00:00:01 in use by ethernet segment vESI-23_600
```

Multiple vESs can be defined on the same LAG. However, the ranges should not overlap. The following error is raised after attempting to configure an additional range in vES "ESI-23_600" that uses S-tag 600 in combination with a range of C-tags. S-tag 600 is already included in the first range: **s-tag 600 end 602**. The error message points out that this range is of a different type: the existing range defines only S-tags, whereas the new range defines a range of C-tags for S-tag 600.

```
*[ex:/configure service system bgp evpn ethernet-segment "vESI-23_600" association lag "lag-1"
 virtual-ranges qinq s-tag-c-tag 600 c-tag-start 100]
A:admin@PE-2# commit
MINOR: SVCMGR #1003: configure service system bgp evpn ethernet-segment "vESI-23_600"
 association lag "lag-1" virtual-ranges qinq s-tag-c-tag 600 c-tag-start 100 - Inconsistent
 value - range overlaps with range of a different type in this ethernet-segment
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

When attempting to define **s-tag 1** in "vESI-23_2", when S-tag 1 is already defined in "vESI-23_1", the following error is raised:

```
*[ex:/configure service system bgp evpn ethernet-segment "vESI-23_600" association lag "lag-1"
 virtual-ranges qinq s-tag 1]
A:admin@PE-2# commit
MINOR: SVCMGR #1003: configure service system bgp evpn ethernet-segment "vESI-23_600"
 association lag "lag-1" virtual-ranges qinq s-tag 1 - Inconsistent value - range overlaps with
 range in ethernet-segment vESI-23_1
```

On PE-4, the following single-active multi-homing vESs are configured. The configuration on PE-5 contains a different SDP.

```
# on PE-4:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "vESI-45_1" {
                        admin-state enable
                        type virtual
                        esi 01:00:00:00:00:45:01:00:00:01
                        multi-homing-mode single-active
                        df-election {
                            es-activation-timer 3
                        }
                        association {
                            sdp 46 {
                                virtual-ranges {
                                    vc-id 1 {
                                        end 1
                                    }
                                    vc-id 500 {
                                        end 501
                                    }
                                }
                            }
                        }
                    }
                    ethernet-segment "vESI-45_2" {
                        admin-state enable
                        type virtual
                        esi 01:00:00:00:00:45:02:00:00:01
                        multi-homing-mode single-active
                        df-election {
                            es-activation-timer 3
                            service-carving-mode manual
                            manual {
                                evi 2 {
                                    end 2
                                }
                            }
                        }
                        association {
                            sdp 46 {
                                virtual-ranges {
                                    vc-id 2 {
                                        end 2
                                    }
                                }
                            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    }
                }
            }
        }
```

The configured ESs and vESs can be retrieved as follows:

```
[/]
A:admin@PE-2# show service system bgp-evpn ethernet-segment

===============================================================================
Service Ethernet Segment
===============================================================================
Name                            ESI                             Admin     Oper
-------------------------------------------------------------------------------
vESI-23_1                       01:00:00:00:00:23:01:00:00:01 Enabled   Up
vESI-23_600                     01:00:00:00:00:23:06:00:00:01 Enabled   Up
-------------------------------------------------------------------------------
Entries found: 2
===============================================================================
```

The following information for the first entry in the list shows that it is a virtual ES.

```
[/]
A:admin@PE-2# show service system bgp-evpn ethernet-segment name "vESI-23_1"

===============================================================================
Service Ethernet Segment
===============================================================================
Name                    : vESI-23_1
Eth Seg Type            : Virtual
Admin State             : Enabled            Oper State         : Up
ESI                     : 01:00:00:00:00:23:01:00:00:01
Multi-homing            : allActive          Oper Multi-homing  : allActive
ES SHG Label            : 524280
Source BMAC LSB         : <none>
Lag                     : lag-1
ES Activation Timer     : 3 secs
Oper Group              : (Not Specified)
Svc Carving             : auto               Oper Svc Carving   : auto
Cfg Range Type          : primary
===============================================================================
```

Virtual ES "vESI-23_1" on PE-2 has the following S-tag ranges and S/C-tag ranges:

```
[/]
A:admin@PE-2# show service system bgp-evpn ethernet-segment name "vESI-23_1" virtual-ranges

===============================================================================
Q-Tag Ranges
===============================================================================
Q-Tag Start         Q-Tag End           Last Changed
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
No entries found
===============================================================================

===============================================================================
VC-Id Ranges
===============================================================================
VC-Id Start         VC-Id End           Last Changed
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
No entries found
===============================================================================

===============================================================================
S-Tag Ranges
===============================================================================
S-Tag Start         S-Tag End           Last Changed
-------------------------------------------------------------------------------
1                   1                   04/20/2021 16:14:55
500                 501                 04/20/2021 16:14:55
-------------------------------------------------------------------------------
Number of Entries: 2
===============================================================================

===============================================================================
S-Tag C-Tag Ranges
===============================================================================
S-Tag Start         C-Tag Start         C-Tag End       Last Changed
-------------------------------------------------------------------------------
495                 100                 102             04/20/2021 16:14:55
-------------------------------------------------------------------------------
Number of Entries: 1
===============================================================================

===============================================================================
Vxlan Instance Service Ranges
===============================================================================
Svc Range Start          Svc Range End            Last Changed
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
No entries found
===============================================================================
```

The ranges in the vES can be modified while the vES is operationally up, for example, an S-tag range can be added as follows:

```
# on PE-2:
configure {
    service {
        system {
            bgp {
                evpn
                    ethernet-segment "vESI-23_1" {
                        association {
                            lag "lag-1" {
                                virtual-ranges {
                                    qinq {
                                        s-tag 10 {
                                            end 10
                                        }
                                    }
                                }
```

The S-tag ranges can be verified with the following command. Compared with the preceding output, the S-tag 10 has been added:

```
[/]
A:admin@PE-2# show service system bgp-evpn ethernet-segment name "vESI-23_1" virtual-ranges |
 match S-Tag post-lines 8
S-Tag Ranges
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
S-Tag Start         S-Tag End           Last Changed
-------------------------------------------------------------------------------
1                   1                   04/20/2021 16:14:55
10                  10                  04/20/2021 16:17:23
500                 501                 04/20/2021 16:14:55
-------------------------------------------------------------------------------
Number of Entries: 3
===============================================================================
===============================================================================
S-Tag C-Tag Ranges
===============================================================================
S-Tag Start         C-Tag Start         C-Tag End       Last Changed
-------------------------------------------------------------------------------
495                 100                 102             04/20/2021 16:14:55
-------------------------------------------------------------------------------
Number of Entries: 1
===============================================================================
===============================================================================
Vxlan Instance Service Ranges
===============================================================================
```

On PE-4, the same **show** command shows the range of VC-IDs, as follows:

```
[/]
A:admin@PE-4# show service system bgp-evpn ethernet-segment name "vESI-45_1" virtual-ranges

===============================================================================
Q-Tag Ranges
===============================================================================
Q-Tag Start         Q-Tag End           Last Changed
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
No entries found
===============================================================================

===============================================================================
VC-Id Ranges
===============================================================================
VC-Id Start         VC-Id End           Last Changed
-------------------------------------------------------------------------------
1                   1                   04/20/2021 16:15:58
500                 501                 04/20/2021 16:15:58
-------------------------------------------------------------------------------
Number of Entries: 2
===============================================================================

===============================================================================
S-Tag Ranges
===============================================================================
S-Tag Start         S-Tag End           Last Changed
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
No entries found
===============================================================================

===============================================================================
S-Tag C-Tag Ranges
===============================================================================
S-Tag Start         C-Tag Start         C-Tag End       Last Changed
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
No entries found
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Vxlan Instance Service Ranges
===============================================================================
Svc Range Start          Svc Range End            Last Changed
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
No entries found
===============================================================================
```

Connection-profile-vlan SAPs (CP-SAPs) cannot be associated with a vES and cannot be configured on ports where vESs are defined. CP-SAP 10 is created on PE-3, as follows:

```
# on PE-3:
configure {
    connection-profile vlan 10 {
        qtag-range 5 {
            end 100
        }
        qtag-range 495 {
            end 495
        }
```

The following vES is configured on PE-3:

```
# on PE-3:
configure {
    service {
        system {
            bgp {
                evpn {
                    ethernet-segment "vESI-23_10" {
                        admin-state enable
                        type virtual
                        esi 01:00:00:00:00:23:10:00:00:01
                        multi-homing-mode single-active
                        df-election {
                            es-activation-timer 3
                        }
                        association {
                            port 1/2/3 {
                                virtual-ranges {
                                    qinq {
                                        s-tag 100 {
                                            end 100
                                        }
                                    }
                                }
                            }
                        }
                    }
```

This vES can only be configured when no CP-SAPs are defined on port 1/2/3. The following error message is raised when a CP-SAP is configured on port 1/2/3 already and the vES is configured afterward:

```
*[ex:/configure service system bgp evpn ethernet-segment "vESI-23_10" association port 1/2/3
 virtual-ranges qinq s-tag 100]
A:admin@PE-3# commit
MINOR: MGMT_CORE #4001: configure service vpls "VPLS 1" sap 1/2/3:100.cp-10 - connection
 profile saps not allowed on port/lags associated with evpn ethernet-segments - configure
 service system bgp evpn ethernet-segment "vESI-23_10" association
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

When attempting to configure CP-SAP 1/2/3:cp-10 in VPLS 1 with port 1/2/3 associated with a vES, the following error message is raised.

```
*[ex:/configure service vpls "VPLS 1" sap 1/2/3:100.cp-10]
A:admin@PE-3# commit
MINOR: MGMT_CORE #4001: configure service vpls "VPLS 1" sap 1/2/3:100.cp-10 - connection
 profile saps not allowed on port/lags associated with evpn ethernet-segments - configure
 service system bgp evpn ethernet-segment "vESI-23_10" association
```

Conclusion

Regular ESs and vESs can be associated with ports, LAGs, and SDPs; in case of vES, ranges of Q-tags, S-tags, C-tags, or VC-IDs can be defined. The granularity for vES is per AC. Multiple vESs with different ESIs can be defined on the same port, LAG, or SDP.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

VLAN Range SAPs for VPLS and Epipe Services

This chapter provides information about VLAN range SAPs for VPLS and Epipe services.

Topics in this chapter include:

Applicability

This chapter was initially written for SR OS Release 14.0.R6, but the MD-CLI in the current edition is based on SR OS Release 21.2.R1. Connection-Profile VLAN SAPs (CP SAPs) are supported in SR OS Release 14.0.R1, and later.

Overview

Backhaul services through metro Ethernet networks require bundled interface support. In SR OS terminology, bundling refers to Connection-Profile VLAN SAPs (CP SAPs)—special SAPs that capture the traffic of a range of CE VLAN IDs (VIDs) entering an Ethernet port. CP SAPs are fully compatible with Metro Ethernet Forum (MEF) 10.3 bundling service attributes and RFC 7432 EVPN VLAN bundle service interfaces. CP SAPs are supported in Layer 2 services only, and can be configured together with other SAPs and/or SDP-bindings.

For frames with an ingress VID contained in the range configured in the SAP's CP, the behavior is similar to default SAPs, such as 1/1/1:*, where "*" spans the entire VID range from 0 to 4095 and serves as a wildcard. However, unlike a default SAP, a CP SAP cannot co-exist with a VLAN SAP that is in the same range and on the same port or LAG. For example, 1/1/1:* and 1/1/1:100 can co-exist whereas 1/1/2:cp-1 (where cp-1 corresponds to the VLAN range from 1 to 200) and 1/1/2:100 cannot co-exist.

The VLAN manipulation between VLAN SAPs, default SAPs, and CP SAPs is compared in Table 16: VLAN manipulation in SAPs .

Table 16: VLAN manipulation in SAPs

	VLAN SAP	Default SAP	CP SAP
Service-delimiting VLAN	Yes For example: VLAN 100 in 1/1/1:100	No	No
Push/pop VLAN tags in egress/ ingress frames	Yes	No	No
VLAN translation	Yes	No	No

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 279: Customer VID is popped and pushed by VLAN SAPs - VLAN translation shows how dot1q VLAN SAPs pop the customer VLAN tag in ingress frames and push the VLAN tag in egress frames. Therefore, frames are untagged between PE-1 and PE-2. VLAN translation is possible when the VIDs in the VLAN tags that are popped or pushed at the SAPs are different at ingress and egress, as follows.

Figure 279: Customer VID is popped and pushed by VLAN SAPs - VLAN translation

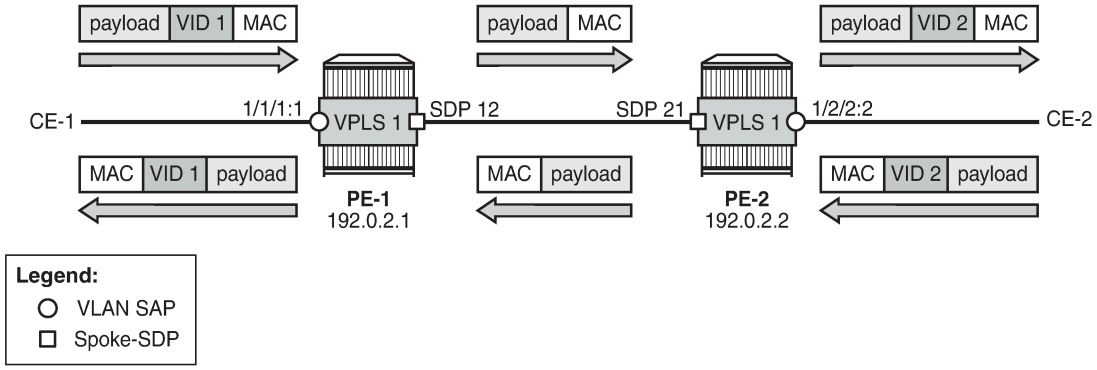

Figure 280: Customer VID is preserved between dot1q CP SAPs - no VLAN translation shows that dot1q CP SAPs do not pop or push the CE VID. Frames keep the same tag end-to-end; therefore, VLAN translation is not possible.

Figure 280: Customer VID is preserved between dot1q CP SAPs - no VLAN translation

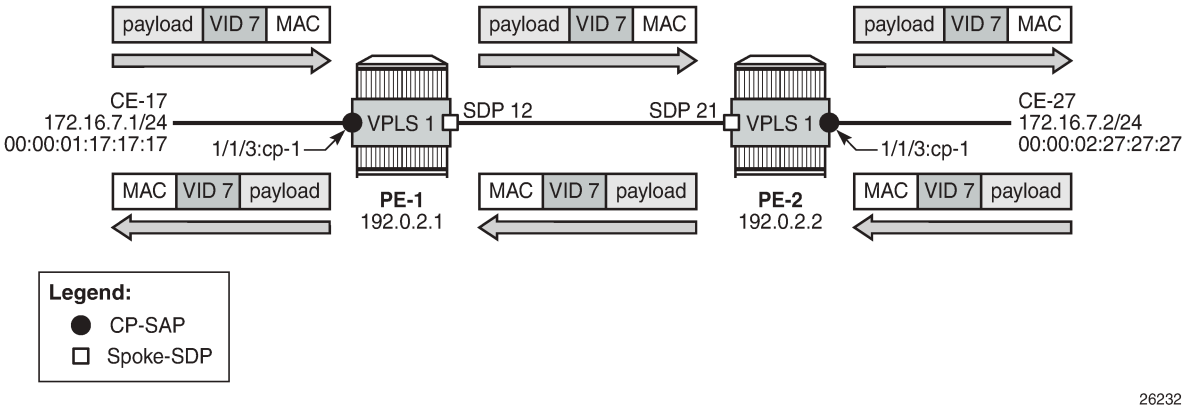

Figure 281: Customer VID is preserved between QinQ CP SAPs - no VLAN translation shows that QinQ CP SAPs only pop or push the service delimiting VID (VID 100), but not the customer VID in the CP range, as follows:

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 281: Customer VID is preserved between QinQ CP SAPs - no VLAN translation

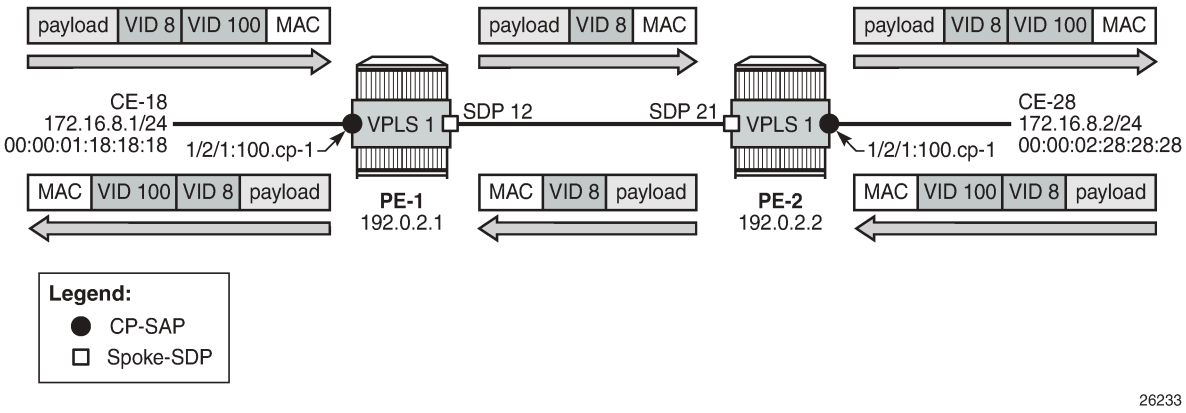

VID 100 is service delimiting and can be different in both SAPs, but the customer VID in the VLAN range of the CP is not.

Connection profile VLAN

Note:
The **connection-profile>vlan** context is different from the connection-profile used for ATM connectivity.

CP SAPs refer to connection profiles that can contain up to 32 ranges of customer VIDs. Connection profiles are configured with the following command:

```
[ex:/configure connection-profile]
A:admin@PE-1# vlan ?

 [connection-profile-id] <number>
 <number>  - <1..8000>

    Identifier of this connection profile
```

VLAN ranges in a CP contain one or more consecutive VIDs, as follows:

```
*[ex:/configure connection-profile vlan 10]
A:admin@PE-1# qtag-range ?

 [start] <number>
 <number>  - <1..4094>

    Lower bound of VLAN range for connection profile
```

```
*[ex:/configure connection-profile vlan 10]
A:admin@PE-1# qtag-range 150 ?

 qtag-range

 Immutable fields      - end

 apply-groups          - Apply a configuration group at this level
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
 apply-groups-exclude  - Exclude a configuration group at this level
 end                   - Upper bound of VLAN range for connection profile
```

Following is an example of a CP configuration containing three non-overlapping VLAN ranges:

```
configure {
    connection-profile {
        vlan 10 {
            qtag-range 5 {
                end 100
            }
            qtag-range 150 {
                end 300
            }
            qtag-range 350 {
                end 350
            }
        }
```

Overlapping ranges are not allowed within the same CP. The following error is raised when attempting to add a VLAN range from 7 to 9 to the preceding CP.

```
*[ex:/configure connection-profile vlan 10 qtag-range 7]
A:admin@PE-1# commit
MINOR: SVCMGR #9012: configure connection-profile vlan 10 - Overlapping range - configure
 connection-profile vlan 10 qtag-range 7 end
```

Additional VLAN ranges can be configured to the CP defined in an existing and operationally up SAP. The CP's VLAN ranges can also be removed on the fly. When a user wants to extend a VLAN range, for example, VLAN range 350 becoming a range from 350 to 400, the existing VLAN range is overwritten, as follows:

```
[ex:configure connection-profile vlan 10]
A:admin@PE-1# qtag-range 350 {

[ex:configure connection-profile vlan 10 qtag-range 350]
A:admin@PE-1# end 400

*[ex:configure connection-profile vlan 10 qtag-range 350]
A:admin@PE-1# }

*[ex:configure connection-profile vlan 10]
A:admin@PE-1# commit

[ex:configure connection-profile vlan 10]
A:admin@PE-1# info
    qtag-range 5 {
        end 100
    }
    qtag-range 150 {
        end 300
    }
    qtag-range 350 {
        end 400
    }
```

The following example shows three VLAN ranges in CP 10, with a timestamp of the last change for each VLAN range:

```
[/]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
A:admin@PE-1# show connection-profile-vlan 10

===============================================================================
Connection Profile 10 Information
===============================================================================
Description : (Not Specified)
Last Change : 03/31/2021 09:02:03


===============================================================================
Connection Profile Vlan Eth Information
===============================================================================
Range Start         Range End                     Last Change
-------------------------------------------------------------------------------
5                   100                           03/31/2021 09:11:59
150                 300                           03/31/2021 09:11:59
350                 400                           03/31/2021 09:12:08
===============================================================================
===============================================================================
```

If a VLAN tag combination matches different SAPs, the highest priority SAP will be picked regardless of the operational status. For completeness, the following two tables show the SAP lookup matching order for dot1q and QinQ ports.

Table 17: SAP lookup order for dot1q ports

Incoming frame qtag VID value	**SAP lookup precedence order (:0 and :* are mutually exclusive on the same port)**			
	:X	**:CP**	**:0**	**:***
X (belongs to the CP range)	1st	1st		2nd
0			1st	1st
<untagged>			1st	1st

Table 18: SAP lookup order for QinQ ports

Incoming frame qtag1.qtag2	**System/port settings = new-qinq-untagged-sap SAP lookup precedence order (assumption: X and Y are defined in CP ranges)**							
	:X.Y	**:X.0**	**:X.CP**	**:CP.***	**:X.***	**:0.***	**:*.null**	**:*.***
X.Y	1st		1st	2nd	2nd			3rd
X.0		1st		2nd	2nd			3rd
0.Y						1st		2nd
0.0						1st		2nd
X		1st		2nd	2nd		3rd	4th
0						1st	2nd	3rd

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Incoming frame qtag1.qtag2	System/port settings = new-qinq-untagged-sap SAP lookup precedence order (assumption: X and Y are defined in CP ranges)							
	:X.Y	:X.0	:X.CP	:CP.*	:X.*	:0.*	:*.null	:*.*
<untagged>						1st	2nd	3rd

For example, ingress frames with VIDs 100.20 are classified as part of CP SAP 1/2/1:100.cp-10, not of CP SAP 1/2/3:cp-10.*. Only when SAP 1/2/1:100.cp-10 is removed from the configuration, frames with VIDs 100.20 will go to SAP 1/2/3:cp-10.*.

Assign CP SAPs to VPLS or Epipe services

Like ordinary SAPs, CP SAPs can be assigned to VPLS or Epipe services, as follows. The VPLS and Epipe can be EVPN services or not. In the following example, VPLS 1 has BGP-EVPN enabled, whereas Epipe 2 does not:

```
# on PE-1:
configure {
    service {
        epipe "Epipe 2" {
            admin-state enable
            service-id 2
            customer "1"
            spoke-sdp 12:2 {
            }
            sap 1/2/1:200.cp-10 {
            }
        }
        sdp 12 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.2
            }
        }
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
            }
            bgp-evpn {
                evi 1
                mpls 1 {
                    admin-state enable
                    ingress-replication-bum-label true
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            sap 1/1/3:cp-10 {
            }
            sap 1/2/1:1.11 {
            }
            sap 1/2/1:100.cp-10 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
            sap 1/2/3:cp-10.* {
            }
        }
```

CP SAPs are configured in the same way as VLAN SAPs and default SAPs, with the following restrictions:

- A CP can be defined for inner or outer tags as shown in the preceding configuration, but not both at the same time, as follows:

```
*[ex:/configure service vpls "VPLS 1" sap 1/2/1:cp-3.cp-10]
A:admin@PE-1# commit
MINOR: MGMT_CORE #4001: configure service vpls "VPLS 1" sap 1/2/1:cp-3.cp-10 - SAP and port
 encapsulation values are incompatible - configure port 1/2/1 ethernet encap-type
```

- If a CP is defined for the outer VID, the inner VID cannot be a specific VID, as follows. The inner VID can only be a "*" (where the inner tag can have any value) or a "0" (where the inner tag can be 0 or null).

```
*[ex:/configure service vpls "VPLS 1" sap 1/2/1:cp-3.4]
A:admin@PE-1# commit
MINOR: MGMT_CORE #4001: configure service vpls "VPLS 1" sap 1/2/1:cp-3.4 - SAP and port
 encapsulation values are incompatible - configure port 1/2/1 ethernet encap-type
```

- No VLAN SAP can be added on a port in dot1q (or a combination of port and service-delimiting VLAN in case of QinQ) when the VLAN is included in the VLAN range in a CP SAP on the same port. One of the VLAN ranges in CP 10 contains all VIDs from 5 to 100. Therefore, it is not allowed to configure a VLAN SAP with VID 100 on port 1/1/3, where a CP SAP is configured with CP 10, as follows:

```
*[ex:/configure service vpls "VPLS 1" sap 1/1/3:100]
A:admin@PE-1# commit
MINOR: COMMON #238: configure service vpls "VPLS 1" sap 1/1/3:100 - Configuration change
 failed validation - sap conflicts with connection-profile-vlan 10
```

- No CP SAPs can be added with overlapping VLAN ranges on the same port for dot1q (or on the same port- and service-delimiting tag for QinQ), as follows. CP 1 contains VLAN range from 7 to 9, which overlaps with VLAN range from 5 to 100 in CP 10.

```
# on PE-1:
configure {
    connection-profile {
        vlan 1 {
            qtag-range 7 {
                end 9
            }
```

```
*[ex:/configure service vpls "VPLS 1" sap 1/1/3:cp-1]
A:admin@PE-1# commit
MINOR: COMMON #238: configure service vpls "VPLS 1" sap 1/1/3:cp-1 - Configuration change
 failed validation - a sap 1/1/3:7 in the connection-profile-vlan conflicts with connect-
profile-vlan 10
MINOR: COMMON #238: configure service vpls "VPLS 1" sap 1/1/3:cp-1 - Configuration change
 failed validation - a sap 1/1/3:8 in the connection-profile-vlan conflicts with connect-
profile-vlan 10
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
MINOR: COMMON #238: configure service vpls "VPLS 1" sap 1/1/3:cp-1 - Configuration change
 failed validation - a sap 1/1/3:9 in the connection-profile-vlan conflicts with connect-
profile-vlan 10
```

```
*[ex:/configure service vpls "VPLS 1" sap 1/2/1:100.cp-1]
A:admin@PE-1# commit
MINOR: COMMON #238: configure service vpls "VPLS 1" sap 1/2/1:100.cp-1 - Configuration
 change failed validation - a sap 1/2/1:100.7 in the connection-profile-vlan conflicts with
 connect-profile-vlan 10
MINOR: COMMON #238: configure service vpls "VPLS 1" sap 1/2/1:100.cp-1 - Configuration
 change failed validation - a sap 1/2/1:100.8 in the connection-profile-vlan conflicts with
 connect-profile-vlan 10
MINOR: COMMON #238: configure service vpls "VPLS 1" sap 1/2/1:100.cp-1 - Configuration
 change failed validation - a sap 1/2/1:100.9 in the connection-profile-vlan conflicts with
 connect-profile-vlan 10
```

However, the CP can be referred to by SAPs on other ports for dot1q or for QinQ on other combinations of port and service-delimiting VLAN.

- CP SAPs can be added when they contain non-overlapping VLAN ranges on the same port, as follows. CP 3 contains one VLAN range with only one VID: 3. This VLAN range (3) does not overlap with any VLAN range in the CP SAPs assigned to VPLS 1.

```
# on PE-1:
configure {
    connection-profile {
        vlan 3 {
            qtag-range 3 {
                end 3
            }
        }
    }
    service {
        vpls "VPLS 1" {
            sap 1/1/3:cp-3 {
            }
            sap 1/2/1:100.cp-3 {
            }
```

VPLS 1 contains the following SAPs. There is no overlap between the VLAN ranges on a port (or port and service-delimiting tag for QinQ).

```
[/]
A:admin@PE-1# show service id 1 sap

===============================================================================
SAP(Summary), Service 1
===============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.  Egr.   Adm  Opr
                                           QoS   Fltr    QoS   Fltr
-------------------------------------------------------------------------------
1/1/3:cp-3                      1          1     none    1     none   Up   Up
1/1/3:cp-10                     1          1     none    1     none   Up   Up
1/2/1:cp-3.0                    1          1     none    1     none   Up   Up
1/2/1:1.11                      1          1     none    1     none   Up   Up
1/2/1:cp-10.*                   1          1     none    1     none   Up   Up
1/2/1:101.cp-1                  1          1     none    1     none   Up   Up
1/2/1:100.cp-3                  1          1     none    1     none   Up   Up
1/2/1:100.cp-10                 1          1     none    1     none   Up   Up
1/2/3:cp-10.*                   1          1     none    1     none   Up   Up
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Number of SAPs : 9
-------------------------------------------------------------------------------
===============================================================================
```

Constraints to be considered when applying CP SAPs in Layer 2 services are described in the Release Notes, section "Known Limitations" - "Services General".

Consumed resources for CP SAPs

The following SAPs are used on PE-1: nine SAPs are used in VPLS 1 and one SAP is used in Epipe 2:

```
[/]
A:admin@PE-1# show service sap-using

===============================================================================
Service Access Points
===============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.  Egr.   Adm  Opr
                                           QoS   Fltr    QoS   Fltr
-------------------------------------------------------------------------------
1/1/3:cp-3                      1          1     none    1     none   Up   Up
1/1/3:cp-10                     1          1     none    1     none   Up   Up
1/2/1:cp-3.0                    1          1     none    1     none   Up   Up
1/2/1:1.11                      1          1     none    1     none   Up   Up
1/2/1:cp-10.*                   1          1     none    1     none   Up   Up
1/2/1:101.cp-1                  1          1     none    1     none   Up   Up
1/2/1:100.cp-3                  1          1     none    1     none   Up   Up
1/2/1:100.cp-10                 1          1     none    1     none   Up   Up
1/2/3:cp-10.*                   1          1     none    1     none   Up   Up
1/2/1:200.cp-10                 2          1     none    1     none   Up   Up
-------------------------------------------------------------------------------
Number of SAPs : 10
-------------------------------------------------------------------------------
===============================================================================
```

Regular and default SAPs consume one SAP instance each, whereas CP SAPs consume a number of SAP instances equal to the number of VLANs in the range. The following shows that there are ten SAP entries (in this example, nine SAPs in VPLS 1 and one SAP in Epipe 2), which can be regular, default, or CP SAP entries:

```
[/]
A:admin@PE-1# tools dump resource-usage system

===============================================================================
Resource Usage Information for System
===============================================================================
                                                 Total   Allocated        Free
-------------------------------------------------------------------------------
                SAP Ingress QoS Policies |        3071           1        3070
                 SAP Egress QoS Policies |        3071           1        3070
            Ingress Queue-Group Templates |        2047           4        2043
             Egress Queue-Group Templates |        2047           5        2042
        Egress Port Queue-Group Instances |      163839           8      163831
         Ingress FP Queue-Group Instances |       16383           0       16383
              Fast Depth Monitored Queues |       50000           0       50000
                        Egress Port VPort |       40959           0       40959
       Dynamic Services Next-Hop Entries +      511999           0      511999
                  IPSec Next-Hop Entries -      500000           0      500000
            Subscriber Next-Hop Entries -      500000           0      500000
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                            SAP Entries +      262143          10      262133
                     (in use by:   Apipe) -                     0
                     (in use by:   Cpipe) -                     0
                     (in use by:   Epipe) -                     1
                     (in use by:   Fpipe) -                     0
                     (in use by:   Ipipe) -                     0
                     (in use by:     Ies) -                     0
                     (in use by:  Mirror) -                     0
                     (in use by:    Vpls) -                     9
                     (in use by:    Vprn) -                     0
===============================================================================
```

However, the number of SAP instances consumed for card 1 FP 1 exceeds the number of SAP entries in the system, as follows:

```
[/]
A:admin@PE-1# tools dump resource-usage card 1 fp 1

===============================================================================
Resource Usage Information for Card Slot #1 FP #1
===============================================================================
                                                 Total   Allocated        Free
-------------------------------------------------------------------------------
---snip---
                              SAP Instances |      63999       1497      62502
---snip---
===============================================================================
```

The calculation of the number of SAP instances is as follows. In this example, CP 10 is used in five SAPs (four in VPLS 1 and one in Epipe 2) and contains the following VLAN ranges:

```
[/]
A:admin@PE-1# show connection-profile-vlan 10

===============================================================================
Connection Profile 10 Information
===============================================================================
Description : (Not Specified)
Last Change : 03/31/2021 09:02:03


===============================================================================
Connection Profile Vlan Eth Information
===============================================================================
Range Start         Range End                     Last Change
-------------------------------------------------------------------------------
5                   100                           03/31/2021 09:11:59
150                 300                           03/31/2021 09:11:59
350                 400                           03/31/2021 09:12:08
===============================================================================
===============================================================================
```

The number of VLANs in the VLAN ranges of CP 10 equals 298. For each of the five SAP entries with CP 10, 298 SAP instances are used, for a total of 1490. As well, there is one CP SAP using CP 1 with three VLANs in the VLAN range from 7 to 9 (for three more SAP instances). Three CP SAPs use CP 3 with only VID 3 in the VLAN range (for three more SAP instances), and one SAP is a regular SAP that consumes one SAP instance. Therefore, the total number of SAP instances is 1497.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Configuration

Figure 282: Example topology shows the example topology used in this chapter.

Figure 282: Example topology

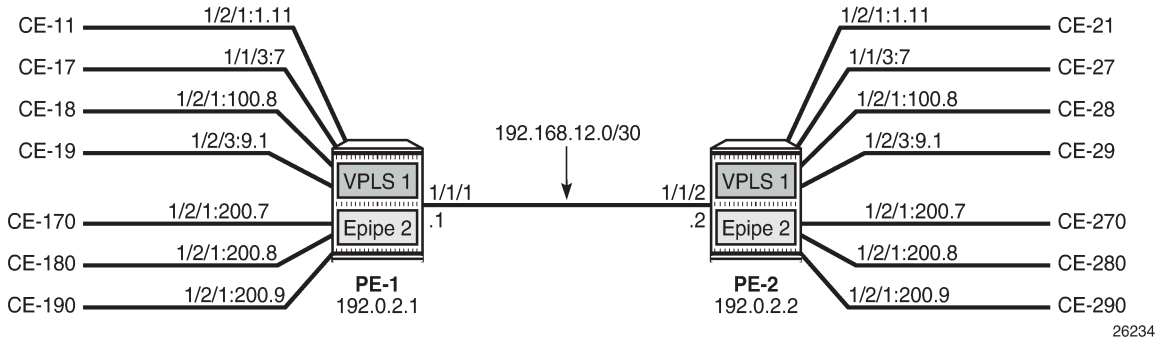

The initial configuration on the PEs includes the following:

- Cards, MDAs, ports
- Router interfaces
- IS-IS (or OSPF) between the PEs
- LDP between the PEs

In this example, no BGP is configured and no BGP-EVPN will be configured in the VPLS and Epipe services. However, VLAN ranges can be applied in EVPN VPLS and EVPN Epipe services.

VLAN ranges in VPLS services

Figure 283: Example topology for VLAN ranges in VPLS 1 shows the example topology for VPLS 1 with a combination on VLAN SAPs and CP SAPs. The port:VID represents the port to which the CE is connected and the VID sent by the CE; for example, CE-17 is connected to port 1/1/3 on PE-1 and sends frames with VID 7. When VLAN ranges are used, the port:VID 1/1/3:7 does not represent the configured SAP, which is 1/1/3:cp-1.

Figure 283: Example topology for VLAN ranges in VPLS 1

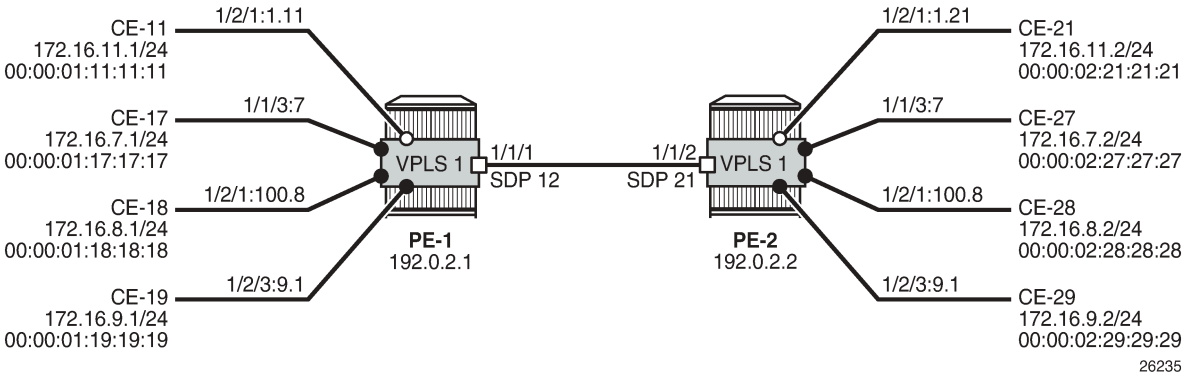

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The service configuration for VPLS 1 on PE-1 is as follows:

```
# on PE-1:
configure {
    service {
        sdp 12 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.2
            }
        }
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            spoke-sdp 12:1 {
            }
            sap 1/1/3:cp-1 {
            }
            sap 1/2/1:1.11 {
            }
            sap 1/2/1:100.cp-1 {
            }
            sap 1/2/3:cp-1.* {
            }
        }
```

The configuration of VPLS 1 on PE-2 is as follows:

```
‘# on PE-2:
configure {
    service {
        sdp 21 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.1
            }
        }
        vpls "VPLS 1" {
            admin-state enable
            service-id 1
            customer "1"
            spoke-sdp 21:1 {
            }
            sap 1/1/3:cp-1 {
            }
            sap 1/2/1:1.21 {
            }
            sap 1/2/1:100.cp-1 {
            }
            sap 1/2/3:cp-1.* {
            }
        }
```

When the CEs send traffic to each other, such as ICMP echo requests, the MAC addresses are learned in the SAPs, and the forwarding database (FDB) on PE-1 is as follows:

```
[/]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
A:admin@PE-1# show service id 1 fdb detail

===============================================================================
Forwarding Database, Service 1
===============================================================================
ServId     MAC               Source-Identifier       Type     Last Change
            Transport:Tnl-Id                         Age
-------------------------------------------------------------------------------
1          00:00:01:11:11:11 sap:1/2/1:1.11          L/90     03/31/21 10:31:01
1          00:00:01:17:17:17 sap:1/1/3:cp-1          L/90     03/31/21 10:26:44
1          00:00:01:18:18:18 sap:1/2/1:100.cp-1      L/90     03/31/21 10:26:44
1          00:00:01:19:19:19 sap:1/2/3:cp-1.*        L/90     03/31/21 10:26:44
1          00:00:02:21:21:21 sdp:12:1                L/90     03/31/21 10:31:01
1          00:00:02:27:27:27 sdp:12:1                L/90     03/31/21 10:26:44
1          00:00:02:28:28:28 sdp:12:1                L/90     03/31/21 10:26:44
1          00:00:02:29:29:29 sdp:12:1                L/90     03/31/21 10:26:44
-------------------------------------------------------------------------------
No. of MAC Entries: 8
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
===============================================================================
```

VLAN manipulation in dot1q SAPs

Figure 284: Customer VIDs are popped and pushed by dot1q VLAN SAPs shows the VLAN manipulation for VLAN SAPs. CE-17 and CE-18 are connected to VLAN SAPs, where the VLAN tag with VID 7 will be popped or pushed. VLAN translation is possible, but does not apply. The configuration of the SAPs in VPLS 1 on PE-1 and PE-2 is modified as follows:

```
# on PE-1, PE-2:
configure {
    service {
        vpls "VPLS 1" {
            delete sap 1/1/3:cp-1
            sap 1/1/3:7 {
            }
```

Figure 284: Customer VIDs are popped and pushed by dot1q VLAN SAPs

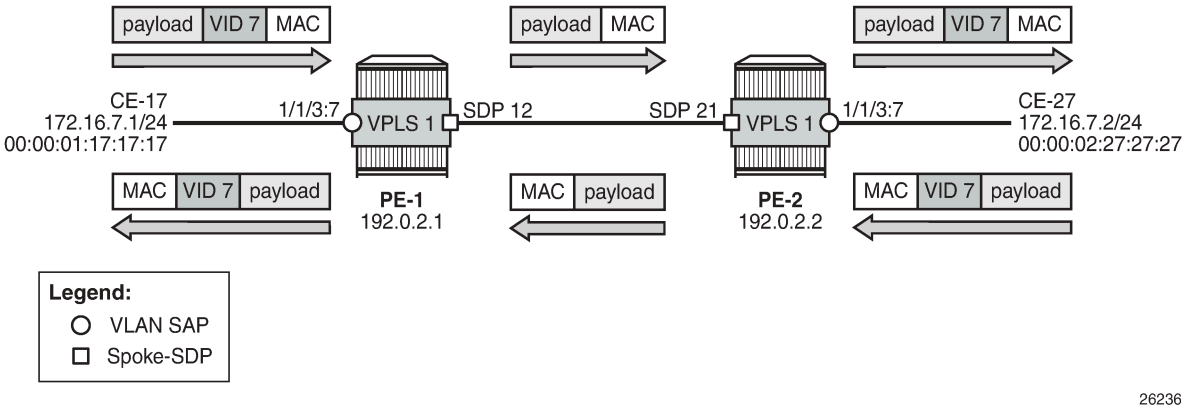

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 285: Customer VID is preserved between two dot1q CP SAPs shows how the customer VID 7 is preserved between CE-17 and CE-27 when CP SAPs are used instead of VLAN SAPs. The configuration for the SAPs is modified as follows:

```
# on PE-1, PE-2:
configure {
    service {
        vpls "VPLS 1" {
            delete sap 1/1/3:7
            sap 1/1/3:cp-1 {
            }
```

CE-17 sends frames with VID 7 to dot1q CP SAP 1/1/3:cp-1 in VPLS 1 on PE-1, and this CP SAP preserves the VLAN tag. When the frames with VID 7 reach the egress CP SAP 1/1/3:cp-1 of VPLS 1 on PE-2, the egress CP SAP preserves the VID, and the frames are forwarded to CE-27. Traffic in the opposite direction is treated in the same way: the customer VID is preserved between the CEs.

Figure 285: Customer VID is preserved between two dot1q CP SAPs

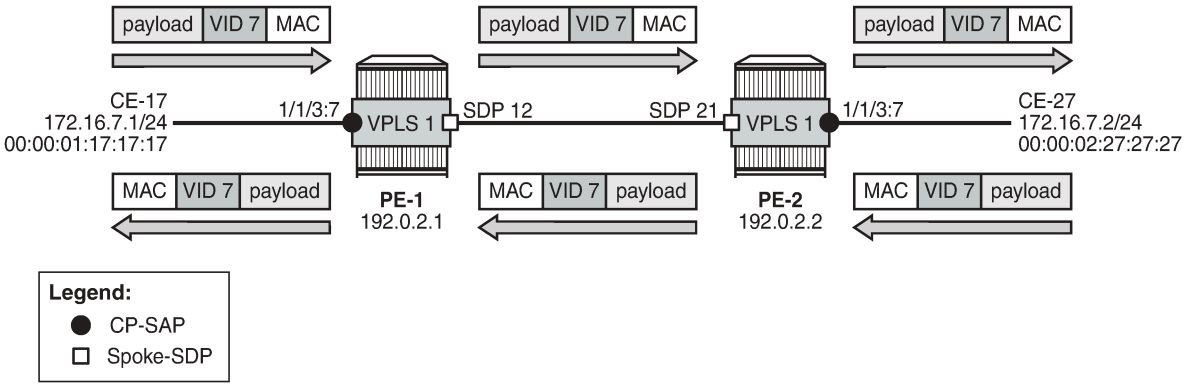

No traffic is possible between a CP SAP in VPLS 1 on PE-1 and a VLAN SAP in VPLS 1 on PE-2, as shown in Figure 286: No traffic between dot1q CP SAP and dot1q VLAN SAP.

Figure 286: No traffic between dot1q CP SAP and dot1q VLAN SAP

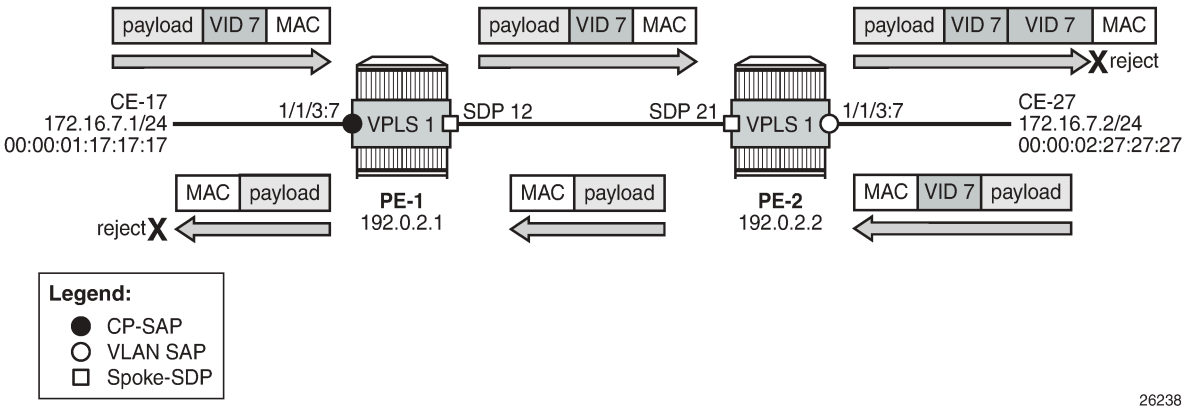

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The CP SAP 1/1/3:cp-1 in VPLS 1 on PE-1 remains unchanged, whereas the SAP in VPLS 1 on PE-2 is reconfigured as VLAN SAP 1/1/3:7 for VLAN 7, as follows:

```
# on PE-2:
configure {
    service {
        vpls "VPLS 1" {
            delete sap 1/1/3:cp-1
            sap 1/1/3:7 {
            }
```

Frames from CE-17 are forwarded by CP SAP 1/1/3:cp-1 in VPLS 1 on PE-1 without any changes to the VLAN tag. The tagged frames reach the VLAN SAP 1/1/3:7, where another VLAN tag with VID 7 is pushed onto the frame. The receiver CE-27 rejects the double-tagged frame. When CE-27 sends traffic to CE-17, the VLAN SAP 1/1/3:7 in VPLS 1 on PE-2 pops the VLAN tag and the frame is forwarded untagged to PE-1. The CP SAP 1/1/3:cp-1 on PE-1 does not push any VLAN tag and the frame is forwarded untagged to CE-17, where it is rejected.

VLAN manipulation in QinQ SAPs

Figure 287: Traffic between two QinQ VLAN SAPs - VLAN translation shows the VLAN manipulation in QinQ VLAN SAPs that pop and push the VLAN labels. In the example, the customer VID is translated.

Figure 287: Traffic between two QinQ VLAN SAPs - VLAN translation

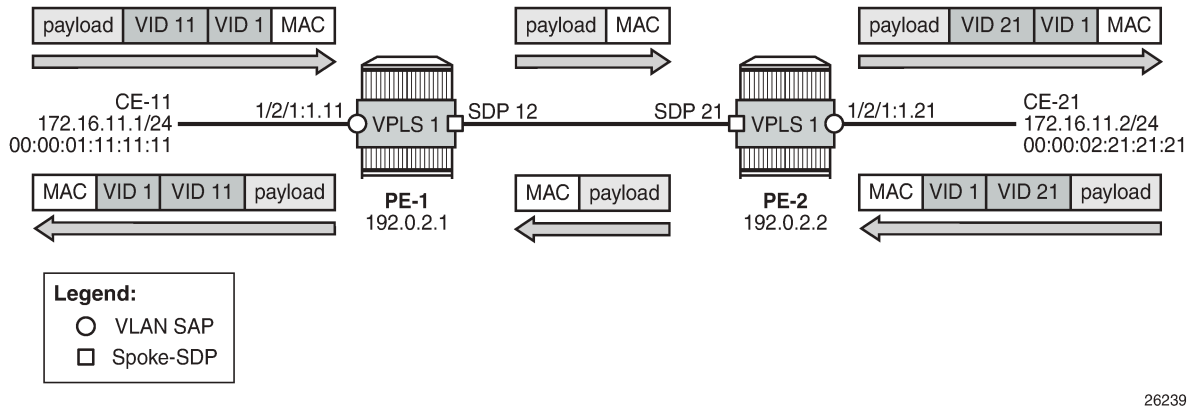

CE-11 sends double-tagged traffic to QinQ VLAN SAP 1/2/1:1.11 in VPLS 1 on PE-1. This VLAN SAP pops both labels and forwards the frame untagged to PE-2. The egress VLAN SAP 1/2/1:1.21 in VPLS 1 on PE-2 pushes a label stack with two labels: the inner label with VID 21 and the outer label with VID 1. Both VIDs can be translated, but in this example, only the inner label gets another VID.

Figure 288: No traffic between two QinQ CP SAPs - VLAN translation not supported shows that VLAN translation is not possible between two QinQ CP SAPs. In the example, the outer tag with VID 1 is popped by the CP SAPs (VLAN translation is possible for this VLAN tag, but not done here) and the inner tag with VID 11 or 21 is preserved by the CP SAPs, which implies that the received frames will be rejected.

In this example, CP 2 is configured on both PE-1 and PE-2 with one VLAN range with one VID (11 or 21), as follows:

```
# on PE-1:
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
configure {
    connection-profile {
        vlan 2 {
            qtag-range 11 {
                end 11
            }
```

```
# on PE-2:
configure {
    connection-profile {
        vlan 2 {
            qtag-range 21 {
                end 21
            }
```

The VLAN SAP 1/2/1:1.11 is replaced by CP SAP 1/2/1:1.cp-2, as follows:

```
# on PE-1:
configure {
    service {
        vpls "VPLS 1" {
            delete sap 1/2/1:1.11
            sap sap 1/2/1:1.cp-2 {
            }
```

Likewise, the VLAN 1/2/1:1.21 is replaced by CP SAP 1/2/1:1.cp-2, as follows:

```
# on PE-2:
configure {
    service {
        vpls "VPLS 1" {
            delete sap 1/2/1:1.21
            sap sap 1/2/1:1.cp-2 {
            }
```

Figure 288: No traffic between two QinQ CP SAPs - VLAN translation not supported

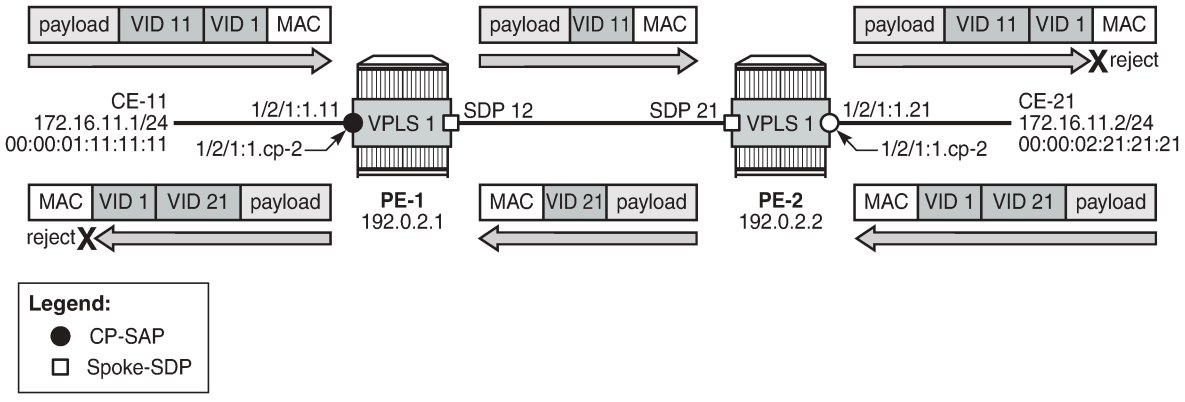

CE-11 sends double-tagged frames to SAP 1/2/1:1.cp-2 in VPLS 1 on PE-1. This CP SAP pops the outer tag with VID 1, but preserves the VLAN tag with VID 11. The single-tagged frame is sent to PE-2 where CP SAP 1/2/1:1.cp-2 pushes an outer tag with VID 1 onto the frame. This double-tagged frame is sent to CE-12 where it is rejected, because an inner label with VID 21 is expected.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

When CE-21 sends frames to CE-11, the frames will be double-tagged with inner tag VID 21 and outer tag 1. The outer tag is popped by the ingress SAP 1/2/1:1.cp-2 in VPLS 1 on PE-2, but the inner tag is preserved. The egress SAP 1/2/1:1.cp-2 in VPLS 1 on PE-1 preserves the inner tag with VID 21 and pushes an outer tag with VID 1. This double-tagged frame is rejected by CE-11, because another inner tag is expected, with VID 11 instead of VID 21.

Figure 289: Traffic between two QinQ CP SAPs - no VLAN translation shows how traffic is sent between two QinQ CP SAPs without VLAN translation. Both CE-18 and CE-28 send double-tagged frames with inner tag VID 8 and outer tag VID 100. The tag with VID 100 need not be the same on both CEs, because it is popped and pushed by the CP SAPs; only the tag with VID 8 from the VLAN range must be unchanged.

Figure 289: Traffic between two QinQ CP SAPs - no VLAN translation

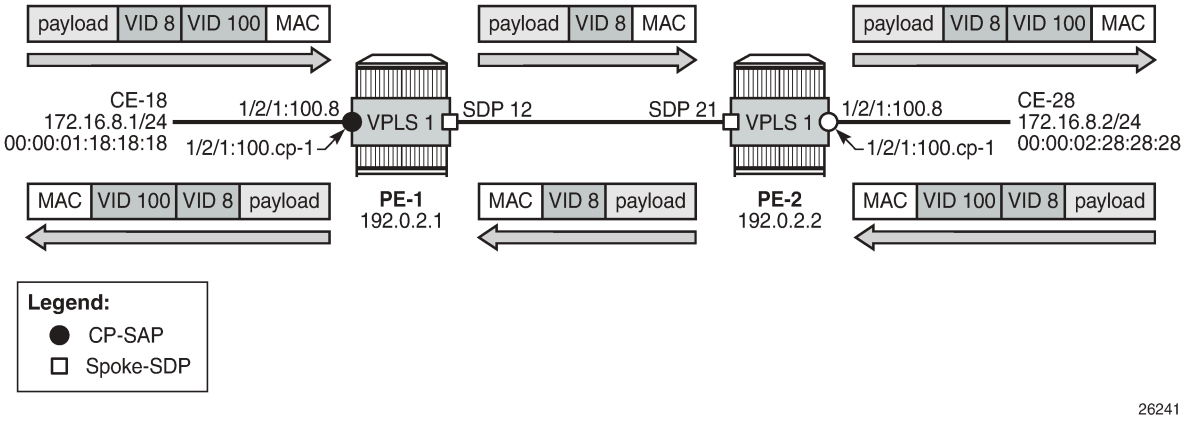

VLAN ranges in Epipe services

Figure 290: Example topology for VLAN ranges in Epipe 2 shows the example topology for VLAN ranges in Epipe 2.

Figure 290: Example topology for VLAN ranges in Epipe 2

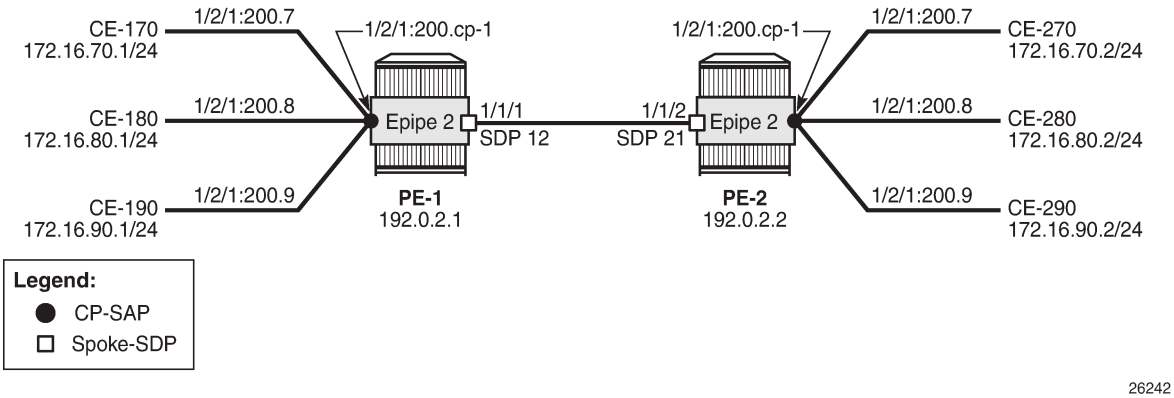

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Epipe 2 is configured with one CP SAP and a spoke-SDP, as follows:

```
# on PE-1:
configure {
    service {
        epipe "Epipe 2" {
            admin-state enable
            service-id 2
            customer "1"
            spoke-sdp 12:2 {
            }
            sap 1/2/1:200.cp-1 {
            }
        }
        sdp 12 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.2
            }
        }
```

CE-170 and CE-270 send double-tagged frames with inner VID 7 and outer VID 200. The inner VID 7 is preserved by the CP SAPs; therefore, CE-170 can only communicate with CE-270, not with any other CE at the other end, because they have different customer VIDs.

Conclusion

CP SAPs can be used to build services that can be bundled as per MEF 10.3 and RFC 7432. Multiple customer VIDs can be mapped to one CP-SAP.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

VXLAN Forwarding Path Extension

This chapter provides information about VXLAN Forwarding Path Extension.

Topics in this chapter include:

Applicability

This chapter was initially written based on SR OS Release 15.0.R4, but the MD-CLI in the current edition corresponds to SR OS Release 21.2.R2. Virtual eXtensible Local Area Network (VXLAN) Forwarding Path Extension (FPE) is supported in SR OS Release 14.0.R4, and later. IPv6 addresses are supported for EVPN-VXLAN BGP peering in SR OS release 15.0.R1, and later.

Overview

Use cases

VXLAN Forwarding Path Extension (FPE) is an SR OS feature that enables VXLAN tunnels to terminate on non-system IPv4 and IPv6 Destination Addresses (DAs). The non-system IPv4/IPv6 VXLAN termination feature can be applied in the following use cases:

- VXLAN Gateway (GW) in Software-Defined VPNs (SD-VPNs)
- VXLAN IPv6 underlay for Data Centers (DCs)

VXLAN GW in SD-VPNs

Traffic transported on a VXLAN is usually connected to a trusted environment through a VPRN running in a private IP/MPLS network. The VXLAN GW system IP address is used for all internal management and MPLS termination in the trusted network. However, in this use case, SR OS routers are expected to be used as a VXLAN GW in SD-VPNs where the VXLAN GW terminates untrusted VXLAN tunnels initiated on the SD-VPN CPEs and forwards packets to a trusted IP/MPLS network, as shown in Figure 291: VXLAN GW in an SD-VPN.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 291: VXLAN GW in an SD-VPN

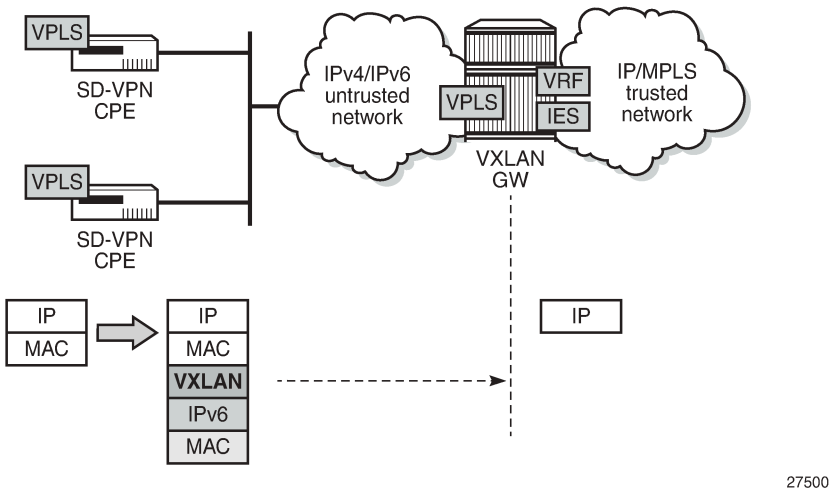

For security reasons, service providers will not expose system IP addresses to the untrusted IP network. Therefore, an IPv4 or IPv6 loopback address will be defined and used for VXLAN termination. The VXLAN tunnel can be terminated in a VPLS, an Epipe, or an R-VPLS service connected to a VPRN.

VXLAN IPv6 underlay for DCs

Some service providers migrate their entire network infrastructure to IPv6, including the DC network, so the DC GW must be able to terminate a VXLAN over an IPv6 infrastructure. Layer 2 (VPLS termination) and Layer 3 (R-VPLS termination) DC interconnect are both supported. Figure 292: VXLAN IPv6 underlay for DC shows the VXLAN IPv6 underlay for DC.

Figure 292: VXLAN IPv6 underlay for DC

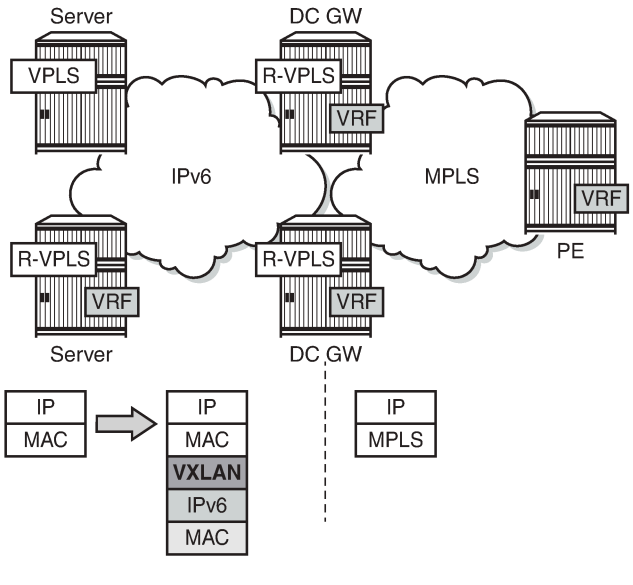

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

VXLAN FPE function

The following applies to VXLAN FPE:

- In an SR OS node, VXLAN tunnels can be terminated in four different VXLAN Tunnel Endpoints (VTEPs):
  - System IPv4 address
  - Up to three non-system IPv4/IPv6 addresses

  This limit is based on the number of supported source IP addresses that can be used for VXLAN encapsulation.
- The preceding four terminating IP addresses can be used in addition to the Assisted Replication IP address (AR IP). The AR IP does not count against this limit of four VTEPs. See chapter "Layer 2 Multicast Optimization for EVPN-VXLAN - Assisted Replication" in the Layer 2 Services and EVPN volume of the *7450 ESS, 7750 SR, and 7950 XRS Advanced Configuration Guide — Part II* for more information about AR.
- VXLAN FPE requires PXC ports; see the "Port Cross-Connect (PCX)" chapter in the Interface Configuration volume in the *7450 ESS, 7750 SR, and 7950 XRS MD-CLI Advanced Configuration Guide - Part I*.
  - Ingress traffic from a VXLAN with an IP DA equal to a loopback address will be redirected to the PXC port where the IP header will get additional processing.
  - Usually, only the ingress traffic from the VXLAN is redirected to the PXC port. The egress traffic to the VXLAN tunnel can go straight out of the egress network port, except for R-VPLS traffic toward an IPv6 VXLAN that is redirected to the PXC port.
- The VPLS/R-VPLS functionality is not impacted by the choice of VTEP termination (system IP address or not).

Provisioning model

Non-system IP VXLAN termination and VXLAN IPv6 underlay are both provisioned as per the following steps:

1. Create an FPE
2. Associate the FPE with VXLAN termination
3. Configure a router loopback interface
4. Configure non-system VXLAN termination VTEP addresses
5. Add the service configuration

Configuration

Figure 293: Example topology for VXLAN FPE shows the example topology with two PEs in an EVPN-VXLAN network. The loopback addresses in the base router will be used for non-system IP VXLAN termination.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 293: Example topology for VXLAN FPE

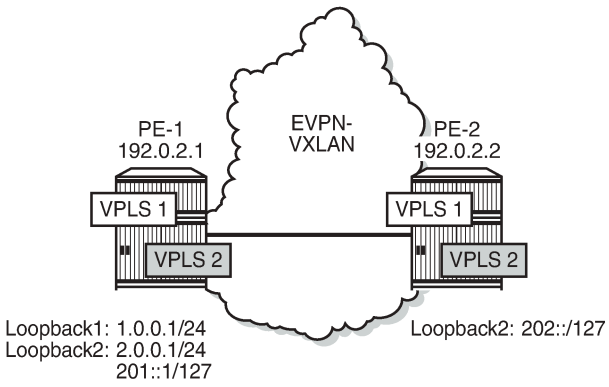

The initial configuration includes the cards, MDAs, ports, router interfaces and IGP. BGP is configured for address family EVPN, for example on PE-1 as follows:

```
# on PE-1:
configure {
    router "Base" {
        autonomous-system 64500
        bgp {
            rapid-withdrawal true
            split-horizon true
            rapid-update {
                evpn true
            }
            group "internal" {
                peer-as 64500
                family {
                    evpn true
                }
            }
            neighbor "192.0.2.2" {
                group "internal"
            }
        }
```

In this example, the BGP peering is IPv4-based, but EVPN-VXLAN routes can also be exchanged between IPv6 BGP peers.

Non-system IP VXLAN termination

Create FPEs

PXC is used as a simple back-to-back cross-connect. An FPE uses the PXC ports assigned in the FPE path, either a PXC port or a LAG-based PXC. For non-system IP VXLAN terminations between VPLSs, the PXC is only required on the ingress (from VXLAN, or from PE-1 to GW PE-2). PXC 1 and PXC 2 are created on PE-1, as follows:

```
# on PE-1:
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
configure {
    port 1/2/1 {
        admin-state enable
        ethernet {
            mode hybrid
            dot1x {
                tunneling true
            }
        }
    }
    port 1/2/2 {
        admin-state enable
        ethernet {
            mode hybrid
            dot1x {
                tunneling true
            }
        }
    }
    port pxc-1.a {
        admin-state enable
    }
    port pxc-1.b {
        admin-state enable
    }
    port pxc-2.a {
        admin-state enable
    }
    port pxc-2.b {
        admin-state enable
    }
    port-xc {
        pxc 1 {
            admin-state enable
            port-id 1/2/1
        }
        pxc 2 {
            admin-state enable
            port-id 1/2/2
        }
    }
```

The sub-ports of PXC 1 are operationally up, as follows.

```
[/]
A:admin@PE-1# show port pxc 1

===============================================================================
Ports on Port Cross Connect 1
===============================================================================
Port          Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id            State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
pxc-1.a       Up    Yes  Up      1574 1574    - hybr dotq xgige
pxc-1.b       Up    Yes  Up      1574 1574    - hybr dotq xgige
===============================================================================
```

The following FPEs use the PXCs.

```
# on PE-1, PE-2:
configure {
    fwd-path-ext {
        fpe 1 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            path {
                pxc 1
            }
        }
        fpe 2 {
            path {
                pxc 2
            }
        }
```

These FPEs are created without defining a range of SDP IDs. SDP IDs are required in case of R-VPLS services terminating IPv6 VXLAN tunnels, where the FPE is also used at the egress and an internal static SDP is created to allow for the required extra processing.

When the FPE has no VXLAN termination associated, no internal router interfaces are created, so the only router interfaces are the system interface and the interface between PE-1 and PE-2, as follows.

```
[/]
A:admin@PE-1# show router interface

===============================================================================
Interface Table (Router: Base)
===============================================================================
Interface-Name                   Adm       Opr(v4/v6)  Mode    Port/SapId
   IP-Address                                                  PfxState
-------------------------------------------------------------------------------
int-PE-1-PE-2                    Up        Up/Down     Network 1/1/1
   192.168.12.1/30                                             n/a
system                           Up        Up/Down     Network system
   192.0.2.1/32                                                n/a
-------------------------------------------------------------------------------
Interfaces : 2
===============================================================================
```

Associate the FPEs with VXLAN termination

The following command associates the FPEs with VXLAN termination.

```
# on PE-1, PE-2:
configure {
    fwd-path-ext {
        sdp-id-range {
            start 10000
            end 10127
        }
        fpe 1 {
            path {
                pxc 1
            }
            application {
                vxlan-termination {
                }
            }
        }
        fpe 2 {
            path {
                pxc 2
            }
            application {
                vxlan-termination {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
            }
        }
```

When attempting to associate the FPEs with VXLAN termination without configuring a range of SDP IDs for FPE, the following error is raised:

```
*[ex:/configure fwd-path-ext fpe 1 application vxlan-termination]
A:admin@PE-1# commit
MINOR: FPE #1021: configure fwd-path-ext fpe 1 - sdp-id-range is not configured - configure
 fwd-path-ext sdp-id-range
```

After the FPEs are associated with VXLAN terminations, the system creates two internal router interfaces per FPE, one per PXC sub-port:

```
[/]
A:admin@PE-1# show router interface

===============================================================================
Interface Table (Router: Base)
===============================================================================
Interface-Name                   Adm       Opr(v4/v6)  Mode    Port/SapId
   IP-Address                                                  PfxState
-------------------------------------------------------------------------------
_tmnx_fpe_1.a                    Up        Up/Up       Network pxc-1.a:1
   fe80::100/64                                                PREFERRED
_tmnx_fpe_1.b                    Up        Up/Up       Network pxc-1.b:1
   fe80::101/64                                                PREFERRED
_tmnx_fpe_2.a                    Up        Up/Up       Network pxc-2.a:1
   fe80::200/64                                                PREFERRED
_tmnx_fpe_2.b                    Up        Up/Up       Network pxc-2.b:1
   fe80::201/64                                                PREFERRED
int-PE-1-PE-2                    Up        Up/Down     Network 1/1/1
   192.168.12.1/30                                             n/a
system                           Up        Up/Down     Network system
   192.0.2.1/32                                                n/a
-------------------------------------------------------------------------------
Interfaces : 6
===============================================================================
```

Configure router loopback interfaces

The following loopback interfaces are configured in PE-1 and added to the IS-IS context:

```
# on PE-1:
configure {
    router "Base"
        interface "loopback1" {
            loopback
            ipv4 {
                primary {
                    address 1.0.0.1
                    prefix-length 24
                }
            }
        }
        interface "loopback2" {
            loopback
            ipv4 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                primary {
                    address 2.0.0.1
                    prefix-length 31
                }
            }
            ipv6 {
                address 201:: {
                    prefix-length 127
                }
            }
        }
        isis 0 {
            interface "loopback1" {
            }
            interface "loopback2" {
            }
        }
```

A non /32 or /128 subnet must be assigned to the loopback interface, because the system cannot terminate VXLAN on a local interface address. In the preceding example, all addresses in the subnet 1.0.0.0/24 can be used for VXLAN tunnel termination, except for 1.0.0.1. The subnet will be advertised by the IGP. The subnet can be as small as /31 or /127, as for example for interface "loopback2".

In this scenario, only one loopback interface with an IPv4 address is sufficient: interface "loopback1" with IPv4 address 1.0.0.1/24. There is no need to configure loopback interfaces in the GW PE-2, because VXLAN FPE is only required in the ingress (from VXLAN to GW).

Configure non-system VTEP addresses

Up to three non-system VTEP addresses can be configured to terminate VXLAN tunnels and their corresponding FPEs; on PE-1 as follows:

```
# on PE-1:
configure {
    service {
        system {
            vxlan {
                tunnel-termination 1.0.0.2 {
                    fpe-id 1
                }
                tunnel-termination 2.0.0.2 {
                    fpe-id 2
                }
                tunnel-termination 201::1 {
                    fpe-id 2
                }
```

No non-system VTEP addresses need to be configured on PE-2.

When attempting to configure the IP address of the loopback interface as a VXLAN tunnel termination, the following error is raised:

```
*[ex:/configure service system vxlan tunnel-termination 1.0.0.1]
A:admin@PE-1# commit
MINOR: MGMT_CORE #4001: configure service system vxlan tunnel-termination 1.0.0.1 - IP address
 matches a local interface IP address
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

When attempting to configure more than three non-system VTEP addresses, the following error is raised:

```
*[ex:/configure service system vxlan tunnel-termination 1.0.0.100]A:admin@PE-1# commit
MINOR: MGMT_CORE #232: configure service system vxlan tunnel-termination 1.0.0.100 - Reached
 maximum number of entries - maximum is 3 but has 4
```

When the non-system VTEP addresses are configured, an internal loopback interface "_tmnx_vli_vxlan_1_131077" is created that can respond to ICMP requests.

```
[/]
A:admin@PE-1# show router interface

===============================================================================
Interface Table (Router: Base)
===============================================================================
Interface-Name                   Adm       Opr(v4/v6)  Mode    Port/SapId
   IP-Address                                                  PfxState
-------------------------------------------------------------------------------
_tmnx_fpe_1.a                    Up        Up/Up       Network pxc-1.a:1
   fe80::100/64                                                PREFERRED
_tmnx_fpe_1.b                    Up        Up/Up       Network pxc-1.b:1
   fe80::101/64                                                PREFERRED
_tmnx_fpe_2.a                    Up        Up/Up       Network pxc-2.a:1
   fe80::200/64                                                PREFERRED
_tmnx_fpe_2.b                    Up        Up/Up       Network pxc-2.b:1
   fe80::201/64                                                PREFERRED
_tmnx_vli_vxlan_1_131077         Up        Up/Up       Network loopback
   1.0.0.2/32                                                  n/a
   2.0.0.2/32                                                  n/a
   201::1/128                                                  PREFERRED
   fe80::f:ffff:fe00:0/64                                      PREFERRED
int-PE-1-PE-2                    Up        Up/Down     Network 1/1/1
   192.168.12.1/30                                             n/a
loopback1                        Up        Up/Down     Network loopback
   1.0.0.1/24                                                  n/a
loopback2                        Up        Up/Up       Network loopback
   2.0.0.1/31                                                  n/a
   201::/127                                                   PREFERRED
   fe80::f:ffff:fe00:0/64                                      PREFERRED
system                           Up        Up/Down     Network system
   192.0.2.1/32                                                n/a
-------------------------------------------------------------------------------
Interfaces : 9
===============================================================================
```

The system does not verify whether there is a local base router loopback interface with a subnet corresponding to the VTEP address. If a tunnel termination address is configured and the FPE is up, the system will start terminating VXLAN traffic and responding using ICMP for that address, regardless of the presence of a loopback interface in the base router. It is also possible that a non-loopback interface has an IP address in the configured subnet.

Configure the VPLS

A VPLS will be configured with EVPN-VXLAN enabled. By default, the system IP address will be used as the source VTEP of the VXLAN-encapsulated frames. This default behavior can be overruled by the

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

source-vtep command in the VPLS. The IP address corresponds to the non-system VTEP address configured in the preceding step (VXLAN tunnel termination). VPLS 1 is configured on PE-1 as follows:

```
# on PE-1:
configure {
    service {
        vpls "EVI-1" {
            admin-state enable
            service-id 1
            customer "1"
            vxlan {
                source-vtep 1.0.0.2
                instance 1 {
                    vni 1
                }
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 1
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
            sap 1/1/2:1 {
            }
```

When attempting to configure an IP address different from the VTEP addresses, the following error is raised:

```
[ex:/configure service vpls "EVI-1" vxlan]
A:admin@PE-1# source-vtep 1.0.0.99

*[ex:/configure service vpls "EVI-1" vxlan]
A:admin@PE-1# commit
MINOR: MGMT_CORE #224: configure service vpls "EVI-1" vxlan source-vtep - Entry does not
 exist - configure service system vxlan tunnel-termination 1.0.0.99
```

A different VTEP address can be configured as **source-vtep** in different services on the same PE, as follows:

```
# on PE-1:
configure {
    service {
        vpls "EVI-2" {
            admin-state enable
            service-id 2
            customer "1"
            vxlan {
                source-vtep 201::1
                instance 1 {
                    vni 2
                }
            }
            routed-vpls {
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 2
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
        }
```

The configuration of VPLS 1 on PE-2 does not include any VTEP address, because it is not required in the egress, as follows:

```
# on PE-2:
configure {
    service {
        vpls "EVI-1" {
            admin-state enable
            service-id 1
            customer "1"
            vxlan {
                instance 1 {
                    vni 1
                }
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 1
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
        }
```

When a source VTEP is configured in VPLS 1 on PE-1, this VTEP address will be used as the IP source VTEP for VPLS 1 and BGP will use this VTEP to the BGP NLRI next-hop, as shown in the following BGP route update messages.

The following BGP EVPN inclusive multicast route sent by PE-1 shows the configured source VTEP address 1.0.0.2 as NLRI next-hop, as originator address, and as tunnel endpoint.

```
# on PE-1:
1 2021/05/06 09:40:06.914 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 77
    Flag: 0x90 Type: 14 Len: 28 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 1.0.0.2
        Type: EVPN-INCL-MCAST Len: 17 RD: 192.0.2.1:1, tag: 0, orig_addr len: 32,
                              orig_addr: 1.0.0.2
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:1
        bgp-tunnel-encap:VXLAN
    Flag: 0xc0 Type: 22 Len: 9 PMSI:
        Tunnel-type Ingress Replication (6)
        Flags: (0x0)[Type: None BM: 0 U: 0 Leaf: not required]
        MPLS Label 1
        Tunnel-Endpoint 1.0.0.2
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
"
```

The following BGP EVPN-MAC route sent by PE-1 shows the configured VTEP for VPLS 1 as NLRI next-hop:

```
# on PE-1:
8 2021/05/06 09:41:43.212 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 81
    Flag: 0x90 Type: 14 Len: 44 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 1.0.0.2
        Type: EVPN-MAC Len: 33 RD: 192.0.2.1:1 ESI: ESI-0, tag: 0, mac len: 48
                       mac: ca:fe:01:10:10:10, IP len: 0, IP: NULL, label1: 1
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 16 Extended Community:
        target:64500:1
        bgp-tunnel-encap:VXLAN
"
```

A BGP peer policy might override the NLRI next-hop created due to the **source-vtep** configuration.

The following shows that the source VTEP address on PE-1 is 1.0.0.2:

```
[/]
A:admin@PE-1# show service id 1 vxlan
===============================================================================
VPLS VXLAN
===============================================================================
Vxlan Src Vtep IP: 1.0.0.2

===============================================================================
Vxlan Instance
===============================================================================
VXLAN Instance                VNI         AR          Oper-flags   VTEP
                                                                   security
-------------------------------------------------------------------------------
1                             1           none        none         disabled
-------------------------------------------------------------------------------
Number of Entries : 1
-------------------------------------------------------------------------------
===============================================================================
```

The following command on PE-1 shows that the egress VTEP is 192.0.2.2:

```
[/]
A:admin@PE-1# show service id 1 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
Instance    VTEP Address                          Egress VNI  EvpnStatic Num
 Mcast       Oper State                            L2 PBR     SupBcasDom MACs
-------------------------------------------------------------------------------
1           192.0.2.2                             1           evpn       0
 BUM         Up                                    No          No
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN-VXLAN Ethernet Segment Dest
===============================================================================
Instance  Eth SegId                       Num. Macs     Last Change
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

The following shows that no source VTEP address is configured on PE-2:

```
[/]
A:admin@PE-2# show service id 1 vxlan
===============================================================================
VPLS VXLAN
===============================================================================
Vxlan Src Vtep IP: N/A

===============================================================================
Vxlan Instance
===============================================================================
VXLAN Instance                VNI         AR          Oper-flags   VTEP
                                                                   security
-------------------------------------------------------------------------------
1                             1           none        none         disabled
-------------------------------------------------------------------------------
Number of Entries : 1
-------------------------------------------------------------------------------
===============================================================================
```

The following command on PE-2 shows that the egress VTEP is 1.0.0.2:

```
[/]
A:admin@PE-2# show service id 1 vxlan destinations

===============================================================================
Egress VTEP, VNI
===============================================================================
Instance    VTEP Address                            Egress VNI  EvpnStatic Num
 Mcast       Oper State                              L2 PBR     SupBcasDom MACs
-------------------------------------------------------------------------------
1           1.0.0.2                                 1           evpn       1
 BUM         Up                                      No          No
-------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 1
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
BGP EVPN-VXLAN Ethernet Segment Dest
===============================================================================
Instance  Eth SegId                       Num. Macs     Last Change
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Underlay IPv6 VXLAN termination

The configuration for underlay IPv6 VXLAN termination is similar to the non-system IP VXLAN termination. In the following example, R-VPLS 2 is configured; therefore, non-system VTEP addresses are configured in PE-2 as well as in PE-1. The changes required in PE-1 are as follows.

- IPv6 must be enabled on the router interfaces
- IPv6 native routing is configured in IS-IS
- IPv6 addresses are loopback address 201::/127 and VTEP address 201::1

```
# on PE-1:
configure {
    port-xc {
        pxc 2 {
            admin-state enable
            port-id 1/2/2
        }
    }
    port pxc-2.a {
        admin-state enable
    }
    port pxc-2.b {
        admin-state enable
    }
    port 1/2/2 {
        admin-state enable
        ethernet {
            mode hybrid
            dot1x {
                tunneling true
            }
        }
    }
    fwd-path-ext {
        sdp-id-range {
            start 10000
            end 10127
        }
        fpe 2 {
            path {
                pxc 2
            }
            application {
                vxlan-termination {
                }
            }
        }
    }
    router "Base" {
        interface "int-PE-1-PE-2" {
            port 1/1/1
            ipv4 {
                primary {
                    address 192.168.12.1
                    prefix-length 30
                }
            }
            ipv6 {
            }
        }
        interface "loopback2" {
            loopback
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            ipv4 {
                primary {
                    address 2.0.0.1
                    prefix-length 31
                }
            }
            ipv6 {
                address 201:: {
                    prefix-length 127
                }
            }
        }
        isis 0 {
            ipv6-routing native
            interface "loopback2" {
            }
        }
    }
    service {
        system {
            vxlan {
                tunnel-termination 201::1 {
                    fpe-id 2
                }
            }
        }
        vpls "EVI-2" {
            admin-state enable
            service-id 2
            customer "1"
            vxlan {
                source-vtep 201::1
                instance 1 {
                    vni 2
                }
            }
            routed-vpls {
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 2
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
        }
```

The service configuration on PE-2 is as follows.

```
# on PE-2:
configure {
    service {
        system {
            vxlan {
                tunnel-termination 202:: {
                    fpe-id 2
                }
            }
        }
        vpls "EVI-2" {
            admin-state enable
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            service-id 2
            customer "1"
            vxlan {
                source-vtep 202::
                instance 1 {
                    vni 2
                }
            }
            routed-vpls {
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 2
                vxlan 1 {
                    admin-state enable
                    vxlan-instance 1
                }
            }
        }
```

The routing table for IPv6 on PE-1 shows that an internal static route is configured for the source VTEP 201::1 using the FPE internal interface "_tmnx_fpe_2.a". The route to egress VTEP 202:: is an IS-IS route.

```
[/]
A:admin@PE-1# show router route-table ipv6

===============================================================================
IPv6 Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
201::/127                                       Local   Local     00h16m34s  0
       loopback2                                                       0
201::1/128                                      Remote  Static    00h11m53s  5
       fe80::201-"_tmnx_fpe_2.a"                                       1
202::/127                                       Remote  ISIS      00h00m06s  15
       fe80::14:1ff:fe01:1-"int-PE-1-PE-2"                             10
-------------------------------------------------------------------------------
No. of Routes: 3
```

Likewise, the routing table for IPv6 on PE-2 shows an internal static route for source VTEP 202:: using the FPE internal interface "_tmnx_fpe_2.a":

```
[/]
A:admin@PE-2# show router route-table ipv6

===============================================================================
IPv6 Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
201::/127                                       Remote  ISIS      00h00m06s  15
       fe80::10:1ff:fe01:1-"int-PE-2-PE-1"                             10
202::/127                                       Local   Local     00h00m12s  0
       loopback2                                                       0
202::/128                                       Remote  Static    00h00m13s  5
       fe80::201-"_tmnx_fpe_2.a"                                       1
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
No. of Routes: 3
```

When non-system IPv6 VTEP addresses are used in an R-VPLS, VTEP addresses need to be configured on ingress and egress VXLAN. The system creates an internal SDP binding for the egress processing. A range of SDP IDs has been configured from 10000 to 10127. The following command lists all SDP bindings for FPE:

```
[/]
A:admin@PE-2# show service sdp-using | match "Fpe"
2          10002:2            Fpe    fpe_2.b              Up    524287  524287
```

The internal SDP has ID 10002 and the far-end is fpe_2.b. The following command shows that the SDP source is FPE.

```
[/]
A:admin@PE-2# show service sdp 10002 detail | match "Sdp" pre-lines 4 post-lines 10

===============================================================================
Service Destination Point (Sdp Id : 10002) Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 10002  -fpe_2.b
-------------------------------------------------------------------------------
Description           : (Not Specified)
SDP Id               : 10002SDP Source         : fpe
Admin Path MTU       : 0                    Oper Path MTU      : 1552
Delivery             : MPLS
Far End              : fpe_2.b              Tunnel Far End     : n/a
Oper Tunnel Far End  : n/a
LSP Types            : FPE

Admin State          : Up                   Oper State         : Up
```

The following command on PE-1 shows that the source VTEP is 201::1:

```
[/]
A:admin@PE-1# show service id 2 vxlan
===============================================================================
VPLS VXLAN
===============================================================================
Vxlan Src Vtep IP: 201::1

===============================================================================
Vxlan Instance
===============================================================================
VXLAN Instance               VNI         AR          Oper-flags   VTEP
                                                                  security
-------------------------------------------------------------------------------
1                            2           none        none         disabled
-------------------------------------------------------------------------------
Number of Entries : 1
-------------------------------------------------------------------------------
===============================================================================
```

The following command on PE-1 shows that the egress VTEP is 202::.

```
[/]
A:admin@PE-1# show service id 2 vxlan destinations

===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Egress VTEP, VNI
==============================================================================
Instance    VTEP Address                          Egress VNI  EvpnStatic Num
 Mcast       Oper State                            L2 PBR     SupBcasDom MACs
------------------------------------------------------------------------------
1           202::                                 2           evpn       0
 BUM         Up                                    No          No
------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 1
------------------------------------------------------------------------------
==============================================================================

==============================================================================
BGP EVPN-VXLAN Ethernet Segment Dest
==============================================================================
Instance  Eth SegId                     Num. Macs     Last Change
------------------------------------------------------------------------------
No Matching Entries
==============================================================================
```

The following command on PE-2 shows that the source VTEP is 202::.

```
[/]
A:admin@PE-2# show service id 2 vxlan
==============================================================================
VPLS VXLAN
==============================================================================
Vxlan Src Vtep IP: 202::

==============================================================================
Vxlan Instance
==============================================================================
VXLAN Instance                VNI         AR          Oper-flags   VTEP
                                                                   security
------------------------------------------------------------------------------
1                             2           none        none         disabled
------------------------------------------------------------------------------
Number of Entries : 1
------------------------------------------------------------------------------
==============================================================================
```

The following command on PE-2 shows that the egress VTEP is 201::1.

```
[/]
A:admin@PE-2# show service id 2 vxlan destinations

==============================================================================
Egress VTEP, VNI
==============================================================================
Instance    VTEP Address                          Egress VNI  EvpnStatic Num
 Mcast       Oper State                            L2 PBR     SupBcasDom MACs
------------------------------------------------------------------------------
1           201::1                                2           evpn       0
 BUM         Up                                    No          No
------------------------------------------------------------------------------
Number of Egress VTEP, VNI : 1
------------------------------------------------------------------------------
==============================================================================

==============================================================================
BGP EVPN-VXLAN Ethernet Segment Dest
==============================================================================
Instance  Eth SegId                     Num. Macs     Last Change
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

Conclusion

VXLAN FPE is required to terminate VXLAN tunnels on non-system IPv4/IPv6 addresses and to configure IPv6 underlay.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Layer 3 Services

This section provides configuration information for the following topics:

- BGP Best External in a VPRN
- Carrier Supporting Carrier IP VPNs
- Flexible Algorithms for SRv6-based VPRNs
- Inter-AS VPRN Model B
- Inter-AS VPRN Model B Using MPLS over UDP
- Inter-AS VPRN Model C
- Layer 3 VPN: VPRN Type Spoke
- Spoke Termination for IPv6-6VPE
- Traffic Leaking from VPRN to GRT
- Weighted ECMP for VPRN over RSVP-TE or SR-TE LSPs

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

BGP Best External in a VPRN

This chapter provides information about BGP Best External in a VPRN.

Topics in this chapter include:

Applicability

The information and configuration in this chapter was originally written for SR OS Release 14.0.R7. In the current edition, the MD-CLI is updated to SR OS Release 22.2.R2.

Overview

By default, BGP speakers only advertise their best route for a destination. The BGP best external feature allows BGP speakers to advertise their best external route for a prefix Network Layer Reachability Information (NLRI) to their IBGP peers when their best overall route for this prefix NLRI is an internal route. This feature provides additional path visibility to the IBGP mesh. When two paths are available to reach a destination, and one is preferred, the availability of an alternate path in the RIB means that only a FIB update is required if the preferred next-hop fails. Also, the presence of two paths can reduce route oscillation.

BGP best external can be enabled in the base router with the following command:

```
[ex:/configure router "Base" bgp]
A:admin@PE-2# advertise-external ?

 advertise-external

 ipv4                  - Enable support for unlabeled unicast IPv4 routes
 ipv6                  - Enable support for unlabeled unicast IPv6 routes
 label-ipv4            - Enable support for labeled-unicast IPv4 routes
 label-ipv6            - Enable support for labeled-unicast IPv6 routes
---snip---
```

```
# on PE-2:
configure {
    router "Base" {
        bgp {
            advertise-external {
                ipv4 true
            }
```

Chapter "BGP Add-Path" in the Unicast Routing Protocols volume of *7450 ESS, 7750 SR, and 7950 XRS Advanced Configuration Guide — Book I* describes the use of the add-paths parameter for different address families. Chapter "BGP Fast Reroute" in the Unicast Routing Protocols volume in the *7450 ESS, 7750 SR, and 7950 XRS MD-CLI Advanced Configuration Guide - Part I* includes a configuration example

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

with BGP best external enabled in the base router, whereas this chapter focuses on BGP best external in a VPRN context.

VPRN BGP best external can be configured with the following command:

```
[ex:/configure service]
A:admin@PE-2# vprn "VPRN 1" ?
---snip---

 export-inactive-bgp   - Export preferred BGP route even if inactive
---snip---
```

```
# on PE-2:
configure {
    service {
        vprn "VPRN 1" {
            export-inactive-bgp true
```

VPRN BGP best external allows the best EBGP IPv4/IPv6 route learned by a VPRN to be exported as a BGP VPN-IPv4/IPv6 route, even when that EBGP IPv4/IPv6 route is inactive due to the presence of a preferred BGP VPN-IPv4/IPv6 route from another PE. This best external route advertisement is useful in active/standby multi-homing scenarios because it can ensure that all PEs have knowledge of the backup path provided by the standby PE, thus reducing convergence times. VPRN BGP best external can also be applied in combination with Equal Cost Multi-Path (ECMP).

Figure 294: CE-4 advertises prefix 10.0.0.0/8 to its EBGP peers PE-1 and PE-2 shows the example topology with CE-4 in Autonomous System (AS) 64500 advertising prefix 10.0.0.0/8 to VPRN 1 in PE-1 and PE-2 in AS 64496.

Figure 294: CE-4 advertises prefix 10.0.0.0/8 to its EBGP peers PE-1 and PE-2

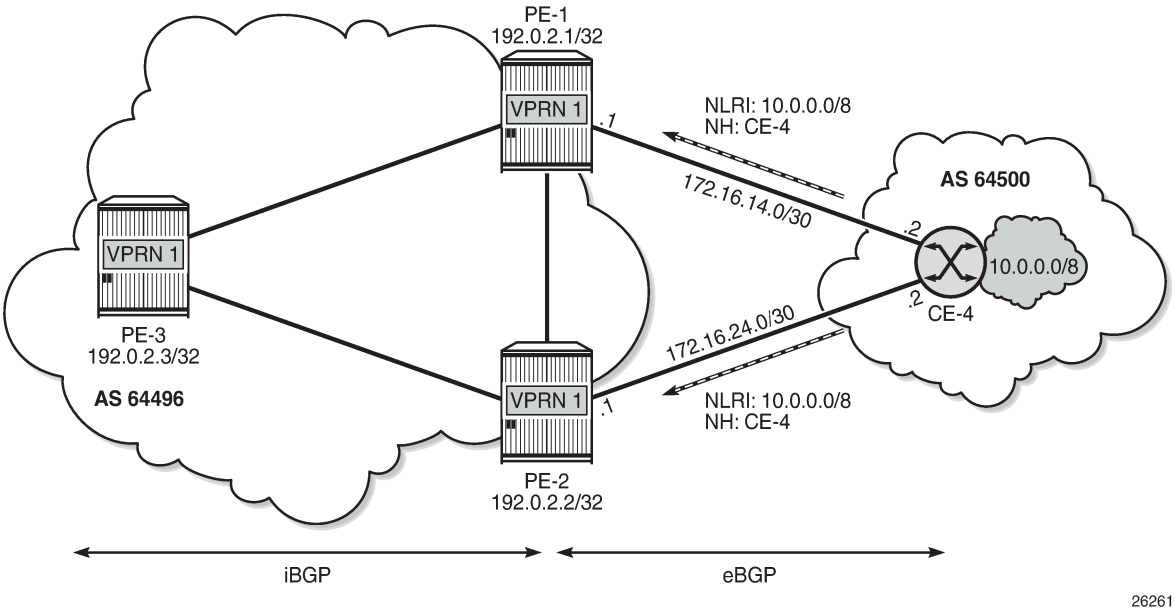

PE-1 is the primary PE for this prefix and it creates a corresponding BGP VPN-IPv4 route with a higher local preference (LP) value (for example, 200) compared to the default LP (100). PE-1 advertises this BGP VPN-IPv4 route to its IBGP peers PE-2 and PE-3. PE-2 imports this BGP VPN-IPv4 route into its

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

VRF, which deactivates the EBGP route received from CE-4, because it has the default LP of 100 (by BGP selection rules, the highest LP wins). By default, BGP prevents PE-2 from exporting its inactive BGP IPv4 route from CE-4 and, therefore, PE-1 and PE-3 cannot learn a BGP VPN-IPv4 backup route for prefix 10.0.0.0/8, as shown in Figure 295: Default BGP behavior: BGP advertises best route only.

Figure 295: Default BGP behavior: BGP advertises best route only

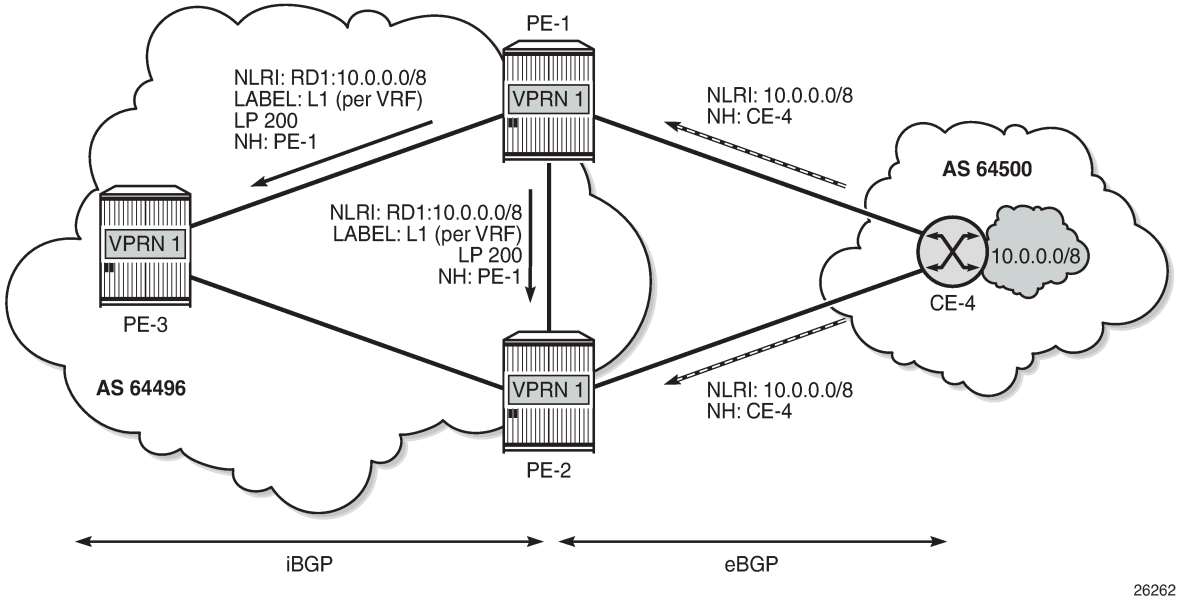

VPRN BGP best external allows PE-2 to advertise its best external route as backup on the following conditions:

- The option **export-inactive-bgp true** is configured in VPRN 1 on PE-2 (or on all PEs in the multi-homed site).
- The BGP route from CE-4 must match the VRF export policy in PE-2.
- The BGP VPN-IPv4 route exported by PE-2 must have a unique NLRI (RD:IP prefix combination) that does not overlap with a BGP VPN-IPv4 route from another PE for the same prefix. Therefore, a different RD can be allocated to the VRF in each PE connected to the multi-homed site. For example, VPRN 1 in PE-1 has RD 64496:11 and VPRN 1 in PE-2 has RD 64496:12.

Figure 296: VPRN BGP best external enabled: BGP advertises active and standby routes shows the BGP route advertisements when VPRN BGP best external is enabled. The BGP VPN-IPv4 route from PE-2 carries a per-next-hop label (meaning pop and forward to CE-4) regardless of the configured label mode of the VPRN service in PE-2.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 296: VPRN BGP best external enabled: BGP advertises active and standby routes

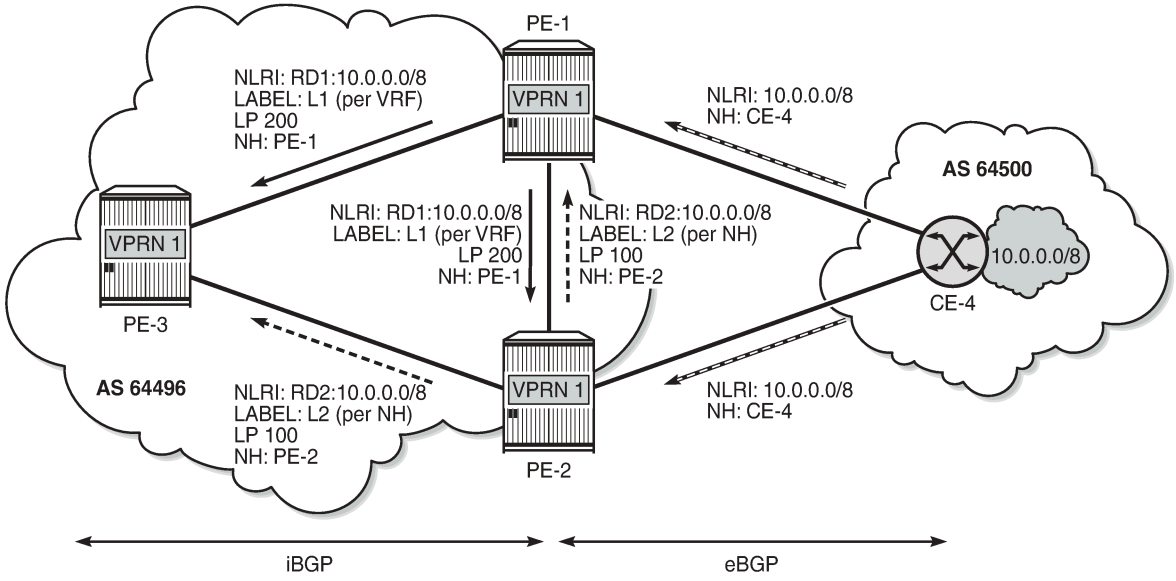

The PEs support BGP Fast Reroute (BGP FRR) using BGP VPN-IPv4 routes; therefore, PE-1 and PE-3 can install the route advertised by PE-2 as a backup path for prefix 10.0.0.0/8 and use that path immediately after detecting that the primary path has failed. When the link between PE-1 and CE-4 fails, PE-1 will detect this link failure typically seconds before the other PEs do. Therefore, PE-3 keeps sending traffic toward the network 10.0.0.0/8 to PE-1 and PE-1 uses the repair path via PE-2, as shown in Figure 297: BGP FRR on PE-1 after failure of active link to CE.

Figure 297: BGP FRR on PE-1 after failure of active link to CE

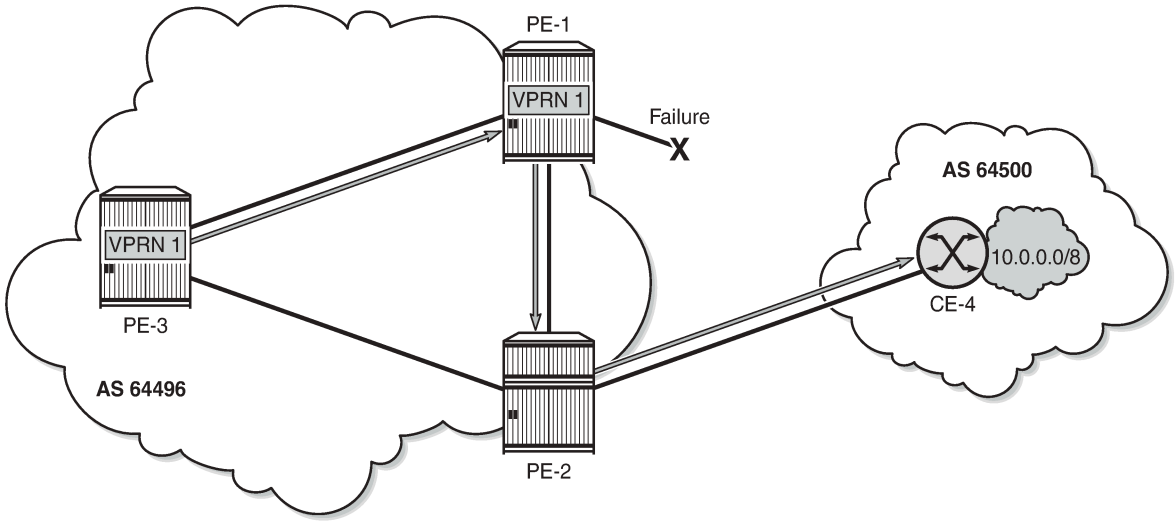

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Even when PE-2 is still unaware of the link failure between PE-1 and CE-4, PE-2 will not loop traffic back to PE-1. The reason is that PE-1 sends traffic to PE-2 with a per-next-hop label so that no FIB lookup occurs in PE-2. Traffic is forwarded correctly to CE-4.

When PE-2 receives the BGP VPN-IPv4 route withdrawal from PE-1 for prefix 10.0.0.0/8, it removes the route from its RIB-IN and reruns the decision process. In this example, the EBGP route to CE-4 becomes the new primary/best path. PE-2 will re-advertise its BGP VPN-IPv4 route for prefix 10.0.0.0/8. The difference is that the BGP VPN-IPv4 route is based on the export of an active/used route and, therefore, the advertised label value is based on the configured label mode of the VPRN service, as shown in Figure 298: PE-2 re-advertises VPN-IPv4 route with label based on VRF for label mode VRF (default).

Figure 298: PE-2 re-advertises VPN-IPv4 route with label based on VRF

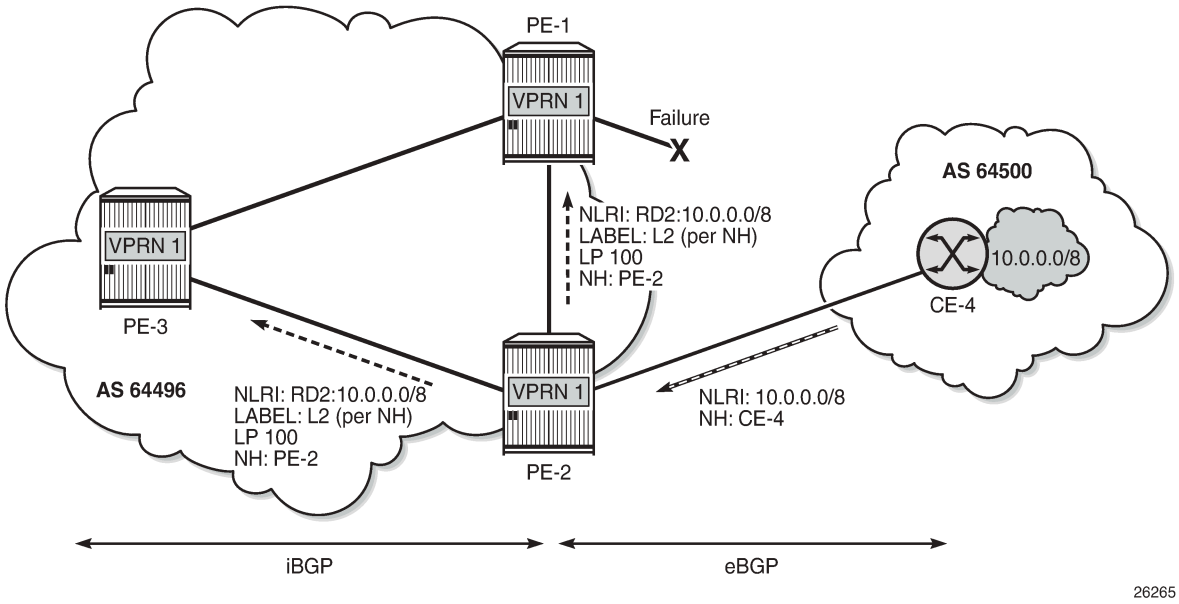

It takes time for this route to reach all ingress routers and for these routers to update their forwarding tables to use the per-VRF label value. For a while, there may still be traffic destined for prefix 10.0.0.0/8 that is received by PE-2 with the per-next-hop label L2. Traffic will be dropped if the per-next-hop label is deleted by the IOM as soon as PE-2 determines there are no more inactive/standby paths with CE-4 as next hop. Traffic loss can be avoided by delaying the deletion of per-next-hop labels in the IOM by configuring label retention for BGP labels with the following command:

```
*[ex:/configure router "Base" mpls-labels]
A:admin@PE-2# bgp-labels-hold-timer ?

 bgp-labels-hold-timer <number>
 <number>  - <0..255>
 Default   - 0

    BGP labels hold timer for the ingress router
```

```
# on PE-2:
configure {
    router "Base" {
        mpls-labels {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            bgp-labels-hold-timer 60
```

Finally, all ingress routers have updated their forwarding tables based on the BGP update sent by PE-2, and PE-3 sends traffic for prefix 10.0.0.0/8 directly toward PE-2, as shown in Figure 299: Traffic destined for prefix 10.0.0.0/8 after control plane convergence.

Figure 299: Traffic destined for prefix 10.0.0.0/8 after control plane convergence

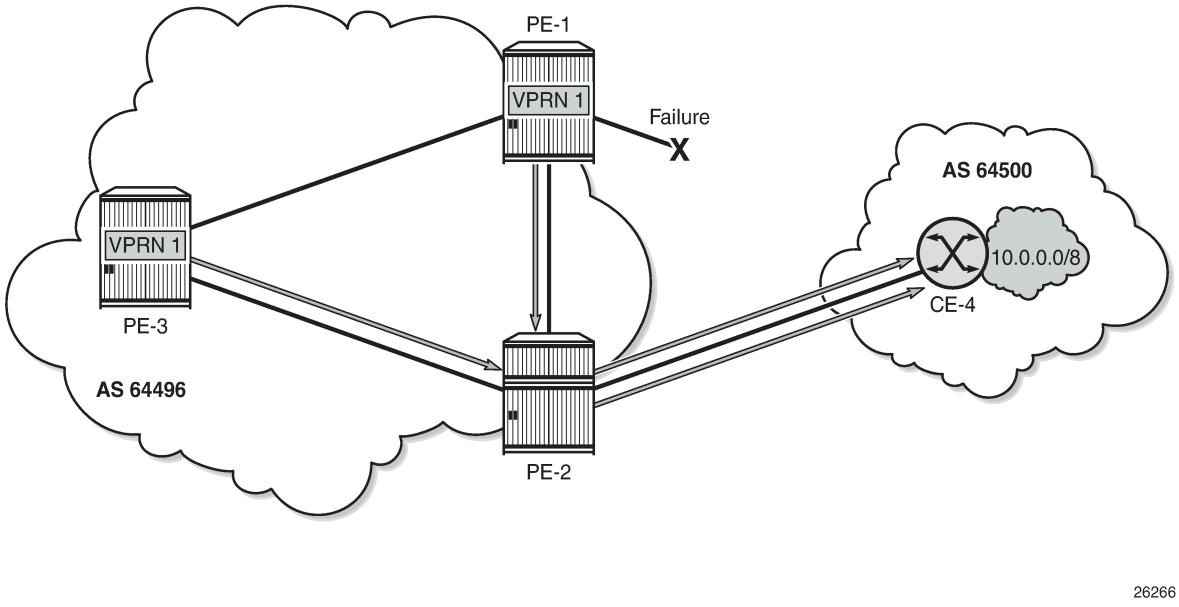

Configuration

Figure 300: Example topology shows the example topology with the used IP addresses.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 300: Example topology

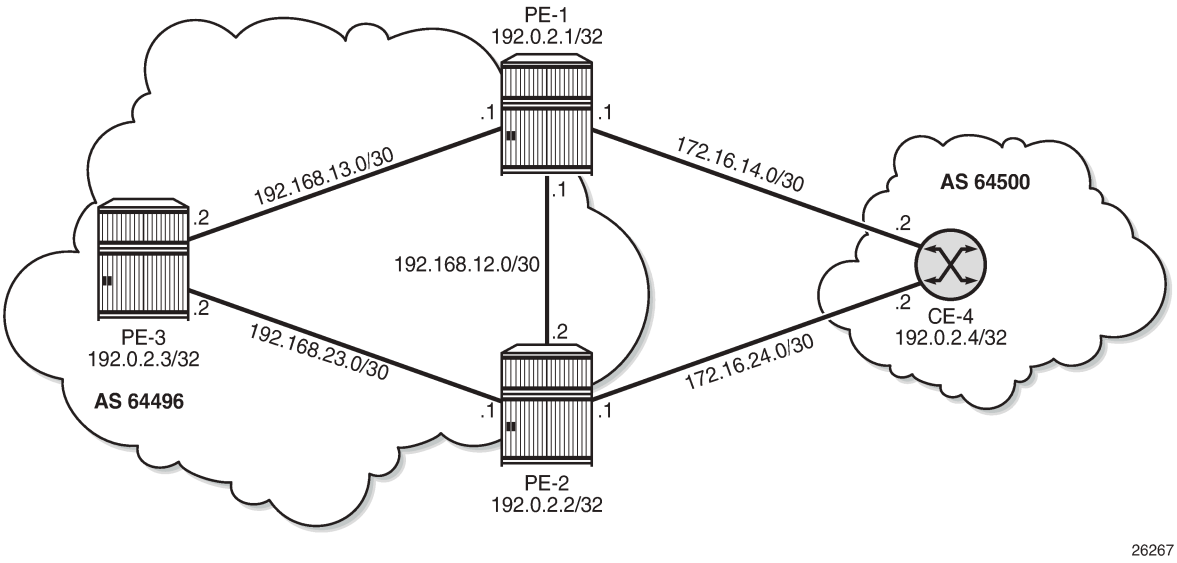

The initial configuration includes the following:

- Cards, MDAs, ports
- Router interfaces
- IS-IS (or OSPF) as IGP within AS 64496
- LDP on all interfaces within AS 64496

BGP is configured in the base router context of all PEs for address family VPN-IPv4; for example, for PE-1 as follows:

```
# on PE-1:
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            rapid-withdrawal true
            group "IBGP" {
                peer-as 64496
                family {
                    vpn-ipv4 true
                }
            }
            neighbor "192.0.2.2" {
                group "IBGP"
            }
            neighbor "192.0.2.3" {
                group "IBGP"
            }
```

The BGP configuration for the base router on the other two PEs is similar and a full mesh is established in AS 64496.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Configure VPRN without BGP best external

VPRN 1 is created on all PEs with the following settings:

- Default label mode: **label-mode vrf**
- Ready for BGP FRR: **bgp-vpn-backup>ipv4 true**
- Different RDs in VPRN 1 for each PE: 64496:11 on PE-1, 64496:12 on PE-2, and 64496:13 on PE-3
- **Auto-bind-tunnel** with **resolution any**. In this example, LDP will be used.
- Loopback interface "lo0" with IP address 172.31.2.1/32 on PE-1, which is also defined as the router ID in VPRN 1. The same approach is used on PE-2 and PE-3: 172.31.2.2/32 and 172.31.2.3/32.
- IBGP between all PEs (full mesh) for address family IPv4
- EBGP between PE-1 and CE-4 and between PE-2 and CE-4
- BGP best external is disabled, by default.

The configuration of VPRN 1 on PE-3 is as follows:

```
# on PE-3:
configure {
    service {
        vprn "VPRN 1" {
            admin-state enable
            service-id 1
            customer "1"
            autonomous-system 64496
            router-id 172.31.2.3
            # label-mode vrf    # default
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:13"
                    vrf-target {
                        community "target:64496:1"
                    }
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            bgp {
                rapid-withdrawal true
                group "IBGP" {
                    peer-as 64496
                }
                neighbor "172.31.2.1" {
                    group "IBGP"
                }
                neighbor "172.31.2.2" {
                    group "IBGP"
                }
            }
            interface "lo0" {
                loopback true
                ipv4 {
                    primary {
                        address 172.31.2.3
                        prefix-length 32
                    }
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
            bgp-vpn-backup {        # enable BGP FRR
                ipv4 true
            }
```

On PE-1 and PE-2, the VPRN configuration includes an external interface toward CE-4, and EBGP is defined toward peer CE-4. The VPRN 1 configuration on PE-2 is as follows:

```
# on PE-2:
configure {
    policy-options {
        policy-statement "import-bgp-LP100" {
            default-action {
                action-type accept
                local-preference 100
            }
        }
    }
    service {
        vprn "VPRN 1" {
            admin-state enable
            service-id 1
            customer "1"
            autonomous-system 64496
            router-id 172.31.2.2
            # label-mode vrf    # default
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:12"
                    vrf-target {
                        community "target:64496:1"
                    }
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            bgp {
                rapid-withdrawal true
                split-horizon true
                group "EBGP" {
                    peer-as 64500
                    import {
                        policy ["import-bgp-LP100"]
                    }
                }
                group "IBGP" {
                    peer-as 64496
                }
                neighbor "172.16.24.2" {
                    group "EBGP"
                }
                neighbor "172.31.2.1" {
                    group "IBGP"
                }
                neighbor "172.31.2.3" {
                    group "IBGP"
                }
            }
            interface "int-PE-2-CE-4_VPRN1" {
                ipv4 {
                    primary {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        address 172.16.24.1
                        prefix-length 30
                    }
                }
                sap 1/1/3:1 {
                }
            }
            interface "lo0" {
                loopback true
                ipv4 {
                    primary {
                        address 172.31.2.2
                        prefix-length 32
                    }
                }
            }
            bgp-vpn-backup {        # enable BGP FRR
                ipv4 true
            }
```

PE-2 has an import policy that sets the LP to 100, which is the default LP, but without import policy, no EBGP routes are accepted by default.

The VPRN 1 configuration on PE-1 looks similar to the configuration on PE-2, but includes an import policy that assigns an LP of 200 to each prefix that is received from CE-4, as follows:

```
# on PE-1:
configure {
    policy-options {
        policy-statement "import-bgp-LP200" {
            default-action {
                action-type accept
                local-preference 200
            }
        }
    }
    service {
        vprn "VPRN 1" {
            admin-state enable
            service-id 1
            customer "1"
            autonomous-system 64496
            router-id 172.31.2.1
            # label-mode vrf    # default
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:11"
                    vrf-target {
                        community "target:64496:1"
                    }
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            bgp {
                rapid-withdrawal true
                split-horizon true
                group "EBGP" {
                    peer-as 64500
                    import {
                        policy ["import-bgp-LP200"]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
            }
            group "IBGP" {
                peer-as 64496
            }
            neighbor "172.16.14.2" {
                group "EBGP"
            }
            neighbor "172.31.2.2" {
                group "IBGP"
            }
            neighbor "172.31.2.3" {
                group "IBGP"
            }
        }
        interface "int-PE-1-CE-4_VPRN1" {
            ipv4 {
                primary {
                    address 172.16.14.1
                    prefix-length 30
                }
            }
            sap 1/1/3:1 {
            }
        }
        interface "lo0" {
            loopback true
            ipv4 {
                primary {
                    address 172.31.2.1
                    prefix-length 32
                }
            }
        }
        bgp-vpn-backup {        # enable BGP FRR
            ipv4 true
        }
    }
```

CE-4 has EBGP configured toward PE-1 and PE-2. CE-4 exports the prefix 10.0.0.0/8, as defined in export policy "export-bgp" that is applied in the **bgp** context:

```
# on CE-4:
configure {
    policy-options {
        prefix-list "10.0.0.0/8" {
            prefix 10.0.0.0/8 type longer {
            }
        }
        policy-statement "export-bgp" {
            entry 10 {
                from {
                    prefix-list ["10.0.0.0/8"]
                }
                action {
                    action-type accept
                }
            }
        }
    }
    router "Base" {
        autonomous-system 64500
        interface "int-CE-4-PE-1_VPRN1" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            port 1/1/1:1
            ipv4 {
                primary {
                    address 172.16.14.2
                    prefix-length 30
                }
            }
        }
        interface "int-CE-4-PE-2_VPRN1" {
            port 1/1/2:1
            ipv4 {
                primary {
                    address 172.16.24.2
                    prefix-length 30
                }
            }
        }
        interface "system" {
            ipv4 {
                primary {
                    address 192.0.2.4
                    prefix-length 32
                }
            }
        }
        interface "test_connectedNW" {
            loopback
            ipv4 {
                primary {
                    address 10.0.0.1
                    prefix-length 8
                }
            }
        }
        bgp {
            rapid-withdrawal true
            split-horizon true
            group "EBGP" {
                peer-as 64496
                export {
                    policy ["export-bgp"]
                }
            }
            neighbor "172.16.14.1" {
                group "EBGP"
            }
            neighbor "172.16.24.1" {
                group "EBGP"
            }
        }
```

Initially, VPRN BGP best external is disabled and, so only the best BGP route will be advertised and IBGP peers will not learn backup paths. The following section shows which routes are exchanged. Afterward, VPRN BGP best external will be enabled and the same show commands will be used.

Verification - VPRN without BGP best external

PE-1 imports prefix 10.0.0.0/8, assigns LP 200 to it, and advertises a corresponding VPN-IPv4 route to its IBGP peers (PE-2 and PE-3). Toward PE-2, this is as follows:

```
# on PE-1:
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
15 2022/05/02 11:24:40.478 CEST MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 65
    Flag: 0x90 Type: 14 Len: 30 Multiprotocol Reachable NLRI:
        Address Family VPN_IPV4
        NextHop len 12 NextHop 192.0.2.1
        10.0.0.0/8 RD 64496:11 Label 524283 (Raw label 0x7fffb1)
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 6 AS Path:
        Type: 2 Len: 1 < 64500 >
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 200
    Flag: 0xc0 Type: 16 Len: 8 Extended Community:
        target:64496:1
"
```

The NLRI includes the prefix 10.0.0.0/8 and the RD 64496:11, and the label is 524283. BGP prevents PE-2 from sending a similar BGP update for prefix 10.0.0.0/8 because that route is not active on PE-2. PE-3 receives a BGP VPN-IPv4 route for network 64496:11:10.0.0.0/8, and this route has PE-1 as next hop and LP 200. No route is received from PE-2 for network 64496:12:10.0.0.0/8; as follows:

```
[/]
A:admin@PE-3# show router bgp routes vpn-ipv4
===============================================================================
 BGP Router ID:192.0.2.3        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv4 Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop (Router)                                 Path-Id     IGP Cost
      As-Path                                                      Label
-------------------------------------------------------------------------------
u*>i  64496:11:10.0.0.0/8                              200         None
      192.0.2.1                                        None        10
      64500                                                        524283
u*>i  64496:11:172.16.14.0/30                          100         None
      192.0.2.1                                        None        10
      No As-Path                                                   524283
u*>i  64496:11:172.31.2.1/32                           100         None
      192.0.2.1                                        None        10
      No As-Path                                                   524283
u*>i  64496:12:172.16.24.0/30                          100         None
      192.0.2.2                                        None        10
      No As-Path                                                   524283
u*>i  64496:12:172.31.2.2/32                           100         None
      192.0.2.2                                        None        10
      No As-Path                                                   524283
-------------------------------------------------------------------------------
Routes : 5
===============================================================================
```

In a similar way, the list of BGP VPN-IPv4 routes on PE-2 includes prefix 64496:11:10.0.0.0/8 with LP 200 and next hop PE-1, as follows:

```
[/]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
A:admin@PE-2# show router bgp routes vpn-ipv4
===============================================================================
 BGP Router ID:192.0.2.2        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv4 Routes
===============================================================================
Flag  Network                                           LocalPref   MED
      Nexthop (Router)                                  Path-Id     IGP Cost
      As-Path                                                       Label
-------------------------------------------------------------------------------
u*>i  64496:11:10.0.0.0/8                               200         None
      192.0.2.1                                         None        10
      64500                                                         524283
u*>i  64496:11:172.16.14.0/30                           100         None
      192.0.2.1                                         None        10
      No As-Path                                                    524283
u*>i  64496:11:172.31.2.1/32                            100         None
      192.0.2.1                                         None        10
      No As-Path                                                    524283
u*>i  64496:13:172.31.2.3/32                            100         None
      192.0.2.3                                         None        10
      No As-Path                                                    524283
-------------------------------------------------------------------------------
Routes : 4
===============================================================================
```

The list of BGP IPv4 routes in VPRN 1 on PE-2 has two entries for prefix 10.0.0.0/8, but none of them is best or used, as follows:

```
[/]
A:admin@PE-2# show router 1 bgp routes
===============================================================================
 BGP Router ID:172.31.2.2       AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP IPv4 Routes
===============================================================================
Flag  Network                                           LocalPref   MED
      Nexthop (Router)                                  Path-Id     IGP Cost
      As-Path                                                       Label
-------------------------------------------------------------------------------
*i    10.0.0.0/8                                        100         None
      172.16.24.2                                       None        0
      64500                                                         -
i     10.0.0.0/8                                        200         None
      172.16.14.2                                       None        0
      64500                                                         -
-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The routing table for VPRN 1 on PE-2 and PE-3 for prefix 10.0.0.0/8 shows that the next hop is PE-1 and the protocol is BGP VPN, as follows:

```
[/]
A:admin@PE-2# show router 1 route-table 10.0.0.0/8

===============================================================================
Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
10.0.0.0/8                                    Remote  BGP VPN   00h02m00s  170
       192.0.2.1 (tunneled)                                        10
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

Instead of using an external route to CE-4, the route for prefix 10.0.0.0/8 is internal (BGP VPN), using an LDP transport tunnel to PE-1. There are no non-active routes, as can be shown by adding the keyword all to the preceding show command, as follows:

```
[/]
A:admin@PE-2# show router 1 route-table 10.0.0.0/8 all

===============================================================================
Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                           Active     Metric
-------------------------------------------------------------------------------
10.0.0.0/8                                    Remote  BGP VPN   00h02m23s  170
       192.0.2.1 (tunneled)                                Y           10
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
       E = Inactive best-external BGP route
===============================================================================
```

There are no standby routes, because BGP only advertises the best used route.

On PE-1, the following BGP IPv4 route with next hop CE-4 is used for prefix 10.0.0.0/8 in VPRN 1:

```
[/]
A:admin@PE-1# show router 1 bgp routes
===============================================================================
 BGP Router ID:172.31.2.1       AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP IPv4 Routes
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop (Router)                                 Path-Id     IGP Cost
      As-Path                                                      Label
-------------------------------------------------------------------------------
u*>i  10.0.0.0/8                                       200         None
      172.16.14.2                                      None        0
      64500                                                        -
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

The route for prefix 10.0.0.0/8 in the routing table of VPRN 1 has next hop 172.16.14.2 on CE-4, as follows:

```
[/]
A:admin@PE-1# show router 1 route-table 10.0.0.0/8 all

===============================================================================
Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                             Type    Proto     Age        Pref
      Next Hop[Interface Name]                          Active     Metric
-------------------------------------------------------------------------------
10.0.0.0/8                                     Remote  BGP       00h03m20s  170
       172.16.14.2                                      Y            0
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
       E = Inactive best-external BGP route
===============================================================================
```

There is no backup route because BGP prevents PE-2 from sending a standby route for prefix 10.0.0.0/8 to its IBGP peers.

PE-2 has advertised two VPN-IPv4 routes in the base router (the last number in Rcv/Act/Sent = Received/Active/Sent), as follows:

```
*A:PE-2# show router bgp summary family vpn-ipv4
===============================================================================
 BGP Router ID:192.0.2.2        AS:64496       Local AS:64496
===============================================================================
BGP Admin State         : Up          BGP Oper State              : Up
---snip---

===============================================================================
BGP VPN-IPv4 Summary
===============================================================================
Legend : D - Dynamic Neighbor
===============================================================================
Neighbor
                  AS PktRcvd PktSent  InQ OutQ Up/Down   State|Recv/Actv/Sent
-------------------------------------------------------------------------------
192.0.2.1
           64496           24      23    0    0 00h08m42s 3/3/2
192.0.2.3
           64496           22      23    0    0 00h08m37s 1/1/2
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Enable BGP best external in VPRN

VPRN BGP best external is configured on PE-2 (or on all PEs in the multi-homing site) as follows:

```
# on PE-2:
configure {
    service {
        vprn "VPRN 1" {
            export-inactive-bgp true
```

When configured, this command causes all IPv4 and IPv6 VPRN BGP best external routes to be exported in the multi-protocol BGP (MP-BGP) domain. Best external routes are BGP routes for which all the following conditions are met:

- The BGP route is matched by the VRF export policy.
- The BGP route is inactive because a more preferred BGP VPN route for the same prefix is present in the route table manager (RTM).
- This BGP route is best and valid considering only VPRN BGP routes.

PE-2 is advertising a best external route and is called the standby PE for prefix 10.0.0.0/8. PEs can be active for some IP prefixes and standby for other IP prefixes.

Best external routes are advertised to the BGP VPN-IPv4 neighbors. In this example, the BGP VPN-IPv4 neighbors are IBGP neighbors, but they can also be EBGP neighbors. The RD must be unique across the PEs exporting a BGP VPN-IP route for the same prefix; otherwise, the best external route may not be advertised. The advertised VPRN label is based on the next hop IP of the best external route, regardless of the label mode of the VPRN in the standby PE.

Verification - VPRN with BGP best external - BGP FRR

VPRN with BGP best external BGP FRR results in the following. VPRN BGP best external is enabled (BGP Export Inactv) in VPRN 1 on PE-2:

```
[/]
A:admin@PE-2# show service id 1 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 1                   Vpn Id            : 0
Service Type      : VPRN
MACSec enabled    : no
Name              : VPRN 1
Description       : (Not Specified)
Customer Id       : 1                   Creation Origin   : manual
---snip---

Max IPv6 Routes   : No Limit
Ignore NH Metric  : Disabled
Hash Label        : Disabled
Entropy Label     : Disabled
Vrf Target        : target:64496:1
---snip---

Label mode        : vrf
BGP VPN Backup    : ipv4
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
BGP Export Inactv : Enabled
LOG all events    : Disabled

SAP Count         : 1                   SDP Bind Count    : 0
VSD Domain        : <none>

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                               Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/3:1                              q-tag        1578    1578    Up   Up
===============================================================================
```

After VPRN BGP best external is enabled, PE-2 advertises its standby route for prefix 10.0.0.0/8 to its IBGP peers, as follows:

```
# on PE-2:
15 2022/05/02 11:28:51.976 CEST MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Send BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 65
    Flag: 0x90 Type: 14 Len: 30 Multiprotocol Reachable NLRI:
        Address Family VPN_IPV4
        NextHop len 12 NextHop 192.0.2.2
        10.0.0.0/8 RD 64496:12 Label 524284 (Raw label 0x7fffc1)
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 6 AS Path:
        Type: 2 Len: 1 < 64500 >
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 8 Extended Community:
        target:64496:1
"
```

The RD is 64496:12, the LP is 100, and the label is 524284. The BGP update shown is sent by PE-2 toward PE-3; the BGP update sent by PE-2 toward PE-1 is similar.

The number of BGP VPN-IPv4 routes sent by PE-2 to each IBGP peer increased from 2 to 3, as follows:

```
[/]
A:admin@PE-2# show router bgp summary all

===============================================================================
BGP Summary
===============================================================================
Legend : D - Dynamic Neighbor
===============================================================================
Neighbor
Description
ServiceId          AS PktRcvd InQ  Up/Down   State|Rcv/Act/Sent (Addr Family)
                      PktSent OutQ
-------------------------------------------------------------------------------
192.0.2.1
Def. Inst       64496      24    0 00h08m42s 3/3/3 (VpnIPv4)
                           23    0
192.0.2.3
Def. Inst       64496      22    0 00h08m37s 1/1/3 (VpnIPv4)
                           23    0
---snip---
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

PE-3 has two BGP VPN-IPv4 routes for prefix 10.0.0.0/8: one for network 64496:11:10.0.0.0/8 with LP 200 and next hop PE-1, and one for network 64496:12:10.0.0.0/8 with LP 100 and next hop PE-2, as follows:

```
[/]
A:admin@PE-3# show router bgp routes 10.0.0.0/8 vpn-ipv4
===============================================================================
 BGP Router ID:192.0.2.3        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv4 Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop (Router)                                 Path-Id     IGP Cost
      As-Path                                                      Label
-------------------------------------------------------------------------------
u*>i  64496:11:10.0.0.0/8                              200         None
      192.0.2.1                                        None        10
      64500                                                        524283
u*>i  64496:12:10.0.0.0/8                              100         None
      192.0.2.2                                        None        10
      64500                                                        524284
-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

PE-1 has one BGP VPN-IPv4 route for network 64496:12:10.0.0.0/8 with LP 100 and next hop PE-2; PE-2 has one BGP VPN-IPv4 route for network 64496:11:10.0.0.0/8 with LP 200 and next hop PE-1.

All PEs are ready for BGP FRR and the "B" flag indicates that a BGP VPN-IPv4 backup route is available. This flag is present when the VPRN is configured for BGP FRR (**bgp-vpn-backup>ipv4 true**) and a standby route has been received, as follows. The B flag was not present in the output for the routing table when VPRN BGP best external was disabled, as shown earlier.

```
[/]
A:admin@PE-1# show router 1 route-table 10.0.0.0/8

===============================================================================
Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
10.0.0.0/8 [B]                                Remote  BGP       00h06m47s  170
       172.16.14.2                                                 0
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

The active route on PE-1 has next hop 172.16.14.2 on CE-4.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

On PE-3, the active BGP VPN-IPv4 route for prefix 10.0.0.0/8 uses an LDP transport tunnel to PE-1; a BGP VPN-IPv4 backup route is also available, as follows:

```
[/]
A:admin@PE-3# show router 1 route-table 10.0.0.0/8

===============================================================================
Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
10.0.0.0/8 [B]                                  Remote  BGP VPN   00h06m39s  170
       192.0.2.1 (tunneled)                                         10
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

The active BGP VPN-IPv4 route on PE-2 uses an LDP transport tunnel to PE-1, but no BGP backup route is available:

```
[/]
A:admin@PE-2# show router 1 route-table 10.0.0.0/8

===============================================================================
Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
10.0.0.0/8                                      Remote  BGP VPN   00h06m34s  170
       192.0.2.1 (tunneled)                                         10
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

PE-2 has a standby BGP IPv4 route that is displayed with the following show command:

```
[/]
A:admin@PE-2# show router 1 route-table 10.0.0.0/8 all

===============================================================================
Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                           Active     Metric
-------------------------------------------------------------------------------
10.0.0.0/8 [E]                                  Remote  BGP       00h03m11s  170
       172.16.24.2                                        N            0
10.0.0.0/8                                      Remote  BGP VPN   00h07m14s  170
       192.0.2.1 (tunneled)                               Y            10
-------------------------------------------------------------------------------
No. of Routes: 2
Flags: n = Number of times nexthop is repeated
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
       E = Inactive best-external BGP route
===============================================================================
```

The "E" flag indicates that this route is an inactive best external BGP route.

VPRN 1 on PE-1 and PE-3 is ready for BGP FRR (**bgp-vpn-backup>ipv4 true**) and PE-2 advertised a standby BGP VPN-IPv4 route for prefix 10.0.0.0/8; therefore, PE-1 and PE-3 can add an alternative route to the routing table of VPRN 1, as follows:

```
[/]
A:admin@PE-1# show router 1 route-table 10.0.0.0/8 alternative

===============================================================================
Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
      Alt-NextHop                                                 Alt-
                                                                  Metric
-------------------------------------------------------------------------------
10.0.0.0/8                                    Remote  BGP       00h07m47s  170
       172.16.14.2                                                   0
10.0.0.0/8 (Backup)                           Remote  BGP VPN   00h07m47s  170
       192.0.2.2 (tunneled)                                       10
-------------------------------------------------------------------------------
No. of Routes: 2
Flags: n = Number of times nexthop is repeated
       Backup = BGP backup route
       LFA = Loop-Free Alternate nexthop
       S = Sticky ECMP requested
===============================================================================
```

```
[/]
A:admin@PE-3# show router 1 route-table 10.0.0.0/8 alternative

===============================================================================
Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
      Alt-NextHop                                                 Alt-
                                                                  Metric
-------------------------------------------------------------------------------
10.0.0.0/8                                    Remote  BGP VPN   00h08m05s  170
       192.0.2.1 (tunneled)                                       10
10.0.0.0/8 (Backup)                           Remote  BGP VPN   00h08m05s  170
       192.0.2.2 (tunneled)                                       10
-------------------------------------------------------------------------------
No. of Routes: 2
Flags: n = Number of times nexthop is repeated
       Backup = BGP backup route
       LFA = Loop-Free Alternate nexthop
       S = Sticky ECMP requested
===============================================================================
```

The alternative BGP VPN-IPv4 route for prefix 10.0.0.0/8 in VPRN 1 uses an LDP transport tunnel toward PE-2.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Configure ECMP

Because BGP best external allows advertising of an alternative path, it can also be used for load-sharing. ECMP is configured with value 2 in VPRN 1 on all PEs, as follows:

```
# on PE-1, PE-2, PE-3:
configure {
    service {
        vprn "VPRN 1" {
            ecmp 2
```

Other than the ECMP configuration, the VPRN configuration is the same as in the previous example. If ECMP is configured, BGP FRR is not needed anymore:

```
# on PE-1, PE-2, PE-3:
configure {
    service {
        vprn "VPRN 1" {
            delete bgp-vpn-backup
```

On PE-3, the BGP decision process will prefer the route with the highest LP and, therefore, only the route via PE-1 with LP 200 will be used and there will be no load-sharing. To ensure that the routes via PE-1 and PE-2 have the same cost, the import policy in VPRN 1 on PE-1 that sets the LP to 200 is replaced by an import policy that sets the LP to 100, as follows:

```
# on PE-1:
configure {
    policy-options {
        policy-statement "import-bgp-LP100" {
            default-action {
                action-type accept
                local-preference 100
            }
        }
    }
    service {
        vprn "VPRN 1" {
            bgp {
                group "EBGP"
                    delete import
                    import {
                        policy ["import-bgp-LP100"]
                    }
                }
```

BGP best external is enabled (on PE-1 and) PE-2, as follows:

```
# on PE-2:
configure {
    service {
        vprn "VPRN 1" {
            export-inactive-bgp true
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Verification - VPRN with BGP best external - ECMP

VPRN with BGP best external ECMP results in the following. With BGP best external enabled on the PEs in the multi-homing site (PE-2 and PE-3), the following two BGP VPN-IPv4 routes are used on PE-3:

```
[/]
A:admin@PE-3# show router bgp routes 10.0.0.0/8 vpn-ipv4
===============================================================================
 BGP Router ID:192.0.2.3        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv4 Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop (Router)                                 Path-Id     IGP Cost
      As-Path                                                      Label
-------------------------------------------------------------------------------
u*>i  64496:11:10.0.0.0/8                              100         None
      192.0.2.1                                        None        10
      64500                                                        524283
u*>i  64496:12:10.0.0.0/8                              100         None
      192.0.2.2                                        None        10
      64500                                                        524283
-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

The following BGP IPv4 routes are learned in VPRN 1 on PE-3, but they are not used:

```
[/]
A:admin@PE-3# show router 1 bgp routes 10.0.0.0/8
===============================================================================
 BGP Router ID:172.31.2.3       AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP IPv4 Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop (Router)                                 Path-Id     IGP Cost
      As-Path                                                      Label
-------------------------------------------------------------------------------
i     10.0.0.0/8                                       100         None
      172.16.14.2                                      None        0
      64500                                                        -
i     10.0.0.0/8                                       100         None
      172.16.24.2                                      None        0
      64500                                                        -
-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

When ECMP is enabled and the routes have the same LP, the routing table on PE-3 has two active routes for prefix 10.0.0.0/8, each using an LDP transport tunnel, as follows:

```
[/]
A:admin@PE-3# show router 1 route-table 10.0.0.0/8

===============================================================================
Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
10.0.0.0/8                                    Remote  BGP VPN   00h02m42s  170
       192.0.2.1 (tunneled)                                         10
10.0.0.0/8                                    Remote  BGP VPN   00h02m42s  170
       192.0.2.2 (tunneled)                                         10
-------------------------------------------------------------------------------
No. of Routes: 2
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

Figure 301: Loadsharing for traffic from PE-3 destined to 10.0.0.0/8 shows that traffic from VPRN 1 on PE-3 destined to prefix 10.0.0.0/8 is sprayed over two paths: one via PE-1 and one via PE-2.

Figure 301: Loadsharing for traffic from PE-3 destined to 10.0.0.0/8

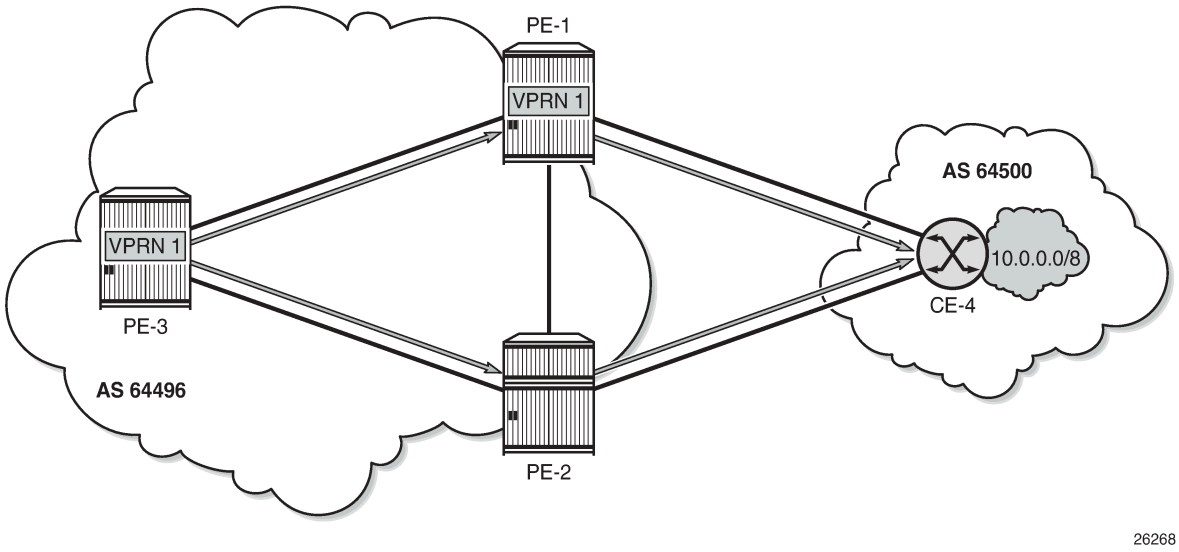

Conclusion

VPRNs can be configured with the option **export-inactive-bgp true**, which allows a BGP speaker to advertise its best external BGP route to its BGP peers even if that route is inactive due to the presence of a more preferred BGP VPN route from another PE. BGP best external in VPRN is useful in active/standby multi-homing scenarios because it allows the standby PE to advertise a backup path. The traffic failover

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

time can be reduced when all PE routers have advance knowledge of the potential backup paths and do not have to wait for BGP route advertisements and/or withdrawals to reprogram their forwarding tables. VPRN BGP best external can also be used in combination with ECMP.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Carrier Supporting Carrier IP VPNs

This chapter provides information about carrier supporting carrier IP VPN configurations.

Topics in this chapter include:

Applicability

This chapter was initially written for on SR OS Release 11.0.R1. The MD-CLI in the current edition corresponds to SR OS Release 22.2.R1. Carrier Supporting Carrier is supported on the 7750 SR and 7950 XRS.

Overview

Carrier Supporting Carrier (CSC) is a solution that allows one service provider (the Customer Carrier) to use the IP VPN service of another service provider (the Super Carrier) for some or all of its backbone transport. RFC 4364 defines a Carrier Supporting Carrier solution for BGP/MPLS IP VPNs that uses MPLS at the interconnection points between the two service providers to provide a scalable and secure solution.

A simplified CSC network topology is shown in Figure 302: CSC network topology. A CSC deployment involves the following types of devices:

- CE — Customer premises equipment dedicated to one enterprise.
- PE — Edge router managed and operated by the Customer Carrier that connects to CEs to provide business VPN or Internet services.
- CSC-CE — Peering router managed and operated by the Customer Carrier that is connected to CSC-PEs for purposes of using the associated CSC IP VPN services for backbone transport. The CSC-CE may attach directly to CEs if it is also configured to be a PE for business VPN services.
- CSC-PE — A PE router managed and operated by the Super Carrier that supports one or more CSC IP VPN services possibly in addition to other traditional PE services.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 302: CSC network topology

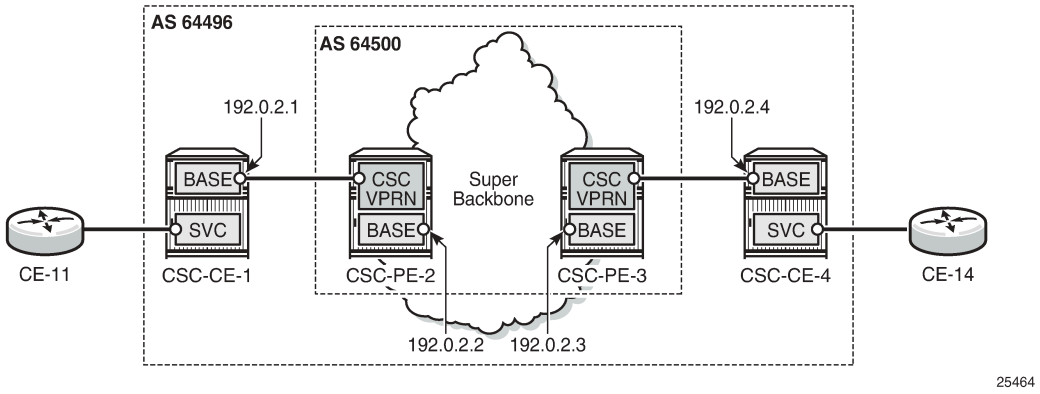

In the CSC solution, the CSC-CE and CSC-PE are directly connected by a link that supports MPLS. The CSC-CE distributes an MPLS label for every /32 IPv4 prefix it and any downstream PE uses as a BGP next-hop in routes associated with services offered by the Customer Carrier. BGP must be used as the label distribution protocol between CSC-CE and CSC-PE if the latter device is an SR OS node. Typically, the Customer Carrier and Super Carrier operate as two different Autonomous Systems (ASs) and therefore BGP, more specifically EBGP, is the best label distribution protocol, even if other options are available. The BGP session between CSC-CE and CSC-PE must be single-hop EBGP (or IBGP) if either device is an SR OS node.

In an SR OS CSC-PE, the interface to a CSC-CE is a special type of IP/MPLS interface that belongs to a VPRN configured for CSC mode. This special type of interface is called a CSC VPRN interface throughout the remainder of this chapter. The CSC VPRN interface has many of the same characteristics as a network interface of the base router but its association with a Virtual Routing and Forwarding (VRF) ensures that the traffic and control plane routes of the Customer Carrier are kept separate from other services.

When an SR OS CSC-PE receives a labeled-IPv4 route (with label L1, next-hop N1) from a CSC-CE BGP peer, the following actions take place in the CSC-PE:

1. The BGP route is installed into the routing table of the CSC VPRN (assuming the BGP route is the best route to the destination).
2. If the BGP route matches the VRF export policy, it is advertised to the core Multi-Protocol Border Gateway Protocol (MP-BGP) peers as a VPN-IPv4 route. The advertised label value is changed to label value L2.
3. BGP programs the line cards with an MPLS forwarding entry that swaps label value L2 for L1 and sends the MPLS packet over the CSC VPRN interface associated with next-hop N1.

When an SR OS CSC-PE receives a VPN-IPv4 route (with label L2, next-hop N2) the following actions take place in the CSC-PE:

1. If the VPN-IPv4 route matches the VRF import policy of a CSC VPRN, it is installed into the routing table of that CSC VPRN.
2. If the imported BGP-VPN route matches the BGP export policy associated with a CSC-CE BGP peer, it is advertised to that peer as a labeled-IPv4 route. The advertised label value is changed to label value L3.
3. BGP programs the line cards with an MPLS forwarding entry that swaps label value L3 for L2 and sends the packet inside the MPLS tunnel to next-hop N2.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Once a CSC-CE has learned a labeled-IPv4 route for a remote CSC-CE and vice versa, the two CSC-CEs can set up a BGP session between themselves and exchange VPN routes over this session if they are both PEs with services. Typically, this BGP session will be an IBGP session because the local and remote CSC-CEs belong to the same AS. The Layer 2 VPN and Layer 3 VPN routes exchanged by the CSC-CEs are resolved by the labeled-IPv4 routes they have for each other's /32 IPv4 address.

Configuration

This section will walk through the steps to configure the CSC solution shown in Figure 302: CSC network topology. The IPv4 addresses in Figure 302: CSC network topology are the system IP addresses of the routers. The steps are the following:

- Configure CSC-CE-1
- Configure CSC service on CSC-PE-2
- Verify exchange of labeled IPv4 routes between CSC-CE-1 and CSC-PE-2
- Configure core connectivity for CSC-PE-2
- Configure core connectivity for CSC-PE-3
- Configure CSC service on CSC-PE-3
- Verify exchange of VPN-IPv4 routes between CSC-PE-2 and CSC-PE-3
- Configure CSC-CE-4
- Verify exchange of labeled IPv4 routes between CSC-PE-3 and CSC-CE-4
- Configure BGP session between CSC-CE-1 and CSC-CE-4
- Verify exchange of VPN-IPv4 routes between CSC-CE-1 and CSC-CE-4

Step 1. Configure CSC-CE-1

This example assumes that CSC-CE-1 is a PE router with Layer 2 and Layer 3 VPN services that must extend across the CSC VPN service; assume that there are no further downstream PEs in AS 64496. The configuration of one such Layer 3 VPN service in CSC-CE-1 is as follows:

```
# on CSC-CE-1:
configure {
    service {
        vprn "VPRN1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:11"
                    vrf-target {
                        community "target:64496:1"
                    }
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            interface "loopback-1" {
                loopback true
                ipv4 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    primary {
                        address 10.11.30.2
                        prefix-length 24
                    }
                }
            }
        }
```

For brevity, the preceding configuration sample omits commands related to SAP IP interfaces, spoke-SDP IP interfaces, PE-CE routing protocols, QoS, IP filters, and so on. The loopback interface is used to test whether this prefix is learned at the remote CSC-CE-4.

The base routing instance of the CSC-CE is configured with the appropriate router ID and autonomous system number and the system interface is configured with an IPv4 address (usually the same as the router ID). If the router ID is not configured, by default, the system IP address is used as the router ID. The interface to CSC-PE-2 is created and configured. The base router configuration of CSC-CE-1 is as follows:

```
# on CSC-CE-1:
configure {
    router "Base" {
        autonomous-system 64496
        interface "int-CSC-CE-1-CSC-PE-2" {
            port 1/1/1:1        # connected to VPRN1 network interface on CSC-PE-2
            ipv4 {
                primary {
                    address 192.168.12.1
                    prefix-length 30
                }
            }
        }
        interface "system" {
            ipv4 {
                primary {
                    address 192.0.2.1
                    prefix-length 32
                }
            }
        }
```

On CSC-CE-1, BGP is configured as the control plane protocol running on the interface to CSC-PE-2. The export policy exports the system IP address of CSC-CE-1 as a labeled-IPv4 route to CSC-PE-2; the import policy imports other system IP addresses, in this case, the system IP address of CSC-CE-4.

```
# on CSC-CE-1:
configure {
    policy-options {
        prefix-list "system-IP" {
            prefix 192.0.2.0/29 type longer {
            }
        }
        policy-statement "export-systemIP" {
            entry 10 {
                from {
                    prefix-list ["system-IP"]
                    protocol {
                        name [direct]
                    }
                }
                action {
                    action-type accept
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
            default-action {
                action-type reject
            }
        }
        policy-statement "import-systemIP" {
            entry 10 {
                from {
                    prefix-list ["system-IP"]
                }
                action {
                    action-type accept
                }
            }
            default-action {
                action-type reject
            }
        }
    }
    router "Base" {
        autonomous-system 64496
        bgp {
            group "CSC-PE" {
                peer-as 64500
            }
            neighbor "192.168.12.2" {
                split-horizon true
                group "CSC-PE"
                family {
                    label-ipv4 true
                }
                import {
                    policy ["import-systemIP"]
                }
                export {
                    policy ["export-systemIP"]
                }
            }
```

The peer type is EBGP (**peer-as** is different from the locally configured **autonomous-system**)

The address family for the EBGP session is **label-ipv4** (the **neighbor** address is an IPv4 address). Family label-IPv4 causes MP-BGP negotiation of the address family for AFI=1 and SAFI=4 (IPv4 NLRI with MPLS labels), as can be observed from the following debug message (in this example, debugging is enabled on CSC-CE-1 for BGP OPEN messages using the command **debug router bgp open**).This BGP OPEN message can obviously only be seen when the BGP peer is up. The configuration for CSC-PE-2 will be shown later, but in order to have the debug message, it must be configured already.

```
# on CSC-CE-1:
1 2022/04/06 15:54:53.648 CEST MINOR: DEBUG #2001 Base BGP
"BGP: OPEN
Peer 1: 192.168.12.2 - Send (Passive) BGP OPEN: Version 4
   AS Num 64496: Holdtime 90: BGP_ID 192.0.2.1: Opt Length 20 (ExtOpt F)
   Opt Para: Type CAPABILITY: Length = 18: Data:
     Cap_Code GRACEFUL-RESTART: Length 2
       Bytes: 0x0 0x78
     Cap_Code MP-BGP: Length 4
       Bytes: 0x0 0x1 0x0 0x4
     Cap_Code ROUTE-REFRESH: Length 0
     Cap_Code 4-OCTET-ASN: Length 4
       Bytes: 0x0 0x0 0xfb 0xf0
"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The **split-horizon true** command is optional. It prevents a best BGP route from the CSC-PE peer from being re-advertised back to that peer.

Step 2. Configure CSC service on SCS-PE-2

CSC-PE-2 must be configured with a VPRN in **carrier-carrier-vpn** mode to provide CSC service to CSC-CE-1. VPRN 1 is configured on CSC-PE-2, as follows:

```
# on CSC-PE-2:
configure {
    service {
        vprn "VPRN1" {
            admin-state enable
            service-id 1
            customer "1"
            autonomous-system 64500
            router-id 192.0.2.2
            carrier-carrier-vpn true
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64500:12"
                    vrf-target {
                        community "target:64500:1"
                    }
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            bgp {
                group "CSC-CE" {
                    as-override true
                    peer-as 64496
                    ebgp-default-reject-policy {
                        import false
                    }
                    export {
                        policy ["BGP-VPN-routes"]
                    }
                }
                neighbor "192.168.12.1" {
                    split-horizon true
                    group "CSC-CE"
                    family {
                        label-ipv4 true
                    }
                }
            }
            network-interface "int-CSC-PE-2-CSC-CE-1" {
                port 1/1/2:1
                ipv4 {
                    primary {
                        address 192.168.12.2
                        prefix-length 30
                    }
                }
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The **carrier-carrier-vpn true** command is mandatory. It cannot be configured if the VPRN has any SAP or spoke-SDP access interfaces configured; they must first be deleted.

```
*[ex:/configure service vprn "VPRN1"]
A:admin@CSC-PE-2# carrier-carrier-vpn true

*[ex:/configure service vprn "VPRN1"]
A:admin@CSC-PE-2# commit
MINOR: COMMON #238: configure service vprn "VPRN1" carrier-carrier-vpn - Configuration change
 failed validation - combination of carrier-carrier-vpn and service interfaces present. -
 configure service vprn "VPRN1"
```

The **auto-bind-tunnel** command must be set appropriately for the type of transport desired to other CSC-PEs, but note that GRE is not supported.

```
*[ex:/configure service vprn "VPRN1" bgp-ipvpn mpls auto-bind-tunnel resolution-filter]
A:admin@CSC-PE-2# gre true

*[ex:/configure service vprn "VPRN1" bgp-ipvpn mpls auto-bind-tunnel resolution-filter]
A:admin@CSC-PE-2# commit
INFO: PIP #1195: configure service vprn "VPRN1" carrier-carrier-vpn - Cannot toggle carrier-
carrier-vpn - GRE auto-bind enabled - configure service vprn "VPRN1" bgp-ipvpn mpls auto-bind-
tunnel resolution-filter gre
```

The interface to CSC-CE-1 must be a network interface. A network interface can be associated with an entire Ethernet port, a VLAN sub-interface of an Ethernet port, an entire LAG or a VLAN sub-interface of a LAG. In all cases, the associated Ethernet ports must be configured in network or hybrid mode.

The peer type is EBGP (**peer-as** is different from the locally configured **autonomous-system**).

The address family for the EBGP session is **label-ipv4** (the **neighbor** address is an IPv4 address). Address family label-ipv4 causes MP-BGP negotiation of the address family for AFI=1 and SAFI=4 (IPv4 NLRI with MPLS labels).

The **split-horizon true** command is optional. It prevents a best BGP route from the CSC-CE peer from being re-advertised back to that peer.

The **as-override** command replaces CSC-CE-1's AS number 64496 with CSC-PE-2's AS number 64500 in the AS_PATH attribute of routes advertised to CSC-CE-1. Without this configuration, CSC-CE-1 may reject routes originated by CSC-CE-4 as invalid due to an AS-path loop.

The **export** command applies a BGP export policy to the session. The configuration of the policy is as follows:

```
# on CSC-PE-2:
configure {
    policy-options {
        policy-statement "BGP-VPN-routes" {
            entry 10 {
                from {
                    protocol {
                        name [bgp-vpn]
                    }
                }
                action {
                    action-type accept
                }
            }
            default-action {
                action-type reject
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
```

The effect of the BGP export policy is to re-advertise VPN-IPv4 routes imported into the CSC VPRN (and used for forwarding) to CSC-CE-4.

Step 3. Verify exchange of labeled IPv4 routes

When steps 1 and 2 have been completed, CSC-CE-1 advertises the labeled-IPv4 route for its system IP address 192.0.2.1/32 to CSC-PE-2. This can be checked in the RIB Out of CSC-CE-1, as follows:

```
[/]
A:admin@CSC-CE-1# show router bgp routes 192.0.2.1/32 label-ipv4 hunt
===============================================================================
 BGP Router ID:192.0.2.1        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP LABEL-IPV4 Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
Network        : 192.0.2.1/32
Nexthop        : 192.168.12.1
Path Id        : None
To             : 192.168.12.2
Res. Nexthop   : n/a
Local Pref.    : n/a                   Interface Name : NotAvailable
Aggregator AS  : None                  Aggregator     : None
Atomic Aggr.   : Not Atomic            MED            : None
AIGP Metric    : None                  IGP Cost       : n/a
Connector      : None
Community      : No Community Members
Cluster        : No Cluster Members
Originator Id  : None                  Peer Router Id : 192.0.2.2
IPv4 Label     : 524287                Label Type     : POP
Lbl Allocation : NEXT-HOP
Origin         : IGP
AS-Path        : 64496
Route Tag      : 0
Neighbor-AS    : 64496
Orig Validation: NotFound
Source Class   : 0                     Dest Class     : 0

-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

CSC-CE-1 has advertised a label value of 524287 with the prefix.

The following output shows the received label-IPv4 route in the RIB In for VPRN "VPRN1" on CSC-PE-2:

```
[/]
A:admin@CSC-PE-2# show router 1 bgp routes 192.0.2.1/32 label-ipv4 hunt
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
 BGP Router ID:192.0.2.2        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP LABEL-IPV4 Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : 192.0.2.1/32
Nexthop        : 192.168.12.1
Path Id        : None
From           : 192.168.12.1
Res. Nexthop   : 192.168.12.1
Local Pref.    : None                   Interface Name : int-CSC-PE-2-CSC-CE-1
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : 0
Connector      : None
Community      : No Community Members
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.1
Fwd Class      : None                   Priority       : None
IPv4 Label     : 524287
Flags          : Used Valid Best IGP In-TTM In-RTM
Route Source   : External
AS-Path        : 64496
Route Tag      : 0
Neighbor-AS    : 64496
Orig Validation: NotFound
Source Class   : 0                      Dest Class     : 0
Add Paths Send : Default
RIB Priority   : Normal
Last Modified  : 00h01m30s

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

Step 4. Configure core connectivity for CSC-PE-2

The next step is to configure the base router instance of CSC-PE-2 so that it can exchange VPN-IPv4 routes with CSC-PE-3 (and potentially other CSC-PEs). This requires:

- Router ID and autonomous system configuration.
- Network interface creation and configuration, including assignment of an IPv4 address to the system interface.
- Configuration of the IGP protocol; in this example, IS-IS is used.
- Configuration of the LDP protocol (optional).
- Configuration of RSVP LSPs used to reach remote CSC-PE devices (optional).
- Configuration of the BGP protocol.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The base router configuration of CSC-PE-2 is as follows:

```
# on CSC-PE-2:
configure {
    router "Base" {
        autonomous-system 64500
        interface "int-CSC-PE-2-CSC-PE-3" {
            port 1/1/1:1000
            ipv4 {
                primary {
                    address 192.168.23.1
                    prefix-length 30
                }
            }
        }
        interface "system" {
            ipv4 {
                primary {
                    address 192.0.2.2
                    prefix-length 32
                }
            }
        }
        bgp {
            group "core" {
                type internal
            }
            neighbor "192.0.2.3" {
                group "core"
                family {
                    vpn-ipv4 true
                }
            }
        }
        isis 0 {
            admin-state enable
            level-capability 2
            area-address [49.01]
            interface "int-CSC-PE-2-CSC-PE-3" {
                interface-type point-to-point
            }
            interface "system" {
                passive true
            }
            level 2 {
                wide-metrics-only true
            }
        }
        ldp {
            interface-parameters {
                interface "int-CSC-PE-2-CSC-PE-3" {
                    ipv4 {
                    }
                }
            }
        }
```

The peer type is IBGP (**type internal**. It is also possible to configure this in a similar way as for EBGP, with the same value for **peer-as** as the locally configured **autonomous-system**).

The transport for the IBGP session is IPv4 (the **neighbor** address is an IPv4 address).

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The **family vpn-ipv4** command causes MP-BGP negotiation of the address family for AFI=1 and SAFI=128 (=0x80), as can be observed from the following debug trace of the BGP OPEN message from CSC-PE-2 to CSC-PE-3.

```
1 2022/04/06 16:01:45.985 CEST MINOR: DEBUG #2001 Base BGP
"BGP: OPEN
Peer 1: 192.0.2.3 - Send (Active) BGP OPEN: Version 4
   AS Num 64500: Holdtime 90: BGP_ID 192.0.2.2: Opt Length 20 (ExtOpt F)
   Opt Para: Type CAPABILITY: Length = 18: Data:
     Cap_Code GRACEFUL-RESTART: Length 2
       Bytes: 0x0 0x78
     Cap_Code MP-BGP: Length 4
       Bytes: 0x0 0x1 0x0 0x80
     Cap_Code ROUTE-REFRESH: Length 0
     Cap_Code 4-OCTET-ASN: Length 4
       Bytes: 0x0 0x0 0xfb 0xf4
"
```

Step 5. Configure core connectivity for CSC-PE-3

The next step is to configure the base router instance of CSC-PE-3 so that it can exchange VPN-IPv4 routes with CSC-PE-2 and potentially other CSC-PEs. This requires:

- Router ID and AS configuration.
- Network interface creation and configuration, including assignment of an IPv4 address to the system interface.
- Configuration of the IGP protocol; in this example IS-IS is used.
- Configuration of the LDP protocol (optional).
- Configuration of RSVP LSPs used to reach remote CSC-PE devices (optional).
- Configuration of the BGP protocol.

The base router configuration of CSC-PE-3 is as follows:

```
# on CSC-PE-3:
configure {
    router "Base" {
        autonomous-system 64500
        interface "int-CSC-PE-3-CSC-PE-2" {
            port 1/1/2:1000
            ipv4 {
                primary {
                    address 192.168.23.2
                    prefix-length 30
                }
            }
        }
        interface "system" {
            ipv4 {
                primary {
                    address 192.0.2.3
                    prefix-length 32
                }
            }
        }
        bgp {
            group "core" {
                type internal
                cluster {
                    cluster-id 192.0.2.3
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
            }
            neighbor "192.0.2.2" {
                split-horizon true
                group "core"
                family {
                    vpn-ipv4 true
                }
            }
        }
        isis 0 {
            admin-state enable
            level-capability 2
            area-address [49.01]
            interface "int-CSC-PE-3-CSC-PE-2" {
                interface-type point-to-point
            }
            interface "system" {
                passive true
            }
            level 2 {
                wide-metrics-only true
            }
        }
        ldp {
            interface-parameters {
                interface "int-CSC-PE-3-CSC-PE-2" {
                    ipv4 {
                    }
                }
            }
        }
```

The peer type is IBGP (**type internal**. Can also be configured with **peer-as** equal to the locally configured **autonomous-system**).

The transport for the IBGP session is IPv4 (the **neighbor** address is an IPv4 address).

The **family vpn-ipv4** command causes MP-BGP negotiation of the address family for AFI=1 and SAFI=128.

The **cluster** command configures CSC-PE-2 as a route reflector for clients in the BGP group "core". This is not required and in a more typical deployment, the route reflector would be a separate router from any CSC-PE.

Step 6. Configure CSC service on CSC-PE-3

CSC-PE-3 must be configured with a VPRN in **carrier-carrier-vpn** mode to provide CSC service to CSC-CE-4. The configuration of the VPRN is as follows:

```
# on CSC-PE-3:
configure {
    service {
        vprn "VPRN1" {
            admin-state enable
            service-id 1
            customer "1"
            autonomous-system 64500
            router-id 192.0.2.3
            carrier-carrier-vpn true
            bgp-ipvpn {
                mpls {
                    admin-state enable
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    route-distinguisher "64500:13"
                    vrf-target {
                        community "target:64500:1"
                    }
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            bgp {
                group "CSC-CE" {
                    as-override true
                    peer-as 64496
                    ebgp-default-reject-policy {
                        import false
                    }
                    export {
                        policy ["BGP-VPN-routes"]
                    }
                }
                neighbor "192.168.34.2" {
                    split-horizon true
                    group "CSC-CE"
                    family {
                        label-ipv4 true
                    }
                }
            }
            network-interface "int-CSC-PE-3-CSC-CE-4" {
                port 1/1/1:1
                ipv4 {
                    primary {
                        address 192.168.34.1
                        prefix-length 30
                    }
                }
            }
        }
```

The **carrier-carrier-vpn true** command is mandatory. It cannot be configured if the VPRN has any SAP or spoke-SDP access interfaces configured; they must first be removed.

The **auto-bind-tunnel** command must be set appropriately for the type of transport desired to other CSC-PEs, but GRE is not supported.

The interface to CSC-CE-4 must be a network interface. A network interface can be associated with an entire Ethernet port, a VLAN sub-interface of an Ethernet port, an entire LAG or a VLAN sub-interface of a LAG. In all cases, the associated Ethernet ports must be configured in network or hybrid mode.

The peer type is EBGP (**peer-as** is different from the locally configured **autonomous-system**).

The address family for the EBGP session is **label-ipv4** (the **neighbor** address is an IPv4 address). Address family label-ipv4 causes MP-BGP negotiation of the address family for AFI=1 and SAFI=4 (IPv4 NLRI with MPLS labels).

The **split-horizon true** command is optional. It prevents a best BGP route from the CSC-CE peer from being re-advertised back to that peer.

The **as-override** command replaces CSC-CE-4's AS number 64496 with CSC-PE-3's AS number 64500 in the AS_PATH attribute of routes advertised to CSC-CE-4. Without this configuration, CSC-CE-4 may reject routes originated by CSC-CE-1 as invalid due to an AS-path loop.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The **export** command applies a BGP export policy to the session. The configuration of the policy is as follows:

```
# on CSC-PE-3:
configure {
    policy-options {
        policy-statement "BGP-VPN-routes" {
            entry 10 {
                from {
                    protocol {
                        name [bgp-vpn]
                    }
                }
                action {
                    action-type accept
                }
            }
            default-action {
                action-type reject
            }
        }
```

The effect of the BGP export policy is to re-advertise VPN-IPv4 routes imported into the CSC VPRN (and used for forwarding) to CSC-CE-4.

Step 7. Verify exchange of VPN-IPv4 routes between CSC-PE-2 and CSC-PE-3.

When the preceding steps have been completed, CSC-PE-2 advertises the labeled-IPv4 route for 192.0.2.1/32 (the system IP address of CSC-CE-1) to CSC-PE-3. This can be checked in the RIB Out of CSC-PE-2, as follows:

```
[/]
A:admin@CSC-PE-2# show router bgp routes 192.0.2.1/32 vpn-ipv4 hunt
===============================================================================
 BGP Router ID:192.0.2.2        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv4 Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
Network        : 192.0.2.1/32
Nexthop        : 192.0.2.2
Route Dist.    : 64500:12               VPN Label      : 524285
Path Id        : None
To             : 192.0.2.3
Res. Nexthop   : n/a
Local Pref.    : 100                    Interface Name : NotAvailable
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : n/a
Connector      : None
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Community      : target:64500:1
Cluster        : No Cluster Members
Originator Id  : None                  Peer Router Id : 192.0.2.3
Origin         : IGP
AS-Path        : 64496
Route Tag      : 0
Neighbor-AS    : 64496
Orig Validation: N/A
Source Class   : 0                     Dest Class     : 0

-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

CSC-PE-2 has advertised a VPN label value of 524285 with the prefix.

The following output shows the received route in the RIB In of CSC-PE-3:

```
[/]
A:admin@CSC-PE-3# show router bgp routes 192.0.2.1/32 vpn-ipv4 hunt
===============================================================================
 BGP Router ID:192.0.2.3        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv4 Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : 192.0.2.1/32
Nexthop        : 192.0.2.2
Route Dist.    : 64500:12              VPN Label      : 524285
Path Id        : None
From           : 192.0.2.2
Res. Nexthop   : n/a
Local Pref.    : 100                   Interface Name : int-CSC-PE-3-CSC-PE-2
Aggregator AS  : None                  Aggregator     : None
Atomic Aggr.   : Not Atomic            MED            : None
AIGP Metric    : None                  IGP Cost       : 10
Connector      : None
Community      : target:64500:1
Cluster        : No Cluster Members
Originator Id  : None                  Peer Router Id : 192.0.2.2
Fwd Class      : None                  Priority       : None
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : 64496
Route Tag      : 0
Neighbor-AS    : 64496
Orig Validation: N/A
Source Class   : 0                     Dest Class     : 0
Add Paths Send : Default
Last Modified  : 00h00m25s
VPRN Imported  :  1

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Routes : 1
===============================================================================
```

The label swap entries that BGP programmed in the line cards of CSC-PE-2 based on the received labeled-IPv4 route from CSC-CE-1 (Label Origin = ExtCarCarVpn) and the advertised VPN-IPv4 route to CSC-PE-3, as follows:

```
[/]
A:admin@CSC-PE-2# show router bgp inter-as-label

===============================================================================
BGP Inter-AS labels
Flags: B - entry has backup, P - entry is promoted
===============================================================================
NextHop                       Received       Advertised     Label
                              Label          Label          Origin
-------------------------------------------------------------------------------
192.168.12.1                  524287         524285         ExtCarCarVpn
-------------------------------------------------------------------------------
Total Labels allocated:   1
===============================================================================
```

Step 8. Configure CSC-CE-4

In this example, CSC-CE-4 is a PE router with Layer 2 and Layer 3 VPN services that must extend across the CSC VPN service. The configuration of one such Layer 3 VPN service in CSC-CE-4 is as follows:

```
# on CSC-CE-4:
configure {
    service {
        vprn "VPRN1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:14"
                    vrf-target {
                        community "target:64496:1"
                    }
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            interface "loopback-1" {
                loopback true
                ipv4 {
                    primary {
                        address 10.14.30.2
                        prefix-length 24
                    }
                }
            }
        }
```

For brevity, the preceding configuration sample omits commands related to SAP IP interfaces, spoke-SDP IP interfaces, PE-CE routing protocols, QoS, IP filters, and so on.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The base routing instance of CSC-CE-4 is configured with the appropriate router ID and AS number and the system interface has an IPv4 address (usually the same as the router ID). The interface to CSC-PE-3 is configured. The base router configuration of CSC-CE-4 is as follows:

```
# on CSC-CE-4:
configure {
    router "Base" {
        interface "int-CSC-CE-4-CSC-PE-3" {
            port 1/1/2:1        # connected to VPRN1 network interface on CSC-PE-3
            ipv4 {
                primary {
                    address 192.168.34.2
                    prefix-length 30
                }
            }
        }
        interface "system" {
            ipv4 {
                primary {
                    address 192.0.2.4
                    prefix-length 32
                }
            }
        }
```

BGP is configured as the control plane protocol running on the interface to CSC-PE-3, as follows:

```
# on CSC-CE-4:
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            group "CSC-PE" {
                peer-as 64500
            }
            neighbor "192.168.34.1" {
                split-horizon true
                group "CSC-PE"
                family {
                    label-ipv4 true
                }
                import {
                    policy ["import-systemIP"]
                }
                export {
                    policy ["export-systemIP"]
                }
            }
```

The peer type is EBGP (**peer-as** is different from the locally configured **autonomous-system**).

The address family for the EBGP session is **label-ipv4** (the **neighbor** address is an IPv4 address). Address family label-ipv4 causes MP-BGP negotiation of the address family for AFI=1 and SAFI=4 (IPv4 NLRI with MPLS labels).

The **split-horizon true** command is optional. It prevents a best BGP route from the CSC-PE peer from being re-advertised back to that peer.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The **export** and **import** commands apply BGP export and import policies to the session. The configuration of the policies is as follows:

```
# on CSC-CE-4:
configure {
    policy-options {
        prefix-list "system-IP" {
            prefix 192.0.2.0/29 type longer {
            }
        }
        policy-statement "export-systemIP" {
            entry 10 {
                from {
                    prefix-list ["system-IP"]
                    protocol {
                        name [direct]
                    }
                }
                action {
                    action-type accept
                }
            }
            default-action {
                action-type reject
            }
        }
        policy-statement "import-systemIP" {
            entry 10 {
                from {
                    prefix-list ["system-IP"]
                }
                action {
                    action-type accept
                }
            }
            default-action {
                action-type reject
            }
        }
```

The purpose of the BGP export policy is to advertise the system IP address of CSC-CE-4 as a labeled-IPv4 BGP route toward CSC-PE-3. The import policy imports the system IP address of CSC-CE-1.

Step 9. Verify exchange of labeled IPv4 routes between CSC-PE-3 and CSC-CE-4

When the preceding steps are completed, CSC-PE-3 advertises the labeled-IPv4 route for 192.0.2.1/32 to CSC-CE-4. This can be checked in the RIB Out for CSC VPRN 1 on CSC-PE-3, as follows:

```
[/]
A:admin@CSC-PE-3# show router 1 bgp routes 192.0.2.1/32 label-ipv4 hunt
===============================================================================
 BGP Router ID:192.0.2.3        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP LABEL-IPV4 Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
Network        : 192.0.2.1/32
Nexthop        : 192.168.34.1
Path Id        : None
To             : 192.168.34.2
Res. Nexthop   : n/a
Local Pref.    : n/a                    Interface Name : NotAvailable
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : n/a
Connector      : None
Community      : No Community Members
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.4
IPv4 Label     : 524285                 Label Type     : SWAP
Lbl Allocation : NEXT-HOP
Origin         : IGP
AS-Path        : 64500 64500
Route Tag      : 0
Neighbor-AS    : 64500
Orig Validation: NotFound
Source Class   : 0                      Dest Class     : 0

-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

CSC-PE-3 has advertised a label value of 524285 with the prefix.

The following output shows the received route in the RIB In of CSC-CE-4:

```
[/]
A:admin@CSC-CE-4# show router bgp routes 192.0.2.1/32 label-ipv4 hunt
===============================================================================
 BGP Router ID:192.0.2.4        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP LABEL-IPV4 Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : 192.0.2.1/32
Nexthop        : 192.168.34.1
Path Id        : None
From           : 192.168.34.1
Res. Nexthop   : 192.168.34.1
Local Pref.    : None                   Interface Name : int-CSC-CE-4-CSC-PE-3
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : 0
Connector      : None
Community      : No Community Members
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.3
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Fwd Class      : None                   Priority       : None
IPv4 Label     : 524285
Flags          : Used Valid Best IGP In-TTM In-RTM
Route Source   : External
AS-Path        : 64500 64500
Route Tag      : 0
Neighbor-AS    : 64500
Orig Validation: NotFound
Source Class   : 0                      Dest Class     : 0
Add Paths Send : Default
RIB Priority   : Normal
Last Modified  : 00h01m12s

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

The BGP distributed labels are programmed in the line cards of CSC-PE-3 based on the received VPN-IPv4 routes from CSC-PE-2 (Label Origin = Internal) and the advertised labeled-IPv4 routes to CSC-CE-4:

```
[/]
A:admin@CSC-PE-3# show router 1 bgp inter-as-label

===============================================================================
BGP Inter-AS labels
Flags: B - entry has backup, P - entry is promoted
===============================================================================
NextHop                       Received       Advertised     Label
                              Label          Label          Origin
-------------------------------------------------------------------------------
192.0.2.2                     524284         524283         Internal
192.0.2.2                     524285         524285         Internal
-------------------------------------------------------------------------------
Total Labels allocated:   2
===============================================================================
```

In the preceding output, the second entry for NextHop 192.0.2.2 corresponds to the prefix 192.0.2.1/32; recall from Step 7 that CSC-PE-3 received the VPN-IPv4 route with label value 524285 and it can be seen from this step that it re-advertised the route to CSC-CE-4 with the same label value 524285.

Step 10. Configure BGP session between CSC-CE-1 and CSC-CE-4

The final step in the setup of the CSC solution shown in Figure 302: CSC network topology is the creation of a BGP session between CSC-CE-1 and CSC-CE-4 so that they can exchange routes belonging to VPN services they support. The configuration of this BGP session on CSC-CE-1 is as follows:

```
# on CSC-CE-1:
configure {
    router "Base" {
        bgp {
            group "CSC-CE" {
                type internal
            }
            neighbor "192.0.2.4" {
                group "CSC-CE"
                family {
                    vpn-ipv4 true
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
```

The configuration of the BGP session on CSC-CE-4 is similar, as follows:

```
# on CSC-CE-4:
configure {
    router "Base" {
        bgp {
            group "CSC-CE" {
                type internal
            }
            neighbor "192.0.2.1" {
                group "CSC-CE"
                family {
                    vpn-ipv4 true
                }
            }
```

The configuration of the BGP session between CSC-CE-1 and CSC-CE-4 has the following properties:

- The peer type is IBGP (**type internal**. Alternatively, **peer-as** can be configured with the same value as the locally configured **autonomous-system**).
- The transport for the IBGP session is IPv4 (the **neighbor** address is an IPv4 address).
- The **family vpn-ipv4** command causes MP-BGP negotiation of the address family for AFI=1 and SAFI=128.

Step 11. Verify exchange of VPN-IPv4 routes

When the preceding steps have been completed, CSC-PE-3 can advertise a VPN-IPv4 route for some IP prefix (for example, 10.11.30.0/24) to CSC-CE-4. This can be checked in the RIB In of CSC-CE-4 as follows:

```
[/]
A:admin@CSC-CE-4# show router bgp routes 10.11.30.0/24 vpn-ipv4 hunt
===============================================================================
 BGP Router ID:192.0.2.4        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv4 Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : 10.11.30.0/24
Nexthop        : 192.0.2.1
Route Dist.    : 64496:11              VPN Label      : 524286
Path Id        : None
From           : 192.0.2.1
Res. Nexthop   : n/a
Local Pref.    : 100                   Interface Name : NotAvailable
Aggregator AS  : None                  Aggregator     : None
Atomic Aggr.   : Not Atomic            MED            : None
AIGP Metric    : None                  IGP Cost       : 0
Connector      : None
Community      : target:64496:1
Cluster        : No Cluster Members
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Originator Id  : None                   Peer Router Id : 192.0.2.1
Fwd Class      : None                   Priority       : None
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
Add Paths Send : Default
Last Modified  : 00h02m52s
VPRN Imported  :  1

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

The following command can be used to check that CSC-CE-4 has properly installed the preceding VPN-IPv4 route into the routing table of the importing VPRN service:

```
[/]
A:admin@CSC-CE-4# show router 1 route-table

===============================================================================
Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
10.11.30.0/24                                 Remote  BGP VPN   00h03m33s  170
       192.0.2.1 (tunneled:BGP)                                      1000
10.14.30.0/24                                 Local   Local     00h07m09s  0
       loopback-1                                                    0
-------------------------------------------------------------------------------
No. of Routes: 2
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

Conclusion

Carrier Supporting Carrier is a scalable and secure solution for using an infrastructure IP VPN to transport traffic between dispersed CSC-CE devices belonging to an ISP or other service provider. Many different topology models are supported by SR OS. This chapter has explored one simplified configuration that can serve as the basis for more complicated setups.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Flexible Algorithms for SRv6-based VPRNs

This chapter provides information about flexible algorithms (Flex-Algorithm) for VPRNs that are based on segment routing over IPv6 (SRv6).

Topics in this chapter include:

Applicability

The information and configuration in this chapter are based on SR OS Release 22.10.R1. The Flex-Algorithm for SRv6-based VPRNs feature is supported on FP-based platforms with FP4-based network ports in SR OS Release 21.5.R2 and later.

Overview

The Flex-Algorithm for SRv6-based VPRNs feature allows the computation of constraint-based paths across an SRv6-enabled network, based on metrics other than the default interior gateway protocol (IGP) metrics. The supported metrics are:

- IGP metrics
- link delay metrics
- traffic engineering (TE) metrics

Based on the metrics that are specifically configured for it, each Flex-Algorithm instance computes optimum paths across routers that participate in the Flex-Algorithm instance. For these paths, the IGP protocol automatically creates SRv6 tunnels between every pair of routers participating in the Flex-Algorithm instance. Two or more routers participate in a single Flex-Algorithm instance; a single router may participate in multiple Flex-Algorithm instances.

At least one router advertises (via extensions to the IGP protocol) the flexible algorithm definition (FAD). The **metric-type** command in the **routing-options flexible-algorithm-definitions flex-algo <flex-algo-name>** context configures the metric type that the Flex-Algorithm instance uses: *igp*, *delay*, or *te-metric*. The router that advertises the FAD typically also participates in the Flex-Algorithm instance. The other routers participate in the advertised Flex-Algorithm instance, without also advertising it. For reasons of redundancy, multiple routers may advertise the same FAD. In that case, the configuration of that FAD should be identical on all these routers. If not, all routers that participate in the Flex-Algorithm instance install from conflicting FADs only the FAD that has the highest priority value. If conflicting FADs have the same priority value, all routers that participate in the Flex-Algorithm instance install only the FAD that is advertised by the IS-IS-enabled router with the highest IS-IS system ID (or by the OSPF-enabled router with the highest OSPF router ID).

The **algorithm** command in the **router Base segment-routing segment-routing-v6 locator <locator-name>** context, associates each Flex-Algorithm instance (algorithm 128 to algorithm 255) with one specific SRv6 locator, which must be different from the base algorithm (algorithm 0) SRv6 locator. This algorithm

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

identifier is included in the SRv6 locator TLV when advertising the locator into IS-IS. The same FAD may be used in multiple Flex-Algorithm instances. Each Flex-Algorithm instance is associated with, at most, one SRv6 locator. Each SRv6 locator has, at most, one associated Flex-Algorithm instance. The **default-locator** command in the **service vprn <service-name> bgp-ipvpn segment-routing-v6 <bgp-instance> srv6** context configures the SRv6 locator that the VPRN data traffic across SRv6-enabled networks uses.

The further processing of the Flex-Algorithm-based VPRN data traffic across SRv6-enabled networks follows that of the base algorithm-based VPRN data traffic across SRv6-enabled networks, as described in the "Segment Routing over IPv6 for VPRN" chapter in the Segment Routing and PCE volume in the *7450 ESS, 7750 SR, and 7950 XRS MD-CLI Advanced Configuration Guide - Part I*.

Configuration

Figure 303: Example topology shows the example topology with four routers. The SRv6-enabled network that it represents comprises PE-1 and PE-2 in the control and data plane, and P-3 and P-4 in the data plane only. The SRv6-enabled network has only IPv6 addresses and interfaces. IS-IS is configured on all routers. BGP is configured only on PE-1 and PE-2.

Figure 303: Example topology

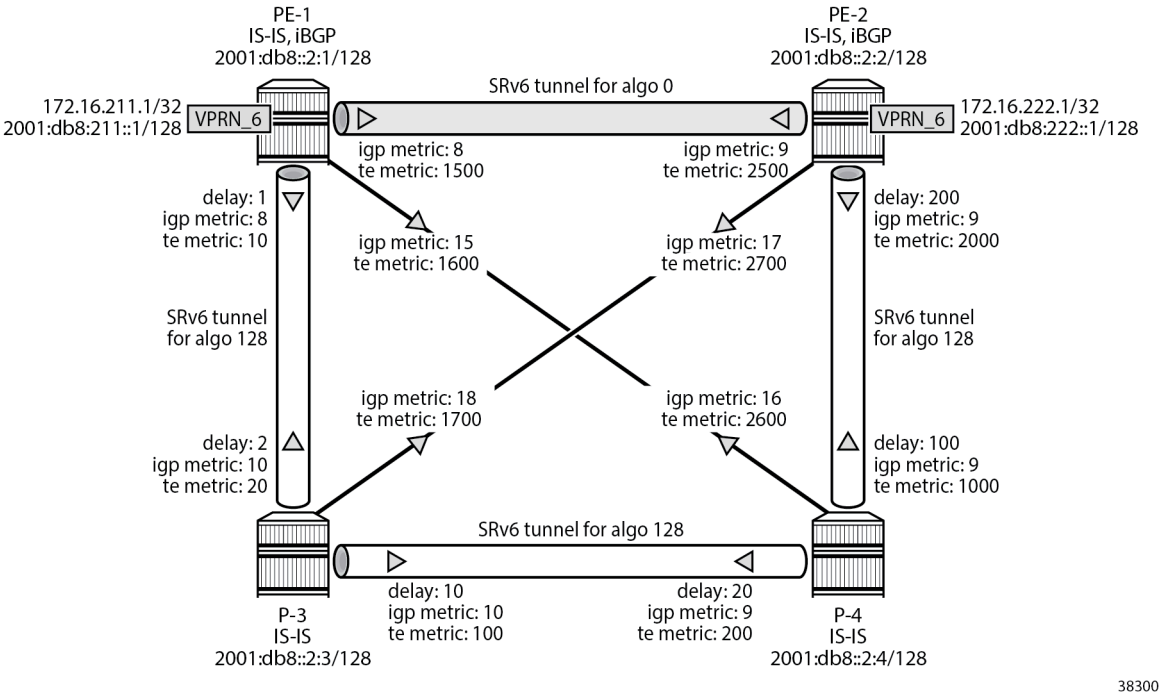

PE-1 and PE-2 are SRv6-enabled routers that each contain a VPRN instance. In this example, bidirectional IPv4 and IPv6 VPRN traffic flows are enabled between PE-1 and PE-2.

To illustrate what IS-IS interface metrics are used and how, the IS-IS interface metrics are configured explicitly (that is, differently from their default values), using the **metric** <*metric*> command (for IPv4 unicast traffic) and the **ipv6-unicast-metric** <*ipv6-unicast-metric*> command (for IPv6 unicast traffic) in the **router Base isis 0 interface <interface-name> level <level-number>** context. Different values can be applied for IS-IS level 1 and IS-IS level 2; for each IS-IS level, a distinction can be made between

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

IPv4 unicast traffic and IPv6 unicast traffic. For each IS-IS level and traffic type, different values can be configured in the two directions between two routers.

As a first example, the link delay metric is chosen for the Flex-Algorithm operation. The link delay metrics are configured explicitly, using the **static** <*value*> command (value in microseconds) in the **router Base interface <interface-name> if-attribute delay** context, only on the links between PE-1 and P-3, P-3 and P-4, and P-4 and PE-2. These metrics are configured differently in the two directions on each link (shown in the example topology as: "delay: <value in microseconds>"). The link delay metric can be configured on links between any pair of routers participating in the Flex-Algorithm instance. Any link that does not have the link delay metric configured is excluded from the Flex-Algorithm instance computation, which may result in no valid path between the ingress and egress routers. The link delay metric values are used for both IPv4 unicast traffic and IPv6 unicast traffic.

As a second example, the TE metric is chosen for the Flex-Algorithm operation. The TE metrics are configured explicitly, using the **te-metric** <*value*> command in the **router Base mpls interface <interface-name>** context, on all links. These metrics are configured differently in the two directions on each link, such that the direct link between the two routers PE-1 and PE-2 is always preferred.

The **ping** and **traceroute** commands between IPv4 and IPv6 loopback addresses in the VPRNs, as described in following sections, are used to simulate data traffic.

SRv6 requires wide metrics to match the 32-bit metric field in SRv6 locator TLV. The example configuration has wide metrics configured only for level 2. So, only the explicitly configured IS-IS level 2 interface metric values are used. Also, multi-topology and multi-protocol are not enabled in the example configuration. So, the explicitly configured IS-IS level 2 interface metric values are used for both IPv4 unicast traffic and IPv6 unicast traffic.

Configure the router

This configuration includes:

- ports and IPv6-only interfaces on PE-1, PE-2, P-3, and P-4, with link delay metrics configured where needed
- port cross-connect (PXC) on PE-1 and PE-2, using internal loopbacks on an FP4 MAC chip, as described in the "Segment Routing over IPv6" chapter in the Segment Routing and PCE volume in the *7450 ESS, 7750 SR, and 7950 XRS MD-CLI Advanced Configuration Guide - Part I*
- IS-IS on PE-1, PE-2, P-3, and P-4, which includes:
  - level 2 capability with wide metrics, and IPv4 metrics on all IS-IS level 2 interfaces
  - native IPv6 routing
  - the **traffic-engineering** and **traffic-engineering-options** commands, as a best practice to advertise the router capability within the autonomous system (AS)
- BGP on PE-1 and PE-2, with internal group "gr_v6_internal", which includes:
  - the IPv4 and IPv6 families
  - **extended-nh-encoding** for IPv4
  - **advertise-ipv6-next-hops** for IPv4
  - BGP neighbor **system** IPv6 addresses
  - **next-hop-self**

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

As the core network topology uses IPv6 for BGP peering (with IPv6 next hop addresses), the commands **advertise-ipv6-next-hops** and **extended-nh-encoding** need to be applied at the BGP, group, or neighbor level, so as to advertise and receive IPv4 routes with IPv6 next hop addresses. The **advertise-ipv6-next-hops** command instructs the system to advertise IPv4 routes with IPv6 next hop addresses. The **extended-nh-encoding** command configures BGP to advertise the capability to receive IPv4 routes with IPv6 next hop addresses.

The following example configuration applies for PE-1. A similar configuration applies for PE-2. P-3 and P-4 have no BGP configuration.

```
A:admin@PE-1# configure {
    router "Base" {
        autonomous-system 64500
        interface "int-PE-1-PE-2" {
            description "interface between PE-1 and PE-2"
            port 1/1/c1/1:1000
            ipv6 {
                address 2001:db8::168:12:1 {
                    prefix-length 126
                }
            }
            admin-state enable
        }
        interface "int-PE-1-P-3" {
            description "interface between PE-1 and P-3"
            port 1/1/c2/1:1000
            ipv6 {
                address 2001:db8::168:13:1 {
                    prefix-length 126
                }
            }
            if-attribute {
                delay {
                    static 1    # microseconds
                }
            }
            admin-state enable
        }
        interface "int-PE-1-P-4" {
            description "interface between PE-1 and P-4"
            port 1/1/c3/1:1000
            ipv6 {
                address 2001:db8::168:14:1 {
                    prefix-length 126
                }
            }
            admin-state enable
        }
        interface "system" {
            description "system interface of PE-1"
            ipv6 {
                address 2001:db8::2:1 {
                    prefix-length 128
                }
            }
            admin-state enable
        }
        isis 0 {
            level-capability 2
            area-address [49.0001]
            traffic-engineering true
            traffic-engineering-options {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                ipv6
                application-link-attributes {
                }
            }
            advertise-router-capability as
            ipv6-routing native
            level 2 {
                wide-metrics-only true    # required for SRv6
            }
            interface "system" {
                passive true
                admin-state enable
            }
            router-id 1.1.1.1
            interface "int-PE-1-PE-2" {
                interface-type point-to-point
                level 2 {
                    metric 8
                }
                admin-state enable
            }
            interface "int-PE-1-P-3" {
                interface-type point-to-point
                level 2 {
                    metric 8
                }
                admin-state enable
            }
            interface "int-PE-1-P-4" {
                interface-type point-to-point
                level 2 {
                    metric 15
                }
                admin-state enable
            }
            admin-state enabe
        }
        bgp {
            min-route-advertisement 1
            router-id 1.1.1.10
            rapid-withdrawal true
            split-horizon true
            ebgp-default-reject-policy {
                import false
                export false
            }
            group "gr_v6_internal" {
                description "internal bgp group on PE-1"
                family {
                    ipv4 true
                    ipv6 true
                }
                next-hop-self true
                type internal
                extended-nh-encoding {
                    ipv4 true
                }
                advertise-ipv6-next-hops {
                    ipv4 true
                }
            }
            neighbor "2001:db8::2:2" {
                group "gr_v6_internal"
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            admin-state enable
        }
    }
}
```

Configure the VPRN services on PE-1 and on PE-2

This configuration includes:

- an IPv4 address and an IPv6 address for a loopback interface "lb_itf_vprn"
- BGP, with external group "gr_v6_vprn", which includes:
  - IPv4 and IPv6 families
  - **extended-nh-encoding** for IPv4
  - **advertise-ipv6-next-hops** for IPv4
  - BGP neighbor **interface** IPv6 addresses, with BGP neighbors in a different external AS

The following example configuration applies for the VPRN on PE-1. A similar configuration applies for the VPRN on PE-2.

```
A:admin@PE-1# configure {
    service {
        vprn "VPRN_6" {
            service-id 6
            customer "1"
            description "VPRN_6 on PE-1"
            autonomous-system 64500
            interface "lb_itf_vprn" {
                ipv4 {
                    primary {
                        address 172.16.211.1
                        prefix-length 32
                    }
                }
                description "VPRN_6 interface on PE-1 for external subnet"
                ipv6 {
                    address 2001:db8:211::1 {
                        prefix-length 128
                    }
                }
                loopback true
            }
            bgp {
                ebgp-default-reject-policy {
                    import false
                    export false
                }
                group "gr_v6_vprn" {
                    description "external bgp group for VPRN_6 on PE-1"
                    family {
                        ipv4 true
                        ipv6 true
                    }
                    extended-nh-encoding {
                        ipv4 true
                    }
                    advertise-ipv6-next-hops {
                        ipv4 true
                    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
                neighbor "2001:db8:101::1" {
                    group "gr_v6_vprn"
                    type external
                    peer-as 64501
                }
                admin-state enable
            }
            admin-state enable
        }
    }
}
```

Configure SRv6 in the router Base context on PE-1 and PE-2

Configure the SRv6 **locator** in the **router Base segment-routing segment-routing-v6** context on PE-2. Perform a similar configuration on PE-1, with **ip-prefix** *2001:db8:aaaa:101::/64* for SRv6 **locator** *"PE-1_loc"*.

```
A:admin@PE-2# configure {
    router "Base" {
        segment-routing {
            segment-routing-v6 {
                locator "PE-2_loc" {
                    block-length 48
                    function-length 20
                    prefix {
                        ip-prefix 2001:db8:aaaa:102::/64
                    }
                    admin-state enable
                }
            }
        }
    }
}
```

Configure the FPEs on PE-1 and PE-2.

```
A:admin@PE-2# configure {
    fwd-path-ext {
        fpe 1 {
            path {
                pxc 1
            }
            application {
                srv6 {
                    type origination
                }
            }
        }
        fpe 2 {
            path {
                pxc 2
            }
            application {
                srv6 {
                    type termination
                }
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    }
}
```

Use FPE 1 as the SRv6 origination FPE in the **router Base segment-routing segment-routing-v6** context and FPE 2 as the SRv6 termination FPE in the **router Base segment-routing segment-routing-v6 locator <locator-name>** context on PE-2. Perform a similar configuration on PE-1, for SRv6 **locator** *"PE-1_loc"*. For more information, see the "Segment Routing over IPv6" chapter in the Segment Routing and PCE volume in the *7450 ESS, 7750 SR, and 7950 XRS MD-CLI Advanced Configuration Guide - Part I*.

```
*A:PE-2# configure {
    router "Base" {
        segment-routing {
            segment-routing-v6 {
                origination-fpe [1]
                locator "PE-2_loc" {
                    termination-fpe [2]
                    admin-state enable
                }
            }
        }
    }
}
```

Configure the SRv6 End function (equivalent to an IPv4 node SID) in the **router Base segment-routing segment-routing-v6 base-routing-instance locator <locator-name>** context on PE-2. Perform a similar configuration on PE-1, for SRv6 **locator** *"PE-1_loc"*.

```
A:admin@PE-2# configure {
    router "Base" {
        segment-routing {
            segment-routing-v6 {
                base-routing-instance {
                    locator "PE-2_loc" {
                        function {
                            end 1 {
                                srh-mode usp
                            }
                        }
                    }
                }
            }
        }
    }
}
```

Advertise the SRv6 locator in IS-IS while ensuring level 2 capability on PE-2. Perform a similar configuration on PE-1, for SRv6 **locator** *"PE-1_loc"*.

```
A:admin@PE-2# configure {
    router "Base" {
        isis 0 {
            segment-routing-v6 {
                locator "PE-2_loc" {
                    level-capability 2
                }
                admin-state enable
            }
        }
    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
}
```

Configure SRv6 for the VPRNs on PE-1 and PE-2

On PE-1 and PE-2, extend the BGP advertisements to include the VPN-IPv4 and VPN-IPv6 families.

```
A:admin@PE-1#/A:admin@PE-2# configure {
    router "Base" {
        bgp {
            rapid-update {
                vpn-ipv4 true
                vpn-ipv6 true
            }
            group "gr_v6_internal" {
                family {
                    ipv4 true
                    ipv6 true
                    vpn-ipv4 true
                    vpn-ipv6 true
                }
                extended-nh-encoding {
                    ipv4 true
                    vpn-ipv4 true
                }
                advertise-ipv6-next-hops {
                    ipv4 true
                    vpn-ipv4 true
                    vpn-ipv6 true
                }
            }
            admin-state enable
        }
    }
}
```

On PE-2, create an SRv6 instance for the VPRN service. Use the SRv6 locator from the **router Base segment-routing segment-routing-v6** context and configure End.DT4 and End.DT6 functions for it.

Use the created SRv6 instance in the **service vprn <service-name> bgp-ipvpn segment-routing-v6 <bgp-instance>** context, with the configured SRv6 locator as the default locator. Ensure a unique route distinguisher. Use the unique PE-2 system IPv6 address as the source address. Perform a similar configuration on PE-1, with the PE-1 SRv6 locator as the default locator, the PE-1 system IPv6 address as the source address, and a different route distinguisher.

```
A:admin@PE-2# configure {
    service {
        vprn "VPRN_6" {
            segment-routing-v6 1 {
                locator "PE-2_loc" {
                    function {
                        end-dt4 {
                        }
                        end-dt6 {
                        }
                    }
                }
            }
            bgp-ipvpn {
                segment-routing-v6 1 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    route-distinguisher "192.0.2.2:6"
                    srv6 {
                        instance 1
                        default-locator "PE-2_loc"
                    }
                    source-address 2001:db8::2:2
                    vrf-target {
                        community "target:64506:6"
                    }
                    admin-state enable
                }
            }
            admin-state enable
        }
    }
}
```

Verify data traffic

At this point, using **ping** and **traceroute** commands, verify that IPv4 data traffic is possible between the local VPRN on PE-1 and the remote VPRN on PE-2:

```
A:admin@PE-1# ping 172.16.222.1 router-instance "VPRN_6"
PING 172.16.222.1 56 data bytes
64 bytes from 172.16.222.1: icmp_seq=1 ttl=64 time=1.72ms.
---snip---
---- 172.16.222.1 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 1.51ms, avg = 1.65ms, max = 1.75ms, stddev = 0.085ms

A:admin@PE-1# traceroute 172.16.222.1 router-instance "VPRN_6"
traceroute to 172.16.222.1, 30 hops max, 40 byte packets
  1  172.16.222.1 (172.16.222.1)    1.64 ms  1.57 ms  1.61 ms
```

This data traffic uses the SRv6 tunnels over the direct link between PE-1 and PE-2.

For IPv4 data traffic, the VPRN routing table shows the next hop and the applicable IGP metric for the route to 172.16.222.1.

```
A:admin@PE-1# show router 6 route-table 172.16.222.1

===============================================================================
Route Table (Service: 6)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
172.16.222.1/32                               Remote  BGP VPN   00h04m21s  170
       2001:db8:aaaa:102:8000:: (tunneled:SRV6)                      8
-------------------------------------------------------------------------------
No. of Routes: 1
---snip---
===============================================================================
```

The next hop is the End.DT4 SRv6 SID of the SRv6 **locator** *"PE-2_loc"* for the VPRN on PE-2, which PE-1 learns from a BGP update from PE-2. The SRv6 tunnel to this next hop has label 524288.

```
A:admin@PE-2# show router segment-routing-v6 local-sid
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Segment Routing v6 Local SIDs
===============================================================================
SID                                             Type          Function
  Locator
  Context
-------------------------------------------------------------------------------
2001:db8:aaaa:102:0:1000::                      End           1
  PE-2_loc
  Base
2001:db8:aaaa:102:7fff:f000::                   End.DT6       524287
  PE-2_loc
  SvcId: 6 Name: VPRN_6
2001:db8:aaaa:102:8000::                        End.DT4       524288
  PE-2_loc
  SvcId: 6 Name: VPRN_6
-------------------------------------------------------------------------------
SIDs : 3
-------------------------------------------------------------------------------
===============================================================================
```

The applicable IGP metric is 8, which corresponds with the IS-IS level 2 "int-PE-1-PE-2" interface metric value.

```
A:admin@PE-1# show router isis 0 interface

===============================================================================
Rtr Base ISIS Instance 0 Interfaces
===============================================================================
Interface                       Level CircID  Oper      L1/L2 Metric     Type
                                              State
-------------------------------------------------------------------------------
system                          L1L2  1       Up        0/0              p2p
int-PE-1-P-3                    L1L2  2       Up        6/8              p2p
int-PE-1-P-4                    L1L2  3       Up        11/15            p2p
int-PE-1-PE-2                   L1L2  4       Up        6/8              p2p
-------------------------------------------------------------------------------
Interfaces : 4
===============================================================================
```

The **show router isis 0 topology** command lists the IS-IS nodes, and for each IS-IS node, the outgoing interface and the next hop. There are only IS-IS nodes at IS-IS level 2. The output of this command shows that data traffic from PE-1 to PE-2 uses interface "int-PE-1-PE-2" to PE-2, while, for example, data traffic from PE-1 to P-3 uses interface "int-PE-1-P-3" to P-3.

```
A:admin@PE-1# show router isis 0 topology

===============================================================================
Rtr Base ISIS Instance 0 Topology Table
===============================================================================
Node                              Interface                 Nexthop
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
IS-IS IP paths (MT-ID 0),   Level 2
-------------------------------------------------------------------------------
PE-2.00                           int-PE-1-PE-2             PE-2
P-3.00                            int-PE-1-P-3              P-3
P-4.00                            int-PE-1-P-4              P-4
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The **show router isis 0 topology detail** command lists the IS-IS nodes, and for each IS-IS node, the next hop, the outgoing interface, and the metric (in this case, IS-IS level 2) that applies.

```
A:admin@PE-1# show router isis 0 topology detail

===============================================================================
Rtr Base ISIS Instance 0 Topology Table
===============================================================================
-------------------------------------------------------------------------------
IS-IS IP paths (MT-ID 0),   Level 2
-------------------------------------------------------------------------------
Node      : PE-2.00
Nexthop   : PE-2
Interface : int-PE-1-PE-2
SNPA      : none                            Metric      : 8

Node      : P-3.00
Nexthop   : P-3
Interface : int-PE-1-P-3
SNPA      : none                            Metric      : 8

Node      : P-4.00
Nexthop   : P-4
Interface : int-PE-1-P-4
SNPA      : none                            Metric      : 15

===============================================================================
```

Verify the IS-IS data base on PE-1 with **show router isis 0 database detail**. The output of this command (shortened here for PE-1 and PE-2, and omitted for P-3 and P-4) provides information about each IS-IS-enabled router. For each uniquely identified IS-IS-enabled router, the SRv6 information indicates:

- the IS-IS-advertised router capabilities
- the IS-IS topology details
- the IPv4 and IPv6 reachability details
- the advertised SRv6 locator TLV
- the advertised configured SRv6 End SID

There is only IS-IS information at IS-IS level 2. On each IS-IS interface, the IS-IS level 2 metrics for IPv4 and IPv6 are identical.

Only the default metric-based SPF algorithm instance 0 is in use, as listed in the SR Alg sub-TLV of the Router Cap TLVs.

On PE-1 and PE-2, the SRv6 locator prefix for algorithm 0, with their End SRv6 SID are present.

```
A:admin@PE-1# show router isis 0 database detail

===============================================================================
Rtr Base ISIS Instance 0 Database (detail)
===============================================================================
---snip---
Displaying Level 2 database
-------------------------------------------------------------------------------
LSP ID    : PE-1.00-00                                Level     : L2
---snip---
TLVs :
  Area Addresses:
    Area Address : (3) 49.0001
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
  Supp Protocols:
    Protocols     : IPv4
    Protocols     : IPv6
  IS-Hostname   : PE-1
  Router ID   :
    Router ID   : 1.1.1.1
  TE Router ID v6  :
    Router ID   : 2001:db8::2:1
  Router Cap : 1.1.1.1, D:0, S:0
    TE Node Cap : B E M  P
    SRv6 Cap: 0x0000
    SR Alg: metric based SPF
    Node MSD Cap: BMI : 0 SRH-MAX-SL : 10 SRH-MAX-END-POP : 9 SRH-MAX-H-ENCAPS : 3 SRH-MAX-END-
D : 9
  I/F Addresses IPv6 :
    IPv6 Address    : 2001:db8::2:1
    IPv6 Address    : 2001:db8::168:12:1
    IPv6 Address    : 2001:db8::168:13:1
    IPv6 Address    : 2001:db8::168:14:1
  TE IS Nbrs   :
    Nbr   : PE-2.00
    Default Metric  : 8
    Sub TLV Len     : 36
    IPv6 Addr : 2001:db8::168:12:1
    Nbr IPv6  : 2001:db8::168:12:2
  TE IS Nbrs   :
    Nbr   : P-3.00
    Default Metric  : 8
    Sub TLV Len     : 36
    IPv6 Addr : 2001:db8::168:13:1
    Nbr IPv6  : 2001:db8::168:13:2
  TE IS Nbrs   :
    Nbr   : P-4.00
    Default Metric  : 15
    Sub TLV Len     : 36
    IPv6 Addr : 2001:db8::168:14:1
    Nbr IPv6  : 2001:db8::168:14:2
  IPv6 Reach:
    Metric: ( I ) 0
    Prefix   : 2001:db8::2:1/128
    Metric: ( I ) 8
    Prefix   : 2001:db8::168:12:0/126
    Metric: ( I ) 8
    Prefix   : 2001:db8::168:13:0/126
    Metric: ( I ) 15
    Prefix   : 2001:db8::168:14:0/126
    Metric: ( I ) 0
    Prefix   : 2001:db8:aaaa:101::/64
  SRv6 Locator  :
    MT ID : 0
    Metric: ( ) 0 Algo:0
    Prefix   : 2001:db8:aaaa:101::/64
    Sub TLV   :
      End-SID   : 2001:db8:aaaa:101:0:1000::, flags:0x0, endpoint:End-USP

-------------------------------------------------------------------------------
LSP ID    : PE-2.00-00                                  Level     : L2
---snip---
TLVs :
  Area Addresses:
    Area Address : (3) 49.0001
  Supp Protocols:
    Protocols     : IPv4
    Protocols     : IPv6
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
  IS-Hostname   : PE-2
  Router ID   :
    Router ID   : 2.2.2.2
  TE Router ID v6  :
    Router ID   : 2001:db8::2:2
  Router Cap : 2.2.2.2, D:0, S:0
    TE Node Cap : B E M  P
    SRv6 Cap: 0x0000
    SR Alg: metric based SPF
    Node MSD Cap: BMI : 0 SRH-MAX-SL : 10 SRH-MAX-END-POP : 9 SRH-MAX-H-ENCAPS : 3 SRH-MAX-END-
D : 9
  ---snip---
  TE IS Nbrs   :
    Nbr   : PE-1.00
    Default Metric  : 9
    ---snip---
  TE IS Nbrs   :
    Nbr   : P-3.00
    Default Metric  : 17
    ---snip---
  TE IS Nbrs   :
    Nbr   : P-4.00
    Default Metric  : 9
    ---snip---
  ---snip---
  SRv6 Locator  :
    MT ID : 0
    Metric: ( ) 0 Algo:0
    Prefix   : 2001:db8:aaaa:102::/64
    Sub TLV   :
      End-SID   : 2001:db8:aaaa:102:0:1000::, flags:0x0, endpoint:End-USP
---snip---
Level (2) LSP Count : 4
-------------------------------------------------------------------------------
---snip---
SABM-flags Flags:  R = RSVP-TE
                   S = SR-TE
                   F = LFA
                   X = FLEX-ALGO
FAD-flags Flags:   M = Prefix Metric
===============================================================================
```

PE-2 advertises to PE-1 the information for network prefix 172.16.222.1/32, as listed in the RIB In Entries section in the following example. PE-1 acts in a similar way as PE-2 for network prefix 172.16.211.1/32, as listed in the RIB Out Entries section.

The following output shows the corresponding VPN-IPv4 BGP routes on PE-1:

```
A:admin@PE-1# show router bgp routes vpn-ipv4 hunt
===============================================================================
 BGP Router ID:1.1.1.10         AS:64500       Local AS:64500
===============================================================================
---snip---
===============================================================================
BGP VPN-IPv4 Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : 172.16.222.1/32
Nexthop        : 2001:db8::2:2
Route Dist.    : 192.0.2.2:6           VPN Label      : 524288
Path Id        : None
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
From           : 2001:db8::2:2
Res. Nexthop   : n/a
Local Pref.    : 100                    Interface Name : int-PE-1-PE-2
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : 8
Connector      : None
Community      : target:64506:6
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 2.2.2.10
Fwd Class      : None                   Priority       : None
Flags          : Valid Best IGP
---snip---
SRv6 TLV Type  : SRv6 L3 Service TLV (5)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:aaaa:102::
Full Sid       : 2001:db8:aaaa:102:8000::
Behavior       : End.DT4 (19)
SRv6 SubSubTLV : SRv6 SID Structure (1)
---snip---
VPRN Imported  : None

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
Network        : 172.16.211.1/32
Nexthop        : 2001:db8::2:1
Route Dist.    : 192.0.2.1:6            VPN Label      : 524288
Path Id        : None
To             : 2001:db8::2:2
Res. Nexthop   : n/a
Local Pref.    : 100                    Interface Name : NotAvailable
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : n/a
Connector      : None
Community      : target:64506:6
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 2.2.2.10
Origin         : IGP
---snip---
SRv6 TLV Type  : SRv6 L3 Service TLV (5)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:aaaa:101::
Full Sid       : 2001:db8:aaaa:101:8000::
Behavior       : End.DT4 (19)
SRv6 SubSubTLV : SRv6 SID Structure (1)
---snip---
-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

Verify the IPv6 route table on PE-1. The IPv6 route table has routes to the local and remote SRv6 locators and to the local SRv6 End function SID. The SRv6 locator prefix of PE-2 is reached via an SRv6 tunnel using IS-IS. The routes with protocol **"SRV6"** correspond with the locally configured SRv6 locator prefix of PE-1 and the locally configured SRv6 End function.

```
A:admin@PE-1# show router route-table ipv6

===============================================================================
IPv6 Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
---snip---
2001:db8:aaaa:101::/64                      Local   SRV6      00h11m44s  3
       fe80::201-"_tmnx_fpe_2.a"                                    0
2001:db8:aaaa:101:0:1000::/128              Local   SRV6      00h11m44s  3
       Black Hole                                                   0
2001:db8:aaaa:102::/64                      Remote  ISIS      00h10m36s  18
       2001:db8:aaaa:102::/64 (tunneled:SRV6-ISIS)                  8
-------------------------------------------------------------------------------
No. of Routes: 13
---snip---
===============================================================================
```

Verify that the tunnel from PE-1 to the SRv6 locator prefix of PE-2 is an SRv6 tunnel that uses the "int-PE-1-PE-2" interface. A similar verification can be performed for the other direction. Interface "int-PE-1-PE-2" is configured on port 1/1/c1/1:1000.

```
A:admin@PE-1# show router fp-tunnel-table 1 ipv6

===============================================================================
IPv6 Tunnel Table Display
---snip---
===============================================================================
Destination                                 Protocol         Tunnel-ID
  Lbl/SID
    NextHop                                                  Intf/Tunnel
  Lbl/SID (backup)
    NextHop   (backup)
-------------------------------------------------------------------------------
2001:db8:aaaa:102::/64                      SRV6              524289
  -
    fe80::60e:1ff:fe01:1-"int-PE-1-PE-2"                    1/1/c1/1:1000
-------------------------------------------------------------------------------
Total Entries : 1
-------------------------------------------------------------------------------
===============================================================================
```

Verify also that IPv6 data traffic is possible between the local VPRN on PE-1 and the remote VPRN on PE-2:

```
A:admin@PE-1# ping 2001:db8:222::1 router-instance "VPRN_6"
PING 2001:db8:222::1 56 data bytes
64 bytes from 2001:db8:222::1 icmp_seq=1 hlim=64 time=1.72ms.
---snip---
---- 2001:db8:222::1 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 1.46ms, avg = 1.74ms, max = 2.21ms, stddev = 0.249ms

A:admin@PE-1# traceroute 2001:db8:222::1 router-instance "VPRN_6"
traceroute to 2001:db8:222::1, 30 hops max, 60 byte packets
  1  2001:db8:222::1 (2001:db8:222::1)    1.63 ms  1.52 ms  1.75 ms
```

For IPv6 data traffic, the VPRN routing table shows the next hop and the applicable IGP metric for the route to 2001:db8:222::1.

```
A:admin@PE-1# show router 6 route-table 2001:db8:222::1

===============================================================================
IPv6 Route Table (Service: 6)
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
2001:db8:222::1/128                             Remote  BGP VPN   00h04m21s  170
       2001:db8:aaaa:102:7fff:f000:: (tunneled:SRV6)                8
-------------------------------------------------------------------------------
No. of Routes: 1
---snip---
===============================================================================
```

The next hop is the End.DT6 SRv6 SID of the SRv6 **locator** *"PE-2_loc"* for the VPRN on PE-2, which PE-1 learns from a BGP update from PE-2. The SRv6 tunnel to this next hop has label 524287.

The applicable IGP metric is also 8, which again corresponds with the IS-IS level 2 "int-PE-1-PE-2" interface metric value.

PE-2 advertises to PE-1 the information for network prefix 2001:db8:222::1/128, as listed in the RIB In Entries section in the following example. PE-1 acts in a similar way as PE-2 for network prefix 2001:db8:211::1/128, as listed in the RIB Out Entries section.

The following output shows the corresponding VPN-IPv6 BGP routes on PE-1:

```
A:admin@PE-1# show router bgp routes vpn-ipv6 hunt
===============================================================================
 BGP Router ID:1.1.1.10         AS:64500       Local AS:64500
===============================================================================
---snip---
===============================================================================
BGP VPN-IPv6 Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : 2001:db8:222::1/128
Nexthop        : 2001:db8::2:2
Route Dist.    : 192.0.2.2:6            VPN Label      : 524287
Path Id        : None
From           : 2001:db8::2:2
Res. Nexthop   : n/a
Local Pref.    : 100                    Interface Name : int-PE-1-PE-2
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : 8
Connector      : None
Community      : target:64506:6
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 2.2.2.10
Fwd Class      : None                   Priority       : None
Flags          : Valid Best IGP
---snip---
SRv6 TLV Type  : SRv6 L3 Service TLV (5)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:aaaa:102::
Full Sid       : 2001:db8:aaaa:102:7fff:f000::
Behavior       : End.DT6 (18)
SRv6 SubSubTLV : SRv6 SID Structure (1)
---snip---
VPRN Imported  : None

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
Network        : 2001:db8:211::1/128
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Nexthop         : 2001:db8::2:1
Route Dist.     : 192.0.2.1:6             VPN Label      : 524287
Path Id         : None
To              : 2001:db8::2:2
Res. Nexthop    : n/a
Local Pref.     : 100                     Interface Name : NotAvailable
Aggregator AS   : None                    Aggregator     : None
Atomic Aggr.    : Not Atomic              MED            : None
AIGP Metric     : None                    IGP Cost       : n/a
Connector       : None
Community       : target:64506:6
Cluster         : No Cluster Members
Originator Id   : None                    Peer Router Id : 2.2.2.10
Origin          : IGP
---snip---
SRv6 TLV Type   : SRv6 L3 Service TLV (5)
SRv6 SubTLV     : SRv6 SID Information (1)
Sid             : 2001:db8:aaaa:101::
Full Sid        : 2001:db8:aaaa:101:7fff:f000::
Behavior        : End.DT6 (18)
SRv6 SubSubTLV : SRv6 SID Structure (1)
---snip---
-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

Configure a delay-based flexible algorithm

Define a Flex-Algorithm definition "FAD_delay" that takes delay as its metric. The FAD can reside on any IS-IS-enabled router. In this example, it resides on PE-1.

```
A:admin@PE-1# configure {
    routing-options {
        flexible-algorithm-definitions {
            flex-algo "FAD_delay" {
                description "FAD_delay_based"
                metric-type delay
                admin-state enable
            }
        }
    }
}
```

Configure PE-1 so as to advertise the FAD "FAD_delay" in Flex-Algorithm instance 128 (possible values between 128 and 255) and to participate in the Flex-Algorithm instance 128.

```
A:admin@PE-1# configure {
    router "Base" {
        isis 0 {
            flexible-algorithms {
                flex-algo 128 {
                    advertise "FAD_delay"
                    participate true
                    loopfree-alternate {
                    }
                }
                admin-state enable
            }
        }
    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
}
```

Ensure that the other IS-IS-enabled routers that must support the Flex-Algorithm instance participate in this Flex-Algorithm instance.

```
A:admin@PE-2#/A:admin@P-3#/A:admin@P-4# configure {
    router "Base" {
        isis 0 {
            flexible-algorithms {
                flex-algo 128 {
                    participate true
                    loopfree-alternate {
                    }
                }
                admin-state enable
            }
        }
    }
}
```

Define a new and unique SRv6 locator prefix to this new Flex-Algorithm instance. A separate SRv6 locator is needed for each Flex-Algorithm instance. So, for the Flex-Algorithm instance 128, configure SRv6 **locator** *"PE-2_loc_FAD128"* on PE-2. Perform a similar configuration for SRv6 **locator** *"PE-1_loc_FAD128"* with **ip-prefix** *2001:db8:a128:101::/64* on PE-1.

```
A:admin@PE-2# configure {
    router "Base" {
        segment-routing {
            segment-routing-v6 {
                locator "PE-2_loc_FAD128" {
                    block-length 48
                    function-length 20
                    algorithm 128
                    prefix {
                        ip-prefix 2001:db8:a128:102::/64
                    }
                    admin-state enable
                }
            }
        }
    }
}
```

For SRv6 **locator** *"PE-2_loc_FAD128"* on PE-2, use FPE 2 as the SRv6 termination FPE in the **router Base segment-routing segment-routing-v6 locator <locator-name>** context and configure the SRv6 End function (equivalent to an IPv4 node SID) in the **router Base segment-routing segment-routing-v6 base-routing-instance locator <locator-name>** context. Perform a similar configuration on PE-1, for SRv6 **locator** *"PE-1_loc_FAD128"*.

```
A:admin@PE-2# configure {
    router "Base" {
        segment-routing {
            segment-routing-v6 {
                locator "PE-2_loc_FAD128" {
                    termination-fpe [2]
                    admin-state enable
                }
                base-routing-instance {
                    locator "PE-2_loc_FAD128" {
                        function {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                            end 1 {
                                srh-mode usp
                            }
                        }
                    }
                }
            }
        }
    }
}
```

In the **router Base isis 0 segment-routing-v6 locator <locator-name>** context on PE-2, configure the IS-IS level capability for SRv6 **locator** *"PE-2_loc_FAD128"* and enable SRv6 in the IS-IS context. Perform a similar configuration on PE-1, for SRv6 **locator** *"PE-1_loc_FAD128"*.

```
A:admin@PE-2# configure {
    router "Base" {
        isis 0 {
            segment-routing-v6 {
                locator "PE-2_loc_FAD128" {
                    level-capability 2
                }
                admin-state enable
            }
        }
    }
}
```

The **show router isis 0 segment-routing-v6 locator** command lists the local and remote SRv6 locator prefixes for all applied algorithms (0 and 128).

```
A:admin@PE-2# show router isis 0 segment-routing-v6 locator

===============================================================================
Rtr Base ISIS Instance 0 SRv6 Locator Table
===============================================================================
Prefix                              AdvRtr                        MT     Lvl/Typ
 AttributeFlags                      Tag                           Flags  Algo
-------------------------------------------------------------------------------
2001:db8:a128:101::/64              PE-1                          0      2/Int.
  -                                   0                             -      128
2001:db8:a128:102::/64              PE-2                          0      2/Int.
  -                                   0                             -      128
2001:db8:aaaa:101::/64              PE-1                          0      2/Int.
  -                                   0                             -      0
2001:db8:aaaa:102::/64              PE-2                          0      2/Int.
  -                                   0                             -      0
-------------------------------------------------------------------------------
No. of Locators: 4
-------------------------------------------------------------------------------
---snip---
===============================================================================
```

On PE-2, in the **service vprn <service-name> segment-routing-v6 <bgp-instance> locator <locator-name>** context, configure End.DT4 and End.DT6 functions for SRv6 **locator** *"PE-2_loc_FAD128"*.

On PE-2, in the **service vprn <service-name> bgp-ipvpn segment-routing-v6 <bgp-instance> srv6** context, use the SRv6 **locator** *"PE-2_loc_FAD128"* as the default locator, instead of the earlier SRv6 **locator** *"PE-2_loc"*. Perform a similar configuration on PE-1, for SRv6 **locator** *"PE-1_loc_FAD128"*.

```
A:admin@PE-2# configure {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    service {
        vprn "VPRN_6" {
            segment-routing-v6 1 {
                locator "PE-2_loc_FAD128" {
                    function {
                        end-dt4 {
                        }
                        end-dt6 {
                        }
                    }
                }
            }
            bgp-ipvpn {
                segment-routing-v6 1 {
                    srv6 {
                        instance 1
                        default-locator "PE-2_loc_FAD128"
                    }
                    admin-state enable
                }
            }
            admin-state enable
        }
    }
}
```

Verify data traffic

At this point, using **ping** and **traceroute** commands, verify that data traffic between the local VPRN on PE-1 and the remote VPRN on PE-2 uses the Flex-Algorithm.

For IPv4 data traffic:

```
A:admin@PE-1# ping 172.16.222.1 router-instance "VPRN_6"
PING 172.16.222.1 56 data bytes
64 bytes from 172.16.222.1: icmp_seq=1 ttl=64 time=2.84ms.
---snip---
---- 172.16.222.1 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.49ms, avg = 2.68ms, max = 2.84ms, stddev = 0.136ms

A:admin@PE-1# traceroute 172.16.222.1 router-instance "VPRN_6"
traceroute to 172.16.222.1, 30 hops max, 40 byte packets
  1  172.16.222.1 (172.16.222.1)    3.04 ms  2.60 ms  2.54 ms
```

For IPv6 data traffic:

```
A:admin@PE-1# ping 2001:db8:222::1 router-instance "VPRN_6"
PING 2001:db8:222::1 56 data bytes
64 bytes from 2001:db8:222::1 icmp_seq=1 hlim=64 time=2.11ms.
---snip---
---- 2001:db8:222::1 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 1.87ms, avg = 2.07ms, max = 2.43ms, stddev = 0.194ms

A:admin@PE-1# traceroute 2001:db8:222::1 router-instance "VPRN_6"
traceroute to 2001:db8:222::1, 30 hops max, 60 byte packets
  1  2001:db8:222::1 (2001:db8:222::1)    2.89 ms  2.70 ms  2.55 ms
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

This data traffic uses the SRv6 tunnels over the links between PE-1 and P-3, P-3 and P-4, and P-4 and PE-2.

For IPv4 data traffic, the VPRN routing table shows the next hop and the applicable metric for the route to 172.16.222.1. In this example, the End.DT4 SID is copied into the IPv6 DA field of the tunneled packet.

```
A:admin@PE-1# show router 6 route-table 172.16.222.1

===============================================================================
Route Table (Service: 6)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
172.16.222.1/32                               Remote  BGP VPN   00h01m59s  170
       2001:db8:a128:102:7fff:e000:: (tunneled:SRV6)                 111
-------------------------------------------------------------------------------
No. of Routes: 1
---snip---
===============================================================================
```

The tunnel next hop is the End.DT4 SRv6 SID of the SRv6 **locator** *"PE-2_loc_FAD128"* for the VPRN on PE-2, which PE-1 learns from a BGP update from PE-2. The SRv6 tunnel to this next hop has label 524286, which is the transposed SRv6 End.DT4 SID function value 0x7fffe.

```
A:admin@PE-2# show router segment-routing-v6 local-sid

===============================================================================
Segment Routing v6 Local SIDs
===============================================================================
SID                                           Type           Function
  Locator
  Context
-------------------------------------------------------------------------------
2001:db8:a128:102:0:1000::                    End            1
  PE-2_loc_FAD128
  Base
2001:db8:a128:102:7fff:d000::                 End.DT6        524285
  PE-2_loc_FAD128
  SvcId: 6 Name: VPRN_6
2001:db8:a128:102:7fff:e000::                 End.DT4        524286
  PE-2_loc_FAD128
  SvcId: 6 Name: VPRN_6
---snip---
-------------------------------------------------------------------------------
SIDs : 6
-------------------------------------------------------------------------------
===============================================================================
```

The **show router isis 0 topology flex-algo 128** command lists the IS-IS nodes in the topology, and for each IS-IS node, the outgoing interface and the next hop. There are only IS-IS nodes at IS-IS level 2. The output of this command shows that data traffic from PE-1 to all IS-IS-enabled routers that participate in the Flex-Algorithm instance 128 (PE-2, P-3, and P-4) uses interface "int-PE-1-P-3" to P-3.

```
A:admin@PE-1# show router isis 0 topology flex-algo flex-algo-id 128

===============================================================================
Rtr Base ISIS Instance 0 Flex-Algo 128 Topology Table
===============================================================================
Node                              Interface                 Nexthop
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
IS-IS IP paths (MT-ID 0),   Level 2
-------------------------------------------------------------------------------
PE-2.00                           int-PE-1-P-3              P-3
P-3.00                            int-PE-1-P-3              P-3
P-4.00                            int-PE-1-P-3              P-3
===============================================================================
```

The applicable metric is 111, which corresponds with the sum of the static link delays that are configured on the router interfaces "int-PE-1-P-3", "int-P-3-P-4", and "int-P-4-PE-2".

The **show router isis 0 topology flex-algo 128 detail** command lists the IS-IS nodes in the topology, and for each IS-IS node, the next hop, the outgoing interface, and the metric (in this case, link delay) that applies.

```
A:admin@PE-1# show router isis 0 topology flex-algo flex-algo-id 128 detail

===============================================================================
Rtr Base ISIS Instance 0 Flex-Algo 128 Topology Table
===============================================================================
-------------------------------------------------------------------------------
IS-IS IP paths (MT-ID 0),   Level 2
-------------------------------------------------------------------------------
Node      : PE-2.00
Nexthop   : P-3
Interface : int-PE-1-P-3
SNPA      : none                          Metric      : 111

Node      : P-3.00
Nexthop   : P-3
Interface : int-PE-1-P-3
SNPA      : none                          Metric      : 1

Node      : P-4.00
Nexthop   : P-3
Interface : int-PE-1-P-3
SNPA      : none                          Metric      : 11

===============================================================================

A:admin@P-3# show router isis 0 topology flex-algo flex-algo-id 128 detail

===============================================================================
Rtr Base ISIS Instance 0 Flex-Algo 128 Topology Table
===============================================================================
-------------------------------------------------------------------------------
IS-IS IP paths (MT-ID 0),   Level 2
-------------------------------------------------------------------------------
---snip---
Node      : PE-2.00
Nexthop   : P-4
Interface : int-P-3-P-4
SNPA      : none                          Metric      : 110

Node      : P-4.00
Nexthop   : P-4
Interface : int-P-3-P-4
SNPA      : none                          Metric      : 10

===============================================================================

A:admin@P-4# show router isis 0 topology flex-algo flex-algo-id 128 detail
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Rtr Base ISIS Instance 0 Flex-Algo 128 Topology Table
===============================================================================
-------------------------------------------------------------------------------
IS-IS IP paths (MT-ID 0),   Level 2
-------------------------------------------------------------------------------
---snip---
Node      : PE-2.00
Nexthop   : PE-2
Interface : int-P-4-PE-2
SNPA      : none                              Metric      : 100
---snip---
===============================================================================
```

The IS-IS database on PE-1 contains more information, which relates to the use of the Flex-Algorithm.

There are additional link delay metrics (identical for IPv4 and IPv6) for each IS-IS-enabled router, as listed in the TE APP LINK ATTR sub-TLVs. The non-legacy Standard Application Bit Mask (SABM) flag value X indicates that they are associated with the Flex-Algorithm. Only the End SRv6 SIDs of the SRv6 locators are present.

Next to the default metric-based SPF, the Flex-Algorithm instance 128 is also in use, as listed in the SR Alg sub-TLV of the Router Cap TLVs.

The FAD sub-TLV of the PE-1 Router Cap TLV contains the delay-based definition for the Flex-Algorithm instance 128, which only router PE-1 advertises. The FAD flag value M indicates that the delay is a prefix metric.

On PE-1 and PE-2, next to the base SRv6 locator (for algorithm 0), with its End SRv6 SID, there is an additional SRv6 locator for the Flex-Algorithm instance 128, with its End SRv6 SID. This additional SRv6 locator indicates the prefix.

The **show router isis 0 database detail** command output on PE-1 is shown below, with separate entries for each IS-IS-enabled router.

For PE-1:

```
A:admin@PE-1# show router isis 0 database detail

===============================================================================
Rtr Base ISIS Instance 0 Database (detail)
===============================================================================
---snip---
Displaying Level 2 database
-------------------------------------------------------------------------------
LSP ID    : PE-1.00-00                                  Level     : L2
---snip---
TLVs :
  Area Addresses:
    Area Address : (3) 49.0001
  Supp Protocols:
    Protocols     : IPv4
    Protocols     : IPv6
  IS-Hostname   : PE-1
  Router ID   :
    Router ID   : 1.1.1.1
  TE Router ID v6  :
    Router ID   : 2001:db8::2:1
  Router Cap : 1.1.1.1, D:0, S:0
    TE Node Cap : B E M  P
    SRv6 Cap: 0x0000
    SR Alg: metric based SPF, 128
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    Node MSD Cap: BMI : 0 SRH-MAX-SL : 10 SRH-MAX-END-POP : 9 SRH-MAX-H-ENCAPS : 3 SRH-MAX-END-
D : 9
    FAD Sub-Tlv:
        Flex-Algorithm   : 128
        Metric-Type      : delay
        Calculation-Type : 0
        Priority         : 100
        Flags: M
  I/F Addresses IPv6 :
    IPv6 Address    : 2001:db8::2:1
    IPv6 Address    : 2001:db8::168:12:1
    IPv6 Address    : 2001:db8::168:13:1
    IPv6 Address    : 2001:db8::168:14:1
  TE IS Nbrs   :
    Nbr   : PE-2.00
    Default Metric  : 8
    Sub TLV Len     : 36
    IPv6 Addr : 2001:db8::168:12:1
    Nbr IPv6  : 2001:db8::168:12:2
  TE IS Nbrs   :
    Nbr   : P-3.00
    Default Metric  : 8
    Sub TLV Len     : 51
    IPv6 Addr : 2001:db8::168:13:1
    Nbr IPv6  : 2001:db8::168:13:2
    TE APP LINK ATTR    :
      SABML-flag:Non-Legacy SABM-flags:   X
        Delay Min : 1 Max : 1
  TE IS Nbrs   :
    Nbr   : P-4.00
    Default Metric  : 15
    Sub TLV Len     : 36
    IPv6 Addr : 2001:db8::168:14:1
    Nbr IPv6  : 2001:db8::168:14:2
  IPv6 Reach:
    Metric: ( I ) 0
    Prefix   : 2001:db8::2:1/128
    Metric: ( I ) 8
    Prefix   : 2001:db8::168:12:0/126
    Metric: ( I ) 8
    Prefix   : 2001:db8::168:13:0/126
    Metric: ( I ) 15
    Prefix   : 2001:db8::168:14:0/126
    Metric: ( I ) 0
    Prefix   : 2001:db8:aaaa:101::/64
  SRv6 Locator  :
    MT ID : 0
    Metric: ( ) 0 Algo:128
    Prefix   : 2001:db8:a128:101::/64
    Sub TLV   :
      End-SID   : 2001:db8:a128:101:0:1000::, flags:0x0, endpoint:End-USP
    Metric: ( ) 0 Algo:0
    Prefix   : 2001:db8:aaaa:101::/64
    Sub TLV   :
      End-SID   : 2001:db8:aaaa:101:0:1000::, flags:0x0, endpoint:End-USP
---snip---
Level (2) LSP Count : 4
-------------------------------------------------------------------------------
---snip---
SABM-flags Flags:  R = RSVP-TE
                   S = SR-TE
                   F = LFA
                   X = FLEX-ALGO
FAD-flags Flags:   M = Prefix Metric
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
```

For PE-2:

```
A:admin@PE-1# show router isis 0 database detail

===============================================================================
Rtr Base ISIS Instance 0 Database (detail)
===============================================================================
---snip---
Displaying Level 2 database
-------------------------------------------------------------------------------
---snip---
LSP ID    : PE-2.00-00                                  Level     : L2
---snip---
TLVs :
  ---snip---
  Router Cap : 2.2.2.2, D:0, S:0
    TE Node Cap : B E M  P
    SRv6 Cap: 0x0000
    SR Alg: metric based SPF, 128
    Node MSD Cap: BMI : 0 SRH-MAX-SL : 10 SRH-MAX-END-POP : 9 SRH-MAX-H-ENCAPS : 3 SRH-MAX-END-
D : 9
  ---snip---
  TE IS Nbrs   :
    Nbr   : PE-1.00
    Default Metric  : 9
    ---snip---
  TE IS Nbrs   :
    Nbr   : P-3.00
    Default Metric  : 17
    ---snip---
  TE IS Nbrs   :
    Nbr   : P-4.00
    Default Metric  : 9
    ---snip---
    TE APP LINK ATTR    :
      SABML-flag:Non-Legacy SABM-flags:   X
        Delay Min : 200 Max : 200
  ---snip---
  SRv6 Locator  :
    MT ID : 0
    Metric: ( ) 0 Algo:128
    Prefix   : 2001:db8:a128:102::/64
    Sub TLV   :
      End-SID   : 2001:db8:a128:102:0:1000::, flags:0x0, endpoint:End-USP
    Metric: ( ) 0 Algo:0
    Prefix   : 2001:db8:aaaa:102::/64
    Sub TLV   :
      End-SID   : 2001:db8:aaaa:102:0:1000::, flags:0x0, endpoint:End-USP
---snip---
===============================================================================
```

For P-3:

```
A:admin@PE-1# show router isis 0 database detail

===============================================================================
Rtr Base ISIS Instance 0 Database (detail)
===============================================================================
---snip---
Displaying Level 2 database
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
---snip---
LSP ID    : P-3.00-00                                   Level     : L2
---snip---
TLVs :
  ---snip---
  Router Cap : 3.3.3.3, D:0, S:0
    TE Node Cap : B E M  P
    SRv6 Cap: 0x0000
    SR Alg: metric based SPF, 128
    Node MSD Cap: BMI : 0 SRH-MAX-SL : 10 SRH-MAX-END-POP : 9 SRH-MAX-H-ENCAPS : 3 SRH-MAX-END-
D : 9
  ---snip---
  TE IS Nbrs   :
    Nbr   : PE-1.00
    Default Metric  : 10
    ---snip---
    TE APP LINK ATTR    :
      SABML-flag:Non-Legacy SABM-flags:   X
        Delay Min : 2 Max : 2
  TE IS Nbrs   :
    Nbr   : PE-2.00
    Default Metric  : 18
    ---snip---
  TE IS Nbrs   :
    Nbr   : P-4.00
    Default Metric  : 10
    Sub TLV Len     : 51
    ---snip---
    TE APP LINK ATTR    :
      SABML-flag:Non-Legacy SABM-flags:   X
        Delay Min : 10 Max : 10
  ---snip---
===============================================================================
```

For P-4:

```
A:admin@PE-1# show router isis 0 database detail

===============================================================================
Rtr Base ISIS Instance 0 Database (detail)
===============================================================================
---snip---
Displaying Level 2 database
-------------------------------------------------------------------------------
---snip---
LSP ID    : P-4.00-00                                   Level     : L2
---snip---
TLVs :
  ---snip---
  Router Cap : 4.4.4.4, D:0, S:0
    TE Node Cap : B E M  P
    SRv6 Cap: 0x0000
    SR Alg: metric based SPF, 128
    Node MSD Cap: BMI : 0 SRH-MAX-SL : 10 SRH-MAX-END-POP : 9 SRH-MAX-H-ENCAPS : 3 SRH-MAX-END-
D : 9
  ---snip---
  TE IS Nbrs   :
    Nbr   : P-3.00
    Default Metric  : 9
    ---snip---
    TE APP LINK ATTR    :
      SABML-flag:Non-Legacy SABM-flags:   X
        Delay Min : 20 Max : 20
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
  TE IS Nbrs   :
    Nbr   : PE-1.00
    Default Metric  : 16
    ---snip---
  TE IS Nbrs   :
    Nbr   : PE-2.00
    Default Metric  : 9
    ---snip---
    TE APP LINK ATTR    :
      SABML-flag:Non-Legacy SABM-flags:   X
        Delay Min : 100 Max : 100
  ---snip---
===============================================================================
```

PE-2 advertises to PE-1 payload prefix 172.16.222.1/32, as listed in the RIB In Entries section in the following example. PE-1 computes the applicable metric 111, which corresponds with the sum of the link delays that are configured on the router interfaces "int-PE-1-P-3", "int-P-3-P-4", and "int-P-4-PE-2". PE-1 advertises to PE-2 payload prefix 172.16.211.1/32, as listed in the RIB Out Entries section in the following example. PE-2 computes the applicable metric 222, which corresponds with the sum of the link delays that are configured on the router interfaces "int-PE-2-P-4", "int-P-4-P-3", and "int-P-3-PE-1".

The following output shows the corresponding VPN-IPv4 BGP routes on PE-1:

```
A:admin@PE-1# show router bgp routes vpn-ipv4 hunt
===============================================================================
 BGP Router ID:1.1.1.10         AS:64500       Local AS:64500
===============================================================================
---snip---
===============================================================================
BGP VPN-IPv4 Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : 172.16.222.1/32
Nexthop        : 2001:db8::2:2
Route Dist.    : 192.0.2.2:6           VPN Label      : 524286
Path Id        : None
From           : 2001:db8::2:2
Res. Nexthop   : n/a
Local Pref.    : 100                   Interface Name : int-PE-1-PE-2
Aggregator AS  : None                  Aggregator     : None
Atomic Aggr.   : Not Atomic            MED            : None
AIGP Metric    : None                  IGP Cost       : 111
Connector      : None
Community      : target:64506:6
Cluster        : No Cluster Members
Originator Id  : None                  Peer Router Id : 2.2.2.10
Fwd Class      : None                  Priority       : None
Flags          : Used Valid Best IGP
---snip---
SRv6 TLV Type  : SRv6 L3 Service TLV (5)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:a128:102::
Full Sid       : 2001:db8:a128:102:7fff:e000::
Behavior       : End.DT4 (19)
SRv6 SubSubTLV : SRv6 SID Structure (1)
---snip---
VPRN Imported  :  6

-------------------------------------------------------------------------------
RIB Out Entries
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
Network        : 172.16.211.1/32
Nexthop        : 2001:db8::2:1
Route Dist.    : 192.0.2.1:6            VPN Label      : 524286
Path Id        : None
To             : 2001:db8::2:2
Res. Nexthop   : n/a
Local Pref.    : 100                    Interface Name : NotAvailable
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : n/a
Connector      : None
Community      : target:64506:6
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 2.2.2.10
Origin         : IGP
---snip---
SRv6 TLV Type  : SRv6 L3 Service TLV (5)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:a128:101::
Full Sid       : 2001:db8:a128:101:7fff:e000::
Behavior       : End.DT4 (19)
SRv6 SubSubTLV : SRv6 SID Structure (1)
---snip---
-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

Verify the IPv6 route table on PE-1. The IPv6 route table has an additional route to the learned remote SRv6 locator for the Flex-Algorithm instance 128. This remotely configured SRv6 locator prefix of PE-2 is reached via an SRv6 tunnel.

```
A:admin@PE-1# show router route-table ipv6

===============================================================================
IPv6 Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
---snip---
2001:db8:a128:101:0:1000::/128                Local   SRV6      00h02m51s  3
       Black Hole                                                 0
2001:db8:a128:102::/64                        Remote  ISIS      00h01m57s  18
       2001:db8:a128:102::/64 (tunneled:SRV6-ISIS)                111
2001:db8:aaaa:101::/64                        Local   SRV6      00h36m30s  3
       fe80::201-"_tmnx_fpe_2.a"                                  0
2001:db8:aaaa:101:0:1000::/128                Local   SRV6      00h36m30s  3
       Black Hole                                                 0
2001:db8:aaaa:102::/64                        Remote  ISIS      00h35m22s  18
       2001:db8:aaaa:102::/64 (tunneled:SRV6-ISIS)                8
-------------------------------------------------------------------------------
No. of Routes: 16
---snip---
===============================================================================
```

Verify that the tunnel from PE-1 to the remote locator is an SRv6 tunnel that uses the "int-PE-1-P-3" interface. Perform a similar verification for the tunnel from PE-2, where the SRv6 tunnel to the remote locator uses the "int-PE-2-P-4" interface. Interface "int-PE-1-P-3" is configured on port 1/1/c2/1:1000.

```
A:admin@PE-1# show router fp-tunnel-table 1 ipv6
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
IPv6 Tunnel Table Display
---snip---
===============================================================================
Destination                                Protocol         Tunnel-ID
  Lbl/SID
    NextHop                                                 Intf/Tunnel
  Lbl/SID (backup)
    NextHop   (backup)
-------------------------------------------------------------------------------
2001:db8:a128:102::/64                     SRV6              524290
  -
    fe80::612:1ff:fe01:b-"int-PE-1-P-3"                    1/1/c2/1:1000
2001:db8:aaaa:102::/64                     SRV6              524289
  -
    fe80::60e:1ff:fe01:1-"int-PE-1-PE-2"                   1/1/c1/1:1000
-------------------------------------------------------------------------------
Total Entries : 2
-------------------------------------------------------------------------------
===============================================================================
```

For IPv6 data traffic, the VPRN routing table shows the next hop and the applicable metric for the route to 2001:db8:222::1.

```
A:admin@PE-1# show router 6 route-table 2001:db8:222::1

===============================================================================
IPv6 Route Table (Service: 6)
===============================================================================
Dest Prefix[Flags]                          Type    Proto     Age        Pref
      Next Hop[Interface Name]                                 Metric
-------------------------------------------------------------------------------
2001:db8:222::1/128                         Remote  BGP VPN   00h01m59s  170
       2001:db8:a128:102:7fff:d000:: (tunneled:SRV6)              111
-------------------------------------------------------------------------------
No. of Routes: 1
---snip---
===============================================================================
```

The next hop is the End.DT6 SRv6 SID of the SRv6 **locator** *"PE-2_loc_FAD128"* for the VPRN on PE-2, which PE-1 learns from a BGP update from PE-2. The SRv6 tunnel to this next hop has label 524285.

The applicable metric is also 111, which again corresponds with the sum of the link delays that are configured on the router interfaces "int-PE-1-P-3", "int-P-3-P-4", and "int-P-4-PE-2".

PE-2 advertises to PE-1 the information for network prefix 2001:db8:222::1/128, as listed in the RIB In Entries section in the following example. PE-1 computes the applicable metric 111, which corresponds with the sum of the link delays that are configured on the router interfaces "int-PE-1-P-3", "int-P-3-P-4", and "int-P-4-PE-2". PE-1 advertises to PE-2 payload prefix 2001:db8:211::1/128, as listed in the RIB Out Entries section in the following example. PE-2 computes the applicable metric 222, which corresponds with the sum of the link delays that are configured on the router interfaces "int-PE-2-P-4", "int-P-4-P-3", and "int-P-3-PE-1".

The following output shows the corresponding VPN-IPv6 BGP routes on PE-1:

```
A:admin@PE-1# show router bgp routes vpn-ipv6 hunt
===============================================================================
 BGP Router ID:1.1.1.10         AS:64500       Local AS:64500
===============================================================================
---snip---
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
BGP VPN-IPv6 Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : 2001:db8:222::1/128
Nexthop        : 2001:db8::2:2
Route Dist.    : 192.0.2.2:6            VPN Label      : 524285
Path Id        : None
From           : 2001:db8::2:2
Res. Nexthop   : n/a
Local Pref.    : 100                    Interface Name : int-PE-1-PE-2
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : 111
Connector      : None
Community      : target:64506:6
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 2.2.2.10
Fwd Class      : None                   Priority       : None
Flags          : Used Valid Best IGP
---snip---
SRv6 TLV Type  : SRv6 L3 Service TLV (5)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:a128:102::
Full Sid       : 2001:db8:a128:102:7fff:d000::
Behavior       : End.DT6 (18)
SRv6 SubSubTLV : SRv6 SID Structure (1)
---snip---
VPRN Imported  :  6

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
Network        : 2001:db8:211::1/128
Nexthop        : 2001:db8::2:1
Route Dist.    : 192.0.2.1:6            VPN Label      : 524285
Path Id        : None
To             : 2001:db8::2:2
Res. Nexthop   : n/a
Local Pref.    : 100                    Interface Name : NotAvailable
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : n/a
Connector      : None
Community      : target:64506:6
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 2.2.2.10
Origin         : IGP
---snip---
SRv6 TLV Type  : SRv6 L3 Service TLV (5)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:a128:101::
Full Sid       : 2001:db8:a128:101:7fff:d000::
Behavior       : End.DT6 (18)
SRv6 SubSubTLV : SRv6 SID Structure (1)
---snip---
-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Configure TE metrics on the router interfaces

For example for PE-1. A similar configuration applies for the other routers.

```
A:admin@PE-1# configure {
    router "Base" {
        mpls {
            interface "int-PE-1-PE-2" {
                te-metric 1500
                admin-state enable
            }
            interface "int-PE-1-P-3" {
                te-metric 10
                admin-state enable
            }
            interface "int-PE-1-P-4" {
                te-metric 1600
                admin-state enable
            }
            admin-state enable
        }
        rsvp {
            admin-state enable
            interface "int-PE-1-P-3" {
            }
            interface "int-PE-1-P-4" {
            }
            interface "int-PE-1-PE-2" {
            }
        }
    }
}
```

Configure a TE-metric-based flexible algorithm

Define a Flex-algorithm definition "FAD_te_metric" that uses TE metric metric type. The FAD can reside on any IS-IS-enabled router. In this example, it resides on PE-1.

```
A:admin@PE-1# configure {
    routing-options {
        flexible-algorithm-definitions {
            flex-algo "FAD_te_metric" {
                description "FAD_te_metric_based"
                metric-type te-metric
                admin-state enable
            }
        }
    }
}
```

Configure PE-1 so as to advertise the FAD "FAD_te_metric" in Flex-Algorithm instance 128 and to participate in the Flex-Algorithm instance 128.

Note: SR OS 22.10 supports multiple concurrent Flex-Algorithm instances, with different instance values between 128 and 255. While "FAD_te_metric" would be associated with a new Flex-Algorithm instance 129 in a deployment scenario, "FAD_te_metric" in the example setup replaces "FAD_delay" for Flex-Algorithm instance 128, so that, for the sake of brevity, the SRv6 locators

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

and labels that were in use earlier for the delay-based flexible algorithm scenario remain in use for the TE-metric-based flexible algorithm scenario.

```
A:admin@PE-1# configure {
    router "Base" {
        isis 0 {
            flexible-algorithms {
                admin-state enable
                flex-algo 128 {
                    advertise "FAD_te_metric"
                    participate true
                }
            }
        }
    }
}
```

Ensure that the IS-IS-enabled routers that must support the Flex-Algorithm instance participate in this Flex-Algorithm instance.

```
A:admin@PE-2#/A:admin@P-3#/A:admin@P-4# configure {
    router "Base" {
        isis 0 {
            flexible-algorithms {
                flex-algo 128 {
                    participate true
                    loopfree-alternate {
                    }
                }
                admin-state enable
            }
        }
    }
}
```

Verify data traffic

At this point, using **ping** and **traceroute** commands, verify that data traffic between the local VPRN on PE-1 and the remote VPRN on PE-2 uses the modified Flex-Algorithm.

For IPv4 data traffic:

```
A:admin@PE-1# ping 172.16.222.1 router-instance "VPRN_6"
PING 172.16.222.1 56 data bytes
64 bytes from 172.16.222.1: icmp_seq=1 ttl=64 time=2.67ms.
---snip---
---- 172.16.222.1 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.51ms, avg = 2.60ms, max = 2.72ms, stddev = 0.083ms

A:admin@PE-1# traceroute 172.16.222.1 router-instance "VPRN_6"
traceroute to 172.16.222.1, 30 hops max, 40 byte packets
  1  172.16.222.1 (172.16.222.1)    2.66 ms  2.87 ms  2.81 ms
```

For IPv6 data traffic:

```
A:admin@PE-1# ping 2001:db8:222::1 router-instance "VPRN_6"
PING 2001:db8:222::1 56 data bytes
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
64 bytes from 2001:db8:222::1 icmp_seq=1 hlim=64 time=2.05ms.
---snip---
---- 2001:db8:222::1 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.05ms, avg = 2.46ms, max = 2.70ms, stddev = 0.242ms

A:admin@PE-1# traceroute 2001:db8:222::1 router-instance "VPRN_6"
traceroute to 2001:db8:222::1, 30 hops max, 60 byte packets
  1  2001:db8:222::1 (2001:db8:222::1)    3.05 ms  2.90 ms  2.64 ms
```

For the example with the TE-metric-based flexible algorithm, the same set of **show** commands as for the example with the delay-based flexible algorithm shows that:

- the metric type changes to *te-metric* (from *delay*)
- the TE APP LINK ATTR sub-TLVs contain a **TE Metric** value
- the applicable metric for the route between the local VPRN on PE-1 and the remote VPRN on PE-2 changes to 1110 (from 111 microseconds), corresponding with the sum of the metric values along the path PE-1, P-3, P-4, PE-2, as shown in the following **show router isis 0 topology flex-algo 128 detail** command output examples.
  On PE-1:

```
A:admin@PE-1# show router isis 0 topology flex-algo flex-algo-id 128 detail

===============================================================================
Rtr Base ISIS Instance 0 Flex-Algo 128 Topology Table
===============================================================================
-------------------------------------------------------------------------------
IS-IS IP paths (MT-ID 0),   Level 2
-------------------------------------------------------------------------------
Node      : PE-2.00
Nexthop   : P-3
Interface : int-PE-1-P-3
SNPA      : none                            Metric      : 1110

Node      : P-3.00
Nexthop   : P-3
Interface : int-PE-1-P-3
SNPA      : none                            Metric      : 10

Node      : P-4.00
Nexthop   : P-3
Interface : int-PE-1-P-3
SNPA      : none                            Metric      : 110

===============================================================================
```

On P-3:

```
A:admin@P-3# show router isis 0 topology flex-algo flex-algo-id 128 detail

===============================================================================
Rtr Base ISIS Instance 0 Flex-Algo 128 Topology Table
===============================================================================
-------------------------------------------------------------------------------
IS-IS IP paths (MT-ID 0),   Level 2
-------------------------------------------------------------------------------
---snip---
Node      : PE-2.00
Nexthop   : P-4
Interface : int-P-3-P-4
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
SNPA      : none                              Metric      : 1100

Node      : P-4.00
Nexthop   : P-4
Interface : int-P-3-P-4
SNPA      : none                              Metric      : 100

===============================================================================
```

On P-4:

```
A:admin@P-4# show router isis 0 topology flex-algo flex-algo-id 128 detail

===============================================================================
Rtr Base ISIS Instance 0 Flex-Algo 128 Topology Table
===============================================================================
-------------------------------------------------------------------------------
IS-IS IP paths (MT-ID 0),   Level 2
-------------------------------------------------------------------------------
---snip---
Node      : PE-2.00
Nexthop   : PE-2
Interface : int-P-4-PE-2
SNPA      : none                              Metric      : 1000
---snip---
===============================================================================
```

Conclusion

The Flex-Algorithm for SRv6-based VPRNs feature allows the computation of constraint-based paths across an SRv6-enabled network, based on metrics other than the default IGP metrics. This allows carrying data traffic over an end-to-end path that is optimized using the best suited metric IGP, delay, or TE).

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Inter-AS VPRN Model B

This chapter describes the Inter-AS VPRN Model B.

Topics in this chapter include:

Applicability

This chapter was initially written for SR OS Release 15.0.R8, but the MD-CLI in the current edition is based on SR OS Release 22.2.R1.

Overview

An inter-AS Virtual Private Routed Network (VPRN) contains sites that are connected to different Autonomous Systems (ASs). Inter-AS is typically used either to provide extended reach through a partnership/trust agreement, as an interim means to interconnect ASs following acquisition, or because of the internal organization of a single Service Provider (SP). Three models for interconnecting ASs are defined in RFC 4364, labeled model A, B, and C. This chapter describes model B.

Inter-AS VPRN model B encompasses EBGP redistributing VPN-IPv4 and VPN-IPv6 routes between neighboring ASs. An Autonomous System Border Router (ASBR) learns VPN routes from within its AS using IBGP, potentially as a client of a Route Reflector (RR), then uses EBGP to redistribute those labeled VPN routes to its adjacent ASBR.

When redistributing the routes into EBGP, the ASBR imposes next-hop-self on the VPN-IPv4 and VPN-IPv6 update messages and generates its own label value when it advertises the update message upstream. Therefore, the ASBR programs a label-swap entry in its FIB and forwards traffic to the neighboring ASBR using a single-level label stack (the VPN label).

A key property of model B is that it eliminates the need for per-VPRN configuration on the ASBRs. However, both ASBRs must have a mechanism to implicitly learn all VPN prefixes within their local AS and selectively advertise some of those prefixes to the neighboring ASBR.

Figure 304: Inter-AS VPRN Model B control and data plane example shows an example of the control plane and corresponding data plane used in model B, where MPLS is used for transport in both ASs. CE-1 is attached to PE-1 in AS 64496 and advertises prefix 172.31.100.0/24, which is propagated between neighboring ASBRs to PE-2 in AS 64510 and upstream to CE-2.

The IP traffic originating from CE-2 and received by PE-2 is received on the VRF interface of VPRN 100 and encapsulated using a two-level label stack; the inner label is the VPN label (300) and the outer label is the LDP transport label used for reaching the local ASBR-2.

ASBR-2 passes the traffic to ASBR-1, removing the LDP transport label and swapping the VPN label (300) with its VPN label (200), resulting in a single-level label stack.

In turn, ASBR-1 swaps the received VPN label (200) with another VPN label (100) and adds an LDP transport label to reach PE-1.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Finally, PE-1 removes the VPN label and delivers the unlabeled IP traffic to CE-1.

Figure 304: Inter-AS VPRN Model B control and data plane example

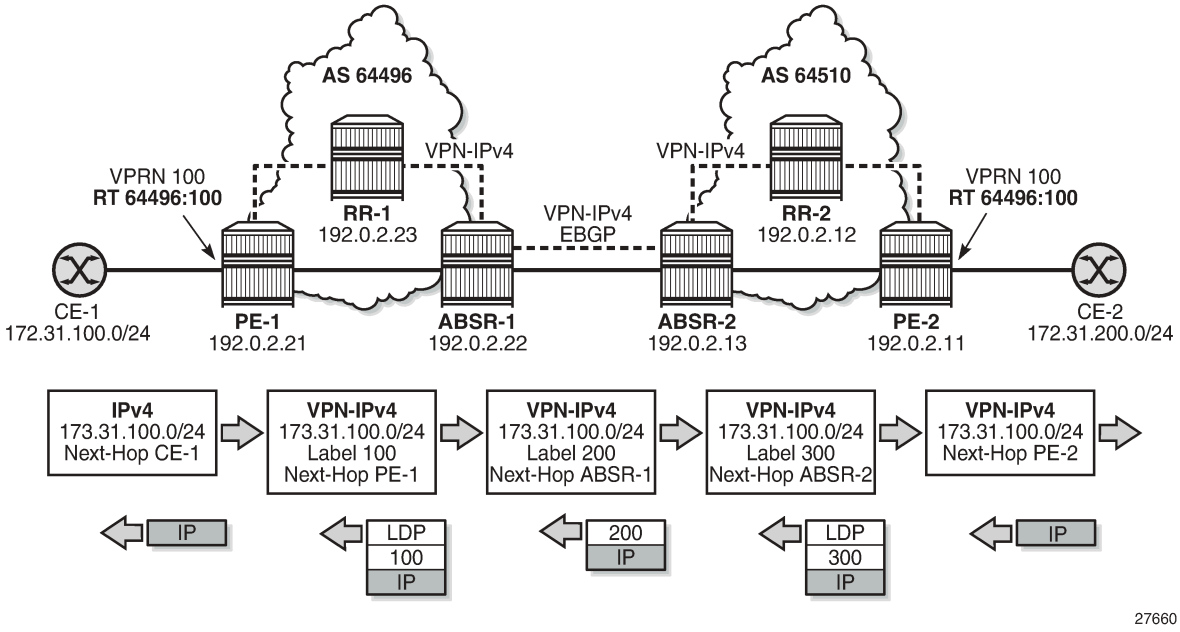

Configuration

In the example shown in Figure 305: Inter-AS VPRN Model B topology, IS-IS is configured in each AS, and MP-IBGP sessions are established between the PEs and the RRs in AS 64496 and 64510, PE-3 and PE-7, respectively. LDP and RSVP-TE is used for transport in AS 64496, whereas AS 64510 uses LDP for its transport. An MP-EBGP session is established between ASBR PE-4 and ASBR PE-8.

Figure 305: Inter-AS VPRN Model B topology

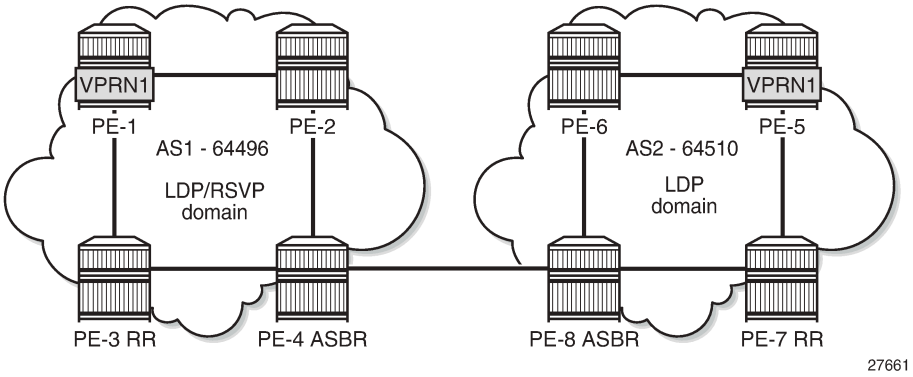

The initial configuration includes:

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- Cards, MDAs, and ports
- Router interfaces
- IS-IS as IGP on all interfaces (alternatively, OSPF can be used), with traffic engineering enabled
- LDP and RSVP-TE configured in AS 64496, LDP configured in AS 64510
- IBGP configured in AS 64496, with PE-3 as RR for clients PE-1, PE-2, and PE-4
- IBGP configured in AS 64510, with PE-7 as RR for clients PE-5, PE-6, and PE-8

Model B configuration

There are no specific requirements on PE routers or RRs for enabling inter-AS VPRN model B; only specific configurations are required on the ASBRs.

First, an ASBR must learn the VPN-IPv4 and VPN-IPv6 routes from the local AS and export these routes to the neighbor AS over an MP-EBGP session. This is achieved on each ASBR by declaring an IBGP group for peering with the local RR, and declaring an EBGP group for peering with the neighboring AS. The IBGP and EBGP groups have included the address family **vpn-ipv4**, **vpn-ipv6**, or both.

Additionally, import and export policies can be used to control the VPN-IPv4 and VPN-IPv6 routes exchanged. The latter requires the **vpn-apply-import** and **vpn-apply-export** commands for SR OS to match the prefixes of the VPN-IPv4 and VPN-IPv6 address families.

The use of the **next-hop-resolution** command is described in the Service configuration section. The BGP configuration on ASBR PE-4 is as follows:

```
# on ASBR PE-4:
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            loop-detect discard-route
            inter-as-vpn true                   # required for inter-as VPRN model B
            split-horizon true
            next-hop-resolution
                labeled-routes
                    transport-tunnel
                        family vpn
                            resolution filter
                            resolution-filter
                                ldp true        # by default enabled for VPN routes
                                rsvp true
                            }
                        }
                    }
                }
            }
            group "vpn-eBGP" {
            }
            group "vpn-iBGP" {
                peer-as 64496
            }
            neighbor "192.0.2.3" {
                group "vpn-iBGP"
                family {
                    vpn-ipv4 true
                    vpn-ipv6 true
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            neighbor "192.168.48.2" {
                group "vpn-eBGP"
                peer-as 64510
                family {
                    vpn-ipv4 true
                    vpn-ipv6 true
                }
                ebgp-default-reject-policy {
                    import false
                    export false
                }
            }
```

The configuration on ASBR PE-8 is similar.

Second, the **inter-as-vpn true** command enables the inter-AS functionality and causes the ASBR to store the received VPN-IPv4 routes in its RIB-In, even though it has no VRF that imports these routes. For a route to be considered valid, the ASBR still needs to resolve the next-hop of this route to a tunnel. The **inter-as-vpn true** command will also change the BGP next-hop of advertised and received VPN-IPv4/ VPN-IPv6 routes. When a route is advertised to an EBGP peer, the BGP next-hop is changed to the local-address used for communicating with the EBGP peer. When a route is received from an EBGP peer and advertised to an IBGP peer, the BGP next-hop is changed to the local-address used for communicating with the IBGP peer.

The configuration of the MP-EBGP session between the ASBRs in the EBGP group allows the ASBR to forward labeled packets over its connection with its peer ASBR.

MPLS LSP configuration

Two LSPs are needed between the end-to-end PEs (PE-1 and PE-5) to exchange service traffic bidirectionally, because LSPs are unidirectional. In AS 64496, this is achieved by configuring a first LSP from the service PE (PE-1) to the local ASBR (PE-4), and a second LSP back from the local ASBR (PE-4) toward the service PE (PE-1). In AS 64510, LDP is enabled on all interfaces; no RSVP LSPs are used.

In AS 64496, both LDP and RSVP are enabled. The LSP (and path) from PE-1 to PE-4 runs via PE-3, as follows:

```
# on PE-1:
configure {
    router "Base" {
        ldp {
            interface-parameters {
                interface "int-PE-1-PE-2" {
                    ipv4 {
                    }
                }
                interface "int-PE-1-PE-3" {
                    ipv4 {
                    }
                }
            }
        }
        mpls {
            admin-state enable
            interface "int-PE-1-PE-2" {
            }
            interface "int-PE-1-PE-3" {
            }
            interface "system" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
            path "path-PE-1-PE-3-PE-4" {
                admin-state enable
                hop 10 {
                    ip-address 192.168.13.2
                    type strict
                }
                hop 20 {
                    ip-address 192.168.34.2
                    type strict
                }
            }
            lsp "lsp-PE-1-PE-4" {
                admin-state enable
                type p2p-rsvp
                to 192.0.2.4
                primary "path-PE-1-PE-3-PE-4" {
                }
            }
        }
        rsvp {
            admin-state enable
            interface "int-PE-1-PE-2" {
            }
            interface "int-PE-1-PE-3" {
            }
        }
```

The LSP (and path) from PE-4 to PE-3 also runs via PE-3, as follows:

```
# on ASBR PE-4:
configure {
    router "Base" {
        ldp {
            interface-parameters {
                interface "int-PE-4-PE-2" {
                    ipv4 {
                    }
                }
                interface "int-PE-4-PE-3" {
                    ipv4 {
                    }
                }
            }
        }
        mpls {
            admin-state enable
            interface "int-PE-4-PE-2" {
            }
            interface "int-PE-4-PE-3" {
            }
            interface "system" {
            }
            path "path-PE-4-PE-3-PE-1" {
                admin-state enable
                hop 10 {
                    ip-address 192.168.34.1
                    type strict
                }
                hop 20 {
                    ip-address 192.168.13.1
                    type strict
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
            lsp "lsp-PE-4-PE-1" {
                admin-state enable
                type p2p-rsvp
                to 192.0.2.1
                primary "path-PE-4-PE-3-PE-1" {
                }
            }
        }
        rsvp {
            admin-state enable
            interface "int-PE-4-PE-2" {
            }
            interface "int-PE-4-PE-3" {
            }
        }
```

Service configuration

VPRN 1 is configured on PE-1 and PE-5. Although the VPRN service IDs used in both ASs do not need to match, in an inter-AS VPRN model B context, the route targets (RTs) used in both ASs must be coordinated. The RT exported by the PE-1 VPRN 1 must be imported by the PE-5 VPRN 1, and vice versa. In this example, specific **import-community** and **export-community** values are used for VPRN 1; the simplified method using a single **vrf-target community** is used for VPRN 33.

To carry the customer data across AS 64496, tunnels must bind to a VPRN service with the **auto-bind-tunnel** command. Resolution is set to filter, indicating that SR OS must select a tunnel using the information defined in the **resolution-filter** context. The keywords **ldp true** and **rsvp true** in the resolution-filter context indicate that LDP or RSVP tunnels can be used, but SR OS prefers the RSVP tunnels because the preference for RSVP (7) is lower than the preference for LDP (9).

In AS 64496, the VPRN service on PE-1 is defined as follows:

```
# on PE-1:
configure {
    service {
        vprn "VPRN1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:1"
                    vrf-target {
                        import-community "target:64510:1"
                        export-community "target:64496:1"
                    }
                    auto-bind-tunnel {
                        resolution filter
                        resolution-filter {
                            ldp true
                            rsvp true
                        }
                    }
                }
            }
            interface "int-S1-1" {
                ipv4 {
                    primary {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        address 10.1.10.1
                        prefix-length 24
                    }
                }
                sap 1/2/1:1 {
                }
                ipv6 {
                    primary {
                        address 2001:db8:1::1:1
                        prefix-length 120
                    }
                }
            }
            interface "int-S1-2" {
                loopback true
                ipv4 {
                    primary {
                        address 10.1.11.1
                        prefix-length 24
                    }
                }
            }
        }
```

In AS 64510, the transport technology is LDP only, so the VPRN service in PE-5 auto-binds using LDP LSPs in the tunnel table to resolve VPN-IPv4 and VPN-IPv6 routes for which the **vrf-target import-community** matches the **vrf-target export-community** value configured in PE-1 and vice versa, as follows:

```
# on PE-5 in AS 64510:
configure {
    service {
        vprn "VPRN1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64510:1"
                    vrf-target {
                        import-community "target:64496:1"
                        export-community "target:64510:1"
                    }
                    auto-bind-tunnel {
                        resolution filter
                        resolution-filter {
                            ldp true
                        }
                    }
                }
            }
            interface "int-S1-1" {
                ipv4 {
                    primary {
                        address 10.1.50.1
                        prefix-length 24
                    }
                }
                sap 1/2/1:1 {
                }
                ipv6 {
                    primary {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        address 2001:db8:1::5:1
                        prefix-length 120
                    }
                }
            }
            interface "int-S1-2" {
                loopback true
                ipv4 {
                    primary {
                        address 10.1.51.1
                        prefix-length 24
                    }
                }
            }
        }
```

A second service is defined on PE-1 and PE-2 (VPRN 33), using loopback addresses 10.33.1.1/32 and 10.33.2.1/32 in PE-1 and PE-2, respectively. These addresses might appear in traces and commands later, but are of no concern because these are used for transporting intra-AS traffic. On PE-1, VPRN 33 is configured as follows:

```
# on PE-1:
configure {
    service {
        vprn "VPRN33" {
            admin-state enable
            service-id 33
            customer "1"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:33"
                    vrf-target {
                        community "target:64496:33"
                    }
                    auto-bind-tunnel {
                        resolution filter
                        resolution-filter {
                            ldp true
                        }
                    }
                }
            }
            interface "int-S1-1" {
                loopback true
                ipv4 {
                    primary {
                        address 10.33.1.1
                        prefix-length 24
                    }
                }
            }
        }
```

For VPRN 1 service traffic to flow in the direction from PE-5 to PE-1, ASBR PE-4 in AS 64496 must offer the possibility to use RSVP-TE tunnels when resolving a BGP next-hop for VPN services. Therefore, ASBR PE-4 must be explicitly configured, as follows:

```
# on ASBR PE-4:
configure {
    router "Base" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        autonomous-system 64496
        bgp {
            next-hop-resolution
                labeled-routes
                    transport-tunnel
                        family vpn
                            resolution filter
                            resolution-filter
                                # ldp true    # LDP by default enabled
                                rsvp true
                            }
                        }
                    }
                }
            }
```

On ASBR PE-8 in AS 64510, no explicit configuration is required because resolving a BGP next-hop for VPN service to LDP tunnels is the default behavior.

Verification

With the configurations from previous sections applied, PE-1 receives three VPN-IPv4 routes and one VPN-IPv6 route, as follows:

```
[/]
A:admin@PE-1# show router bgp summary all

===============================================================================
BGP Summary
===============================================================================
Legend : D - Dynamic Neighbor
===============================================================================
Neighbor
Description
ServiceId          AS PktRcvd InQ  Up/Down   State|Rcv/Act/Sent (Addr Family)
                      PktSent OutQ
-------------------------------------------------------------------------------
192.0.2.3
Def. Inst       64496     421    0 03h27m16s 3/3/3 (VpnIPv4)
                          424    0           1/1/1 (VpnIPv6)
-------------------------------------------------------------------------------
```

PE-1 received the following three VPN-IPv4 routes:

```
[/]
A:admin@PE-1# show router bgp routes vpn-ipv4
===============================================================================
 BGP Router ID:192.0.2.1        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv4 Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop (Router)                                 Path-Id     IGP Cost
      As-Path                                                      Label
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
u*>i  64496:33:10.33.2.0/24                              100         None
      192.0.2.2                                          None        10
      No As-Path                                                     524283
u*>i  64510:1:10.1.50.0/24                               100         None
      192.0.2.4                                          None        20
      64510                                                          524279
u*>i  64510:1:10.1.51.0/24                               100         None
      192.0.2.4                                          None        20
      64510                                                          524279
-------------------------------------------------------------------------------
Routes : 3
===============================================================================
```

PE-1 received the following VPN-IPv6 route:

```
[/]
A:admin@PE-1# show router bgp routes vpn-ipv6
===============================================================================
 BGP Router ID:192.0.2.1        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv6 Routes
===============================================================================
Flag  Network                                            LocalPref   MED
      Nexthop (Router)                                   Path-Id     IGP Cost
      As-Path                                                        Label
-------------------------------------------------------------------------------
u*>i  64510:1:2001:db8:1::5:0/120                        100         None
      ::ffff:192.0.2.4                                   None        20
      64510                                                          524278
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

PE-1 has three LDP tunnels and one RSVP tunnel, and its tunnel table looks as follows:

```
[/]
A:admin@PE-1# show router tunnel-table

===============================================================================
IPv4 Tunnel Table (Router: Base)
===============================================================================
Destination           Owner     Encap TunnelId  Pref   Nexthop        Metric
   Color
-------------------------------------------------------------------------------
192.0.2.2/32          ldp       MPLS  65537     9      192.168.12.2   10
192.0.2.3/32          ldp       MPLS  65538     9      192.168.13.2   10
192.0.2.4/32          rsvp      MPLS  1         7      192.168.13.2   16777215
192.0.2.4/32          ldp       MPLS  65539     9      192.168.12.2   20
-------------------------------------------------------------------------------
Flags: B = BGP or MPLS backup hop available
       L = Loop-Free Alternate (LFA) hop available
       E = Inactive best-external BGP route
       k = RIB-API or Forwarding Policy backup hop
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The IPv4 routing table for VPRN 1 is as follows:

```
[/]
A:admin@PE-1# show router 1 route-table

===============================================================================
Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
10.1.10.0/24                                    Local   Local     00h02m15s  0
       int-S1-1                                                        0
10.1.11.0/24                                    Local   Local     00h02m15s  0
       int-S1-2                                                        0
10.1.50.0/24                                    Remote  BGP VPN   00h01m01s  170
       192.0.2.4 (tunneled:RSVP:1)                                     16777215
10.1.51.0/24                                    Remote  BGP VPN   00h01m01s  170
       192.0.2.4 (tunneled:RSVP:1)                                     16777215
-------------------------------------------------------------------------------
No. of Routes: 4
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

The IPv4 addresses for VPRN 1 on PE-8 are 10.1.50.0/24 and 10.1.51.0/24, and are reachable through RSVP-TE tunnel 1 (tunneled:RSVP:1). The VPN label value for these prefixes is assigned and advertised by ASBR PE-4 and gets to PE-1 via the RR PE-3 in an MP-BGP update message. The 10.33.2.0/24 prefix belongs to a different service and is not relevant for model B because it is used for intra-AS traffic. The VPN-IPv4 routes received on PE-1 are as follows:

```
[/]
A:admin@PE-1# show router bgp neighbor 192.0.2.3 received-routes vpn-ipv4
===============================================================================
 BGP Router ID:192.0.2.1        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv4 Routes
===============================================================================
Flag  Network                                           LocalPref   MED
      Nexthop (Router)                                  Path-Id     IGP Cost
      As-Path                                                       Label
-------------------------------------------------------------------------------
u*>i  64496:33:10.33.2.0/24                             100         None
      192.0.2.2                                         None        10
      No As-Path                                                    524283
u*>i  64510:1:10.1.50.0/24                              100         None
      192.0.2.4                                         None        20
      64510                                                         524279
u*>i  64510:1:10.1.51.0/24                              100         None
      192.0.2.4                                         None        20
      64510                                                         524279
-------------------------------------------------------------------------------
Routes : 3
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
```

The BGP next-hops for the VPN-IPv4 BGP address family are as follows. Service traffic for VPRN 33 uses the LDP tunnel to PE-2 carrying the intra-AS traffic, and service traffic for VPRN 1 uses the RSVP tunnel to ASBR PE-4 carrying the inter-AS traffic.

```
[/]
A:admin@PE-1# show router bgp next-hop vpn-ipv4
===============================================================================
 BGP Router ID:192.0.2.1        AS:64496       Local AS:64496
===============================================================================

===============================================================================
BGP VPN Next Hop
===============================================================================
VPN Next Hop                                                  Owner
   Autobind                                          FibProg  Reason
   Labels (User-labels)                              FlexAlgo Metric
   Admin-tag-policy (strict-tunnel-tagging)
-------------------------------------------------------------------------------
192.0.2.2                                                     LDP
   ldp bgp                                              Y
   -- (2)                                             --       10
   -- (-)
192.0.2.4                                                     RSVP
   ldp rsvp bgp                                         Y
   -- (2)                                             --       16777215
   -- (-)
-------------------------------------------------------------------------------
Next Hops : 2
===============================================================================
```

The IPv6 routing table for VPRN 1 is as follows:

```
[/]
A:admin@PE-1# show router 1 route-table ipv6

===============================================================================
IPv6 Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
2001:db8:1::1:0/120                           Local   Local     00h03m54s  0
       int-S1-1                                                     0
2001:db8:1::5:0/120                           Remote  BGP VPN   00h02m40s  170
       192.0.2.4 (tunneled:RSVP:1)                                  16777215
-------------------------------------------------------------------------------
No. of Routes: 2
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

The VPN-IPv6 routes received on PE-1 are as follows:

```
[/]
A:admin@PE-1# show router bgp neighbor 192.0.2.3 received-routes vpn-ipv6
===============================================================================
 BGP Router ID:192.0.2.1        AS:64496       Local AS:64496
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv6 Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop (Router)                                 Path-Id     IGP Cost
      As-Path                                                      Label
-------------------------------------------------------------------------------
u*>i  64510:1:2001:db8:1::5:0/120                      100         None
      ::ffff:192.0.2.4                                 None        20
      64510                                                        524278
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

The BGP next-hop for the VPN-IPv6 address family is as follows:

```
[/]
A:admin@PE-1# show router bgp next-hop vpn-ipv6
===============================================================================
 BGP Router ID:192.0.2.1        AS:64496       Local AS:64496
===============================================================================

===============================================================================
BGP VPN Next Hop
===============================================================================
VPN Next Hop                                                   Owner
   Autobind                                          FibProg   Reason
   Labels (User-labels)                              FlexAlgo  Metric
   Admin-tag-policy (strict-tunnel-tagging)
-------------------------------------------------------------------------------
::ffff:192.0.2.4                                               RSVP
   ldp rsvp bgp                                       Y
   -- (2)                                             --        16777215
   -- (-)
-------------------------------------------------------------------------------
Next Hops : 1
===============================================================================
```

The forwarding plane is programmed accordingly, as follows:

```
[/]
A:admin@PE-1# show router 1 fib 1 ipv4

===============================================================================
FIB Display
===============================================================================
Prefix [Flags]                                                 Protocol
  NextHop
-------------------------------------------------------------------------------
10.1.10.0/24                                                   LOCAL
  10.1.10.0 (int-S1-1)
10.1.11.0/24                                                   LOCAL
  10.1.11.0 (int-S1-2)
10.1.50.0/24                                                   BGP_VPN
  192.0.2.4 (VPRN Label:524279 Transport:RSVP LSP:1)
10.1.51.0/24                                                   BGP_VPN
  192.0.2.4 (VPRN Label:524279 Transport:RSVP LSP:1)
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
Total Entries : 4
-------------------------------------------------------------------------------
===============================================================================
```

```
[/]
A:admin@PE-1# show router 1 fib 1 ipv6

===============================================================================
FIB Display
===============================================================================
Prefix [Flags]                                          Protocol
  NextHop
-------------------------------------------------------------------------------
2001:db8:1::1:0/120                                     LOCAL
  2001:db8:1::1:0 (int-S1-1)
2001:db8:1::5:0/120                                     BGP_VPN
  192.0.2.4 (VPRN Label:524278 Transport:RSVP LSP:1)
-------------------------------------------------------------------------------
Total Entries : 2
-------------------------------------------------------------------------------
===============================================================================
```

SR OS uses a label-per-VRF mode of label distribution, meaning that the same label is used for different VPN-IPv4 and different VPN-IPv6 prefixes from the same VRF, which saves on MPLS label resources. In this example, the VPRN service label is 524279 for the VPN-IPv4 prefixes 10.1.50.0/24 and 10.1.51.0/24, and 524278 for VPN-IPv6 prefix 2001:db8:1::5:0/120.

The forwarding plane is also programmed with the outer label to be used for transport purposes. Two labels are present: 524282 assigned through RSVP, and 524284 assigned through LDP. Because RSVP takes precedence over LDP, the RSVP label is actively used, as follows:

```
[/]
A:admin@PE-1# show router fp-tunnel-table 1 192.0.2.4/32

===============================================================================
IPv4 Tunnel Table Display

Legend:
label stack is ordered from bottom-most to top-most
B - FRR Backup
===============================================================================
Destination                             Protocol         Tunnel-ID
  Lbl/SID
    NextHop                                              Intf/Tunnel
  Lbl/SID (backup)
    NextHop   (backup)
-------------------------------------------------------------------------------
192.0.2.4/32                            LDP               -
  524284
    192.168.12.2                                          1/1/1:1000
192.0.2.4/32                            RSVP              1
  524282
    192.168.13.2                                          1/1/2:1000
-------------------------------------------------------------------------------
Total Entries : 2
-------------------------------------------------------------------------------
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Traffic over VPRN 1 is generated using a ping command on PE-1 to the remote loopback address, as follows:

```
[/]
A:admin@PE-1# ping 10.1.50.1 router-instance "VPRN1"
PING 10.1.50.1 56 data bytes
64 bytes from 10.1.50.1: icmp_seq=1 ttl=64 time=5.55ms.
64 bytes from 10.1.50.1: icmp_seq=2 ttl=64 time=6.38ms.
64 bytes from 10.1.50.1: icmp_seq=3 ttl=64 time=5.91ms.
64 bytes from 10.1.50.1: icmp_seq=4 ttl=64 time=5.75ms.
64 bytes from 10.1.50.1: icmp_seq=5 ttl=64 time=5.43ms.

---- 10.1.50.1 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 5.43ms, avg = 5.80ms, max = 6.38ms, stddev = 0.331ms
```

On PE-1, the IPv4 VPRN 1 traffic is pushed with VPN label 524279, followed by RSVP-TE transport label 524282. ASBR PE-4 removes the RSVP-TE transport label and swaps the internal (advertised) VPN label 524279 with the external VPN label 524280 received from ASBR PE-8. For IPv6 VPRN 1 traffic, VPN label 524278 is swapped by VPN label 524279. The inter-AS BGP labels stored by ASBR PE-4 are as follows:

```
[/]
A:admin@PE-4# show router bgp inter-as-label

===============================================================================
BGP Inter-AS labels
Flags: B - entry has backup, P - entry is promoted
===============================================================================
NextHop                       Received       Advertised     Label
                              Label          Label          Origin
-------------------------------------------------------------------------------
192.0.2.1                     524281         524282         Internal
192.0.2.1                     524282         524281         Internal
192.0.2.1                     524282         524280         Internal
192.0.2.2                     524283         524277         Internal
192.168.48.2                  524279         524278         External
192.168.48.2                  524280         524279         External
-------------------------------------------------------------------------------
Total Labels allocated:   6
===============================================================================
```

The forward data flow (from AS 64496 to AS 64510) for VPRN 1 uses the labels for which the label origin is external. The VPN labels used for the backward data flow (from AS 64510 to 64496) uses the labels for which the label origin is internal.

For brevity, the commands to display and check VPN prefixes and labels used in AS 64510 are omitted.

By disabling (**admin-state disable**) both RSVP LSPs between PE-1 and ASBR PE-4 in AS 64496, both PE-1 and PE-4 will select LDP tunnels for resolving VPN BGP next-hops. Then, the route table for VPRN 1 is as follows, where **tunneled** indicates an LDP tunnel is used to reach the next hop:

```
[/]
A:admin@PE-1# show router 1 route-table

===============================================================================
Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
10.1.10.0/24                                  Local   Local     04h29m37s  0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
       int-S1-1                                                          0
10.1.11.0/24                                Local   Local     04h29m37s  0
       int-S1-2                                                          0
10.1.50.0/24                                Remote  BGP VPN   00h00m08s  170
       192.0.2.4 (tunneled)                                              20
10.1.51.0/24                                Remote  BGP VPN   00h00m08s  170
       192.0.2.4 (tunneled)                                              20
-------------------------------------------------------------------------------
No. of Routes: 4
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

Only LDP tunnels are available in PE-1 and ASBR PE-4, as follows:

```
[/]
A:admin@PE-1# show router tunnel-table

===============================================================================
IPv4 Tunnel Table (Router: Base)
===============================================================================
Destination           Owner     Encap TunnelId  Pref   Nexthop        Metric
   Color
-------------------------------------------------------------------------------
192.0.2.2/32          ldp       MPLS  65537     9      192.168.12.2   10
192.0.2.3/32          ldp       MPLS  65538     9      192.168.13.2   10
192.0.2.4/32          ldp       MPLS  65539     9      192.168.12.2   20
-------------------------------------------------------------------------------
Flags: B = BGP or MPLS backup hop available
       L = Loop-Free Alternate (LFA) hop available
       E = Inactive best-external BGP route
       k = RIB-API or Forwarding Policy backup hop
===============================================================================
```

```
[/]
A:admin@PE-4# show router tunnel-table

===============================================================================
IPv4 Tunnel Table (Router: Base)
===============================================================================
Destination           Owner     Encap TunnelId  Pref   Nexthop        Metric
   Color
-------------------------------------------------------------------------------
192.0.2.1/32          ldp       MPLS  65538     9      192.168.24.1   20
192.0.2.2/32          ldp       MPLS  65539     9      192.168.24.1   10
192.0.2.3/32          ldp       MPLS  65537     9      192.168.34.1   10
-------------------------------------------------------------------------------
Flags: B = BGP or MPLS backup hop available
       L = Loop-Free Alternate (LFA) hop available
       E = Inactive best-external BGP route
       k = RIB-API or Forwarding Policy backup hop
===============================================================================
```

The BGP next-hop for VPN-IPv4 traffic in PE-1 also indicates that, to reach PE-5 via PE-4, an LDP tunnel is used, as follows:

```
[/]
A:admin@PE-1# show router bgp next-hop vpn-ipv4
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
 BGP Router ID:192.0.2.1        AS:64496       Local AS:64496
===============================================================================

===============================================================================
BGP VPN Next Hop
===============================================================================
VPN Next Hop                                                    Owner
   Autobind                                          FibProg  Reason
   Labels (User-labels)                              FlexAlgo Metric
   Admin-tag-policy (strict-tunnel-tagging)
-------------------------------------------------------------------------------
192.0.2.2                                                       LDP
   ldp bgp                                              Y
   -- (2)                                            --       10
   -- (-)
192.0.2.4                                                       LDP
   ldp rsvp bgp                                         Y
   -- (2)                                            --       20
   -- (-)
-------------------------------------------------------------------------------
Next Hops : 2
===============================================================================
```

The forwarding plane is programmed accordingly, as follows:

```
[/]
A:admin@PE-1# show router 1 fib 1 ipv4

===============================================================================
FIB Display
===============================================================================
Prefix [Flags]                                                  Protocol
  NextHop
-------------------------------------------------------------------------------
10.1.10.0/24                                                    LOCAL
  10.1.10.0 (int-S1-1)
10.1.11.0/24                                                    LOCAL
  10.1.11.0 (int-S1-2)
10.1.50.0/24                                                    BGP_VPN
  192.0.2.4 (VPRN Label:524279 Transport:LDP)
10.1.51.0/24                                                    BGP_VPN
  192.0.2.4 (VPRN Label:524279 Transport:LDP)
-------------------------------------------------------------------------------
Total Entries : 4
-------------------------------------------------------------------------------
===============================================================================
```

```
[/]
A:admin@PE-1# show router fp-tunnel-table 1

===============================================================================
IPv4 Tunnel Table Display

Legend:
label stack is ordered from bottom-most to top-most
B - FRR Backup
===============================================================================
Destination                                Protocol         Tunnel-ID
  Lbl/SID
    NextHop                                                 Intf/Tunnel
  Lbl/SID (backup)
    NextHop   (backup)
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
192.0.2.2/32                                LDP              -
  524287
    192.168.12.2                                           1/1/1:1000
192.0.2.3/32                                LDP              -
  524287
    192.168.13.2                                           1/1/2:1000
192.0.2.4/32                                LDP              -
  524284
    192.168.12.2                                           1/1/1:1000
-------------------------------------------------------------------------------
Total Entries : 3
-------------------------------------------------------------------------------
===============================================================================
```

The details for the LDP tunnel from PE-1 to PE-4 are as follows:

```
[/]
A:admin@PE-1# show router tunnel-table 192.0.2.4/32 detail

===============================================================================
Tunnel Table (Router: Base)
===============================================================================
Destination      : 192.0.2.4/32
NextHop          : 192.168.12.2
Tunnel Flags     : (Not Specified)
Age              : 00h12m46s
CBF Classes      : (Not Specified)
Owner            : ldp                  Encap            : MPLS
Tunnel ID        : 65539                Preference       : 9
Tunnel Label     : 524284               Tunnel Metric    : 20
Tunnel MTU       : 1556                 Max Label Stack  : 1
-------------------------------------------------------------------------------
Number of tunnel-table entries          : 1
Number of tunnel-table entries with LFA : 0
===============================================================================
```

On PE-1, the IPv4 traffic in VPRN 1 is pushed with VPN label 524279, followed by LDP transport label 524284. ASBR PE-4 removes the LDP transport label and swaps the internal (advertised) VPN label 524279 with the external VPN label 524280 received from ASBR PE-8. The inter-AS label mapping between the ASBRs remains unchanged.

On the directly connected interface between the ASBRs, nothing has changed; only a single MPLS label is used to carry the VPN data, as shown in the following capture:

With this configuration, all the VPN-IPv4 and VPN-IPv6 routes known to AS 64496 are advertised by ASBR PE-4 to AS 64510, even the VPN-IPv4 and VPN-IPv6 routes from other AS 64496 VPRN services that do not need to be distributed:

```
[/]
A:admin@PE-4# show router bgp neighbor 192.168.48.2 advertised-routes vpn-ipv4 brief
===============================================================================
 BGP Router ID:192.0.2.4        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv4 Routes
===============================================================================
Flag  Network
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
i     64496:1:10.1.10.0/24
i     64496:1:10.1.11.0/24
i     64496:33:10.33.1.0/24
i     64496:33:10.33.2.0/24
-------------------------------------------------------------------------------
Routes : 4
===============================================================================
```

As already indicated, the 10.33.1.0/24 and 10.33.2.0/24 prefixes belong to VPRN 33. This service exists on PE-1 and PE-2 only, and the corresponding customer traffic must be kept within AS 64496, so there is no need to advertise these prefixes to the peer AS. The "exp-SVC-1" policy is defined at ASBR PE-4 to achieve this, as follows:

```
# on ASBR PE-4:
configure {
    policy-options {
        prefix-list "pfx-SVC-1" {
            prefix 10.1.10.0/24 type longer {
            }
            prefix 10.1.11.0/24 type longer {
            }
            prefix 2001:db8:1::/96 type longer {
            }
        }
        policy-statement "exp-SVC-1" {
            entry 10 {
                from {
                    prefix-list "pfx-SVC-1"
                }
                action {
                    action-type accept
                }
            }
            default-action {
                action-type reject
            }
        }
```

The "exp-SVC-1" policy is applied to ASBR PE-4 as an export policy, but also import policies can be used to control which prefixes are exchanged. This additionally requires the **vpn-apply-export true** (and the **vpn-apply-import true**) command, and the change required at ASBR PE-4 is as follows:

```
# on ASBR PE-4:
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            group "vpn-eBGP" {
                vpn-apply-export true
                export {
                    policy ["exp-SVC-1"]
                }
            }
            neighbor "192.168.48.2" {
                group "vpn-eBGP"
                peer-as 64510
                family {
                    vpn-ipv4 true
                    vpn-ipv6 true
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                ebgp-default-reject-policy {
                    import false
                    delete export
                }
            }
```

Therefore, the PE-4 ASBR will only advertise the VPN-IPv4 and VPN-IPv6 prefixes for VRPN 1, as follows:

```
[/]
A:admin@PE-4# show router bgp neighbor 192.168.48.2 advertised-routes vpn-ipv4 brief
===============================================================================
 BGP Router ID:192.0.2.4        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete
===============================================================================
BGP VPN-IPv4 Routes
===============================================================================
Flag  Network
-------------------------------------------------------------------------------
i     64496:1:10.1.10.0/24
i     64496:1:10.1.11.0/24
-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

```
[/]
A:admin@PE-4# show router bgp neighbor 192.168.48.2 advertised-routes vpn-ipv6 brief
===============================================================================
 BGP Router ID:192.0.2.4        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv6 Routes
===============================================================================
Flag  Network
-------------------------------------------------------------------------------
i     64496:1:2001:db8:1::1:0/120
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

Conclusion

Inter-AS VPRN model B offers service providers a way to interconnect IPv4 and IPv6 VPN sites across different ASs, avoiding the need for dedicated services in the ASBR, which would otherwise consume valuable resources in the ASBR. Model B is useful for scenarios where model C does not apply; for example, when there is no trust agreement with the peer AS, so that exchanging PE system addresses with that peer is not permitted or does not make sense.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Inter-AS VPRN Model B Using MPLS over UDP

This chapter describes Inter-AS VPRN Model B using MPLS over UDP.

Topics in this chapter include:

- Applicability
- Overview
- Configuration
- Conclusion

Applicability

This chapter was initially written based on SR OS Release 15.0.R8, but the MD-CLI in the current edition corresponds to SR OS Release 23.3.R3.

Overview

MPLS over UDP encapsulates MPLS packets into a UDP tunnel that can be transported by any IP network, and is defined in RFC 7510.

With MPLS over UDP, an outer IPv4/UDP or IPv6/UDP header encapsulates the inner MPLS label stack and message body; see Figure 306: IP over MPLS over UDP packet format. In the UDP header, the destination UDP port 6635 identifies the MPLS over UDP format to the egress PE and the source port number in the range from 49152 to 65535 is a source of entropy, because it is based on an ECMP hash calculation by the ingress PE. The entropy in the IP/UDP header ensures that ECMP uses all the available parallel paths between the tunnel endpoints.

Figure 306: IP over MPLS over UDP packet format

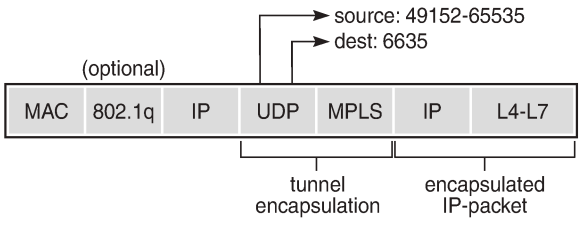

The MPLS over UDP implementation in SR OS has the following characteristics:

- The UDP checksum is set to 0 on transmission and ignored on reception.
- MPLS over UDP packets are processed only if they arrive with an IP destination address matching the system IP address.
- MPLS over UDP packets are originated with the Don't Fragment (DF) bit set in the outer IP header.
- Reassembly of received MPLS over UDP packets is not supported.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

BGP next-hops can be resolved to UDP-based tunnels for L3 VPN and EVPN services, which is useful, for example, in data center (DC) environments where IP is prevalent. This chapter describes the use and configuration of MPLS over UDP in an inter-AS VPRN model B environment. See the Inter-AS VPRN Model B chapter for more information.

Configuration

Figure 307: Inter-AS VPRN model B topology using MPLS over UDP in AS 64496 shows the topology used in this chapter, with two autonomous systems (ASs) and using inter-AS VPRN model B. VPRN "VPRN-1" carries customer traffic between PE-1 and PE-2 in AS 64496, and PE-5 and PE-6 in AS 64497. IS-IS is configured in AS 64496 and AS 64497, and MP-IBGP sessions are established between the PEs of each AS. An MP-EBGP session is established between autonomous system border routers (ASBRs) ASBR-4 and ASBR-8. RR-3 and RR-7 are the route reflectors (RRs) in AS 64496 and 64497, respectively, for the VPN-IPv4 and VPN-IPv6 address families. AS 64496 uses MPLS over UDP for its transport, whereas AS 64497 uses LDP.

Figure 307: Inter-AS VPRN model B topology using MPLS over UDP in AS 64496

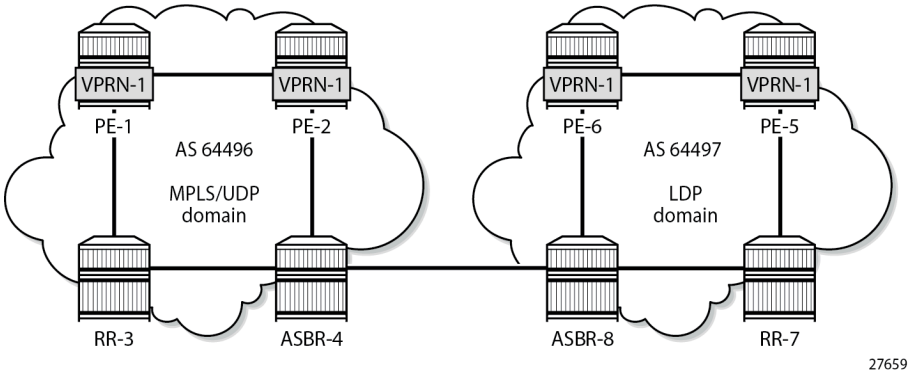

An inter-AS VPRN model B requires configuration of an MP-EBGP session with support for the VPN-IPv4 and VPN-IPv6 address families and enabling the inter-AS functionality on the ASBRs by the **inter-as-vpn** command. No configuration is required to support inter-AS VPRN model B on the PEs and the RRs. See the Inter-AS VPRN Model B chapter for more information.

The initial configuration includes:

- cards, MDAs, and ports
- router interfaces
- IS-IS as IGP on all interfaces (alternatively, OSPF can be used), with traffic engineering enabled.
- LDP configured in AS 64497, but not in AS 64496
- IBGP configured in both ASs
- EBGP configured between ASBR-4 and ASBR-8
- RR-3 and RR-7 configured as RR for VPN-IPv4 and VPN-IPv6 in AS 64496 and AS 64497

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

MPLS over UDP configuration

MPLS over UDP tunnels are UDP-based LSP tunnels that are created dynamically through a BGP import policy, where the action is **create-udp-tunnel true**. This import policy is configured and used on PE-1, PE-2, and ASBR-4. MPLS over UDP tunnels are created when BGP receives an update message where the incoming route has a next hop which matches the import policy. This import policy *create-UDP-tunnel* is defined as follows:

```
# on PE-1, PE-2, ASBR-4:
configure {
    policy-options {
        prefix-list "system-pfxs" {
            prefix 192.0.2.0/24 type range {
                start-length 32
                end-length 32
            }
        }
        policy-statement "create-UDP-tunnel" {
            entry 10 {
                from {
                    family [vpn-ipv4 vpn-ipv6]
                    next-hop {
                        prefix-list "system-pfxs"
                    }
                }
                action {
                    action-type accept
                    create-udp-tunnel true
                }
            }
```

This policy is applied to the IBGP sessions on which the VPN routes are exchanged; on PE-1 and PE-2, this is in the *IBGP-vpn* group, as follows:

```
# on PE-1, PE-2:
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            loop-detect discard-route
            split-horizon true
            next-hop-resolution {
                labeled-routes {
                    transport-tunnel {
                        family vpn {
                            resolution-filter {
                                bgp false
                                ldp false
                                udp true
                            }
                        }
                    }
                }
            }
            group "IBGP-vpn" {
                vpn-apply-import true
                peer-as 64496
                import {
                    policy ["create-UDP-tunnel"]
                }
            }
            neighbor "192.0.2.3" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                group "IBGP-vpn"
                family {
                    vpn-ipv4 true
                    vpn-ipv6 true
                }
            }
```

On RR-3, BGP is configured as follows:

```
# on RR-3:
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            loop-detect discard-route
            split-horizon true
            group "IBGP-vpn" {
                peer-as 64496
                family {
                    vpn-ipv4 true
                    vpn-ipv6 true
                }
                cluster {
                    cluster-id 192.0.2.3
                }
            }
            neighbor "192.0.2.1" {
                group "IBGP-vpn"
            }
            neighbor "192.0.2.2" {
                group "IBGP-vpn"
            }
            neighbor "192.0.2.4" {
                group "IBGP-vpn"
            }
```

For ASBR-4 to advertise VPRN routes to the peer AS, ASBR-4 must know the VPRN routes used within the AS, so it peers with RR-3 through IBGP, using the *IBGP-vpn* group. Although the **inter-as-vpn true** command enables inter-AS VPN model B, the VPN-IPv4 and VPN-IPv6 address family must also resolve to MPLS over UDP tunnels, so the BGP configuration on ASBR-4 is as follows:

```
# on ASBR-4:
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            loop-detect discard-route
            inter-as-vpn true                       # inter-AS VPRN model B
            split-horizon true
            next-hop-resolution {
                labeled-routes {
                    transport-tunnel {
                        family vpn {
                            # resolution filter     # default
                            resolution-filter {
                                bgp false           # default: bgp true
                                ldp false           # default: ldp true
                                udp true            # default: udp false
                            }
                        }
                    }
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
            group "EBGP-vpn" {
                peer-as 64497
            }
            group "IBGP-vpn" {
                vpn-apply-import true
                peer-as 64496
                import {
                    policy ["create-UDP-tunnel"]
                }
            }
            neighbor "192.0.2.3" {
                group "IBGP-vpn"
                family {
                    vpn-ipv4 true
                    vpn-ipv6 true
                }
            }
            neighbor "192.168.48.2" {
                group "EBGP-vpn"
                family {
                    vpn-ipv4 true
                    vpn-ipv6 true
                }
                ebgp-default-reject-policy {
                    import false
                    export false
                }
            }
```

Service configuration

VPRN "VPRN-1" is configured on PE-1, PE-2, PE-5, and PE-6. Although the VPRN service names and IDs used in both ASs do not have to match, when inter-AS VPRN model B applies, the route targets (RTs) used in both ASs must be coordinated. The RT exported by the PE-1 VPRN must be imported by PE-2, PE-5, and PE-6, and vice versa. In this example, no specific VRF import and VRF export policies are used; the **vrf-target** command is used instead.

To carry the customer data across AS 64496, the MPLS over UDP tunnels must bind to a VPRN service with the **auto-bind-tunnel** command. The resolution is set to **filter**, indicating that SR OS must select a tunnel using the information defined in the **resolution-filter** context. The **udp** keyword in the **resolution-filter** context indicates that MPLS over UDP tunnels can be used. If **resolution** is set to **any**, the resolution filter is ignored, and a tunnel from the tunnel table manager (TTM) is selected, based on availability and preference.

In AS 64496, the service "VPRN-1" on PE-1 is defined as follows. The configuration of service "VPRN-1" on PE-2 is similar.

```
# on PE-1:
configure {
    service {
        vprn "VPRN-1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:1"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    vrf-target {
                        community "target:64496:1"
                    }
                    auto-bind-tunnel {
                        resolution filter
                        resolution-filter {
                            bgp false
                            udp true
                        }
                    }
                }
            }
            interface "int-S1-1" {
                ipv4 {
                    primary {
                        address 10.1.10.1
                        prefix-length 24
                    }
                }
                sap 1/1/c4/1:1 {
                }
                ipv6 {
                    address 2001:db8:1::1:1 {
                        prefix-length 120
                    }
                }
            }
```

The transport technology in AS 64497 is LDP, so service "VPRN-1" in PE-5 and PE-6 auto-binds to LDP LSPs in the tunnel table, to resolve VPN-IPv4 routes for which the VRF target matches the VRF target community value configured in PE-1 and PE-2, as follows. The configuration of "VPRN-1" on PE-6 is similar.

```
# on PE-5:
configure {
    service {
        vprn "VPRN-1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64497:1"
                    vrf-target {
                        community "target:64496:1"
                    }
                    auto-bind-tunnel {
                        resolution filter
                        resolution-filter {
                            bgp false
                            ldp true
                        }
                    }
                }
            }
            interface "int-S1-1" {
                ipv4 {
                    primary {
                        address 10.1.50.1
                        prefix-length 24
                    }
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                sap 1/1/c4/1:1 {
                }
                ipv6 {
                    address 2001:db8:1::5:1 {
                        prefix-length 120
                    }
                }
            }
```

Verification

With the configurations from previous subsections applied, so that PE-1, PE-2, PE-5, and PE-6 have a service instance of "VPRN-1" running, PE-1 receives three VPN-IPv4 and three VPN-IPv6 prefixes, as follows:

```
[/]
A:admin@PE-1# show router bgp summary all

===============================================================================
BGP Summary
===============================================================================
Legend : D - Dynamic Neighbor
===============================================================================
Neighbor
Description
ServiceId          AS PktRcvd InQ  Up/Down   State|Rcv/Act/Sent (Addr Family)
                      PktSent OutQ
-------------------------------------------------------------------------------
192.0.2.3
Def. Inst       64496      18    0 00h04m56s 3/3/1 (VpnIPv4)
                           18    0           3/3/1 (VpnIPv6)

-------------------------------------------------------------------------------
```

PE-1 creates two UDP tunnels, and its tunnel table is as follows. The UDP tunnel to 192.0.2.2 is used for intra-AS customer traffic; the UDP tunnel to 192.0.2.4 is used for inter-AS customer traffic.

```
[/]
A:admin@PE-1# show router tunnel-table

===============================================================================
IPv4 Tunnel Table (Router: Base)
===============================================================================
Destination           Owner     Encap TunnelId  Pref   Nexthop        Metric
   Color
-------------------------------------------------------------------------------
192.0.2.2/32          udp       MPLS  786434    254    192.168.12.2   10
192.0.2.4/32          udp       MPLS  786435    254    192.168.12.2   20
-------------------------------------------------------------------------------
Flags: B = BGP or MPLS backup hop available
       L = Loop-Free Alternate (LFA) hop available
       E = Inactive best-external BGP route
       k = RIB-API or Forwarding Policy backup hop
===============================================================================
```

These tunnels are used as BGP next-hops for the VPN-IPv4 (and VPN-IPv6) routes, as follows:

```
[/]
A:admin@PE-1# show router bgp next-hop vpn-ipv4
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
 BGP Router ID:192.0.2.1        AS:64496       Local AS:64496
===============================================================================

===============================================================================
BGP VPN Next Hop
===============================================================================
VPN Next Hop                                                  Owner
   Autobind                                          FibProg  Reason
   Labels (User-labels)                              FlexAlgo Metric
   Admin-tag-policy (strict-tunnel-tagging)                   Last Mod.
-------------------------------------------------------------------------------
192.0.2.2                                                     UDP
   udp                                               Y
   -- (2)                                            --       10
   -- (N)                                                     00h03m44s
192.0.2.4                                                     UDP
   udp                                               Y
   -- (2)                                            --       20
   -- (N)                                                     00h03m14s
-------------------------------------------------------------------------------
Next Hops : 2
===============================================================================
```

The IPv4 and IPv6 routing tables for VPRN 1 are as follows:

```
[/]
A:admin@PE-1# show router 1 route-table ipv4

===============================================================================
Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
10.1.10.0/24                                  Local   Local     00h04m57s  0
       int-S1-1                                                     0
10.1.20.0/24                                  Remote  BGP VPN   00h04m08s  170
       192.0.2.2 (tunneled:UDP)                                     10
10.1.50.0/24                                  Remote  BGP VPN   00h03m38s  170
       192.0.2.4 (tunneled:UDP)                                     20
10.1.60.0/24                                  Remote  BGP VPN   00h03m38s  170
       192.0.2.4 (tunneled:UDP)                                     20
-------------------------------------------------------------------------------
No. of Routes: 4
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

```
[/]
A:admin@PE-1# show router 1 route-table ipv6

===============================================================================
IPv6 Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
2001:db8:1::1:0/120                           Local   Local     00h04m56s  0
       int-S1-1                                                     0
2001:db8:1::2:0/120                           Remote  BGP VPN   00h04m08s  170
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
       192.0.2.2 (tunneled:UDP)                                          10
2001:db8:1::5:0/120                          Remote  BGP VPN   00h03m38s  170
       192.0.2.4 (tunneled:UDP)                                          20
2001:db8:1::6:0/120                          Remote  BGP VPN   00h03m38s  170
       192.0.2.4 (tunneled:UDP)                                          20
-------------------------------------------------------------------------------
No. of Routes: 4
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

In this case, all VPRN remote prefixes are reachable through an MPLS over UDP tunnel (tunneled:UDP). The VPN label values for these prefixes are assigned and advertised by ASBR-4 and get to PE-1 via RR-3 in an MP-BGP update message, and can be displayed as follows:

```
[/]
A:admin@PE-1# show router bgp neighbor 192.0.2.3 received-routes vpn-ipv4
===============================================================================
 BGP Router ID:192.0.2.1        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv4 Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop (Router)                                 Path-Id     IGP Cost
      As-Path                                                      Label
-------------------------------------------------------------------------------
u*>i  64496:1:10.1.20.0/24                             100         None
      192.0.2.2                                        None        10
      No As-Path                                                   524286
u*>i  64497:1:10.1.50.0/24                             100         None
      192.0.2.4                                        None        20
      64497                                                        524287
u*>i  64497:1:10.1.60.0/24                             100         None
      192.0.2.4                                        None        20
      64497                                                        524286
-------------------------------------------------------------------------------
Routes : 3
===============================================================================
```

The prefixes 10.1.50.0/24 and 10.1.60.0/24 have ASBR-4 as next hop: prefix 10.1.50.0/24 gets VPRN label 524287 and prefix 10.1.60.0/24 gets VPRN label 524286.

```
[/]
A:admin@PE-1# show router bgp neighbor 192.0.2.3 received-routes vpn-ipv6
===============================================================================
 BGP Router ID:192.0.2.1        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv6 Routes
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop (Router)                                 Path-Id     IGP Cost
      As-Path                                                      Label
-------------------------------------------------------------------------------
u*>i  64496:1:2001:db8:1::2:0/120                      100         None
      ::ffff:192.0.2.2                                 None        10
      No As-Path                                                   524286
u*>i  64497:1:2001:db8:1::5:0/120                      100         None
      ::ffff:192.0.2.4                                 None        20
      64497                                                        524284
u*>i  64497:1:2001:db8:1::6:0/120                      100         None
      ::ffff:192.0.2.4                                 None        20
      64497                                                        524285
-------------------------------------------------------------------------------
Routes : 3
===============================================================================
```

The prefixes 2001:db8:1::5:0/120 and 2001:db8:1::6:0/120 have ASBR-4 as next hop: prefix 2001:db8:1::5:0/120 gets VPRN label 524284 and prefix 2001:db8:1::6:0/120 gets VPRN label 524285.

The forwarding plane is programmed accordingly, as follows:

```
[/]
A:admin@PE-1# show router 1 fib 1 ipv4

===============================================================================
FIB Display
===============================================================================
Prefix [Flags]                                           Protocol
  NextHop
-------------------------------------------------------------------------------
10.1.10.0/24                                             LOCAL
  10.1.10.0 (int-S1-1)
10.1.20.0/24                                             BGP_VPN
  192.0.2.2 (VPRN Label:524286 Transport:UDP)
10.1.50.0/24                                             BGP_VPN
  192.0.2.4 (VPRN Label:524287 Transport:UDP)
10.1.60.0/24                                             BGP_VPN
  192.0.2.4 (VPRN Label:524286 Transport:UDP)
-------------------------------------------------------------------------------
Total Entries : 4
-------------------------------------------------------------------------------
===============================================================================
```

```
[/]
A:admin@PE-1# show router 1 fib 1 ipv6

===============================================================================
FIB Display
===============================================================================
Prefix [Flags]                                           Protocol
  NextHop
-------------------------------------------------------------------------------
2001:db8:1::1:0/120                                      LOCAL
  2001:db8:1::1:0 (int-S1-1)
2001:db8:1::2:0/120                                      BGP_VPN
  192.0.2.2 (VPRN Label:524286 Transport:UDP)
2001:db8:1::5:0/120                                      BGP_VPN
  192.0.2.4 (VPRN Label:524284 Transport:UDP)
2001:db8:1::6:0/120                                      BGP_VPN
  192.0.2.4 (VPRN Label:524285 Transport:UDP)
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Total Entries : 4
------------------------------------------------------------------------------
==============================================================================
```

Traffic over the VPRN is generated using a **ping** command on PE-1 to the remote address, as follows:

```
[/]
A:admin@PE-1# ping 10.1.50.1 router-instance "VPRN-1"
PING 10.1.50.1 56 data bytes
64 bytes from 10.1.50.1: icmp_seq=1 ttl=64 time=6.21ms.
64 bytes from 10.1.50.1: icmp_seq=2 ttl=64 time=5.93ms.
64 bytes from 10.1.50.1: icmp_seq=3 ttl=64 time=5.88ms.
64 bytes from 10.1.50.1: icmp_seq=4 ttl=64 time=5.98ms.
64 bytes from 10.1.50.1: icmp_seq=5 ttl=64 time=4.92ms.

---- 10.1.50.1 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 4.92ms, avg = 5.78ms, max = 6.21ms, stddev = 0.445ms
```

In contrast with the traditional inter-AS VPRN model B, where customer traffic is pushed with a VPN label followed by a transport label, now customer traffic to destination 10.1.50.1 is pushed with VPN label 524287, followed by an IP/UDP header with IP SA 192.0.2.1 and IP DA 192.0.2.4 and with UDP source and destination port.

The interconnection between the ASBRs carries the VPN data with a single MPLS label, so ASBR-4 removes the IP/UDP header and swaps the VPN label 524287 with the VPN label 524283 received from ASBR-8. The inter-AS label mapping on ASBR-4 is as follows:

```
[/]
A:admin@ASBR-4# show router bgp inter-as-label

===============================================================================
BGP Inter-AS labels
Flags: B - entry has backup, P - entry is promoted
===============================================================================
NextHop                       Received       Advertised     Label
                              Label          Label          Origin
-------------------------------------------------------------------------------
192.0.2.1                     524286         524283         Internal
192.0.2.1                     524286         524281         Internal
192.0.2.2                     524286         524282         Internal
192.0.2.2                     524286         524280         Internal
192.168.48.2                  524280         524285         External
192.168.48.2                  524281         524284         External
192.168.48.2                  524282         524286         External
192.168.48.2                  524283         524287         External
-------------------------------------------------------------------------------
Total Labels allocated:   8
===============================================================================
```

The forward data flow from PE-1 to PE-5 for VPRN 1 uses the labels for which the label origin is external. The VPN labels used for the backward data flow use the labels for which the label origin is internal.

For brevity, the commands to display and check VPN prefixes and labels in AS 64497 are omitted.

Customer traffic destined to the VPN routes received by ASBR-4 can be sent to the correct destination because ASBR-4 has the relevant BGP next-hops resolved to UDP tunnels, as follows:

```
[/]
A:admin@ASBR-4# show router bgp next-hop vpn-ipv4
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
 BGP Router ID:192.0.2.4        AS:64496       Local AS:64496
===============================================================================

===============================================================================
BGP VPN Next Hop
===============================================================================
VPN Next Hop                                                  Owner
   Autobind                                          FibProg  Reason
   Labels (User-labels)                              FlexAlgo Metric
   Admin-tag-policy (strict-tunnel-tagging)                   Last Mod.
-------------------------------------------------------------------------------
192.0.2.1                                                     UDP
   udp                                               Y
   -- (2)                                            --       20
   -- (N)                                                     00h06m30s
192.0.2.2                                                     UDP
   udp                                               Y
   -- (2)                                            --       10
   -- (N)                                                     00h06m30s
192.168.48.2                                                  LOCAL
   udp                                               Y
   -- (2)                                            --       0
   -- (N)                                                     00h06m09s
-------------------------------------------------------------------------------
Next Hops : 3
===============================================================================
```

The forwarding plane is programmed accordingly, as follows:

```
[/]
A:admin@ASBR-4# show router fp-tunnel-table 1

===============================================================================
IPv4 Tunnel Table Display

Legend:
label stack is ordered from bottom-most to top-most
B - FRR Backup
===============================================================================
Destination                                Protocol         Tunnel-ID
  Lbl/SID
    NextHop                                               Intf/Tunnel
  Lbl/SID (backup)
    NextHop   (backup)
-------------------------------------------------------------------------------
192.0.2.1/32                               UDP               -
  -
    192.168.24.1                                         1/1/c2/1:1000
192.0.2.2/32                               UDP               -
  -
    192.168.24.1                                         1/1/c2/1:1000
-------------------------------------------------------------------------------
Total Entries : 2
-------------------------------------------------------------------------------
===============================================================================
```

Changing the BGP next-hop resolution for auto-binding the BGP next-hop on PE-1, and for VPN next-hop resolution on ASBR-4 from **resolution filter** to **resolution any**, does not lead to the tunnel being changed, but to a change of the allowed tunnel types for auto-bind. On PE-1 and PE-2, SR OS selects UDP as the

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

only viable tunnel type, because no tunnels of type LDP, RSVP, SR-ISIS, SR-OSPF, GRE and so on are available:

```
[/]
A:admin@PE-1# show router bgp next-hop vpn-ipv4
===============================================================================
 BGP Router ID:192.0.2.1        AS:64496       Local AS:64496
===============================================================================

===============================================================================
BGP VPN Next Hop
===============================================================================
VPN Next Hop                                                    Owner
   Autobind                                          FibProg    Reason
   Labels (User-labels)                              FlexAlgo   Metric
   Admin-tag-policy (strict-tunnel-tagging)                     Last Mod.
-------------------------------------------------------------------------------
192.0.2.2                                                       UDP
   ldp rsvp sr-isis sr-ospf gre bgp sr-te udp sr-p Y
   olicy rib-api mpls-fwd-policy sr-ospf3
   -- (2)                                          --         10
   -- (N)                                                     00h00m05s
192.0.2.4                                                       UDP
   ldp rsvp sr-isis sr-ospf gre bgp sr-te udp sr-p Y
   olicy rib-api mpls-fwd-policy sr-ospf3
   -- (2)                                          --         20
   -- (N)                                                     00h00m05s
-------------------------------------------------------------------------------
Next Hops : 2
===============================================================================
```

```
[/]
A:admin@ASBR-4# show router bgp next-hop vpn-ipv4
===============================================================================
 BGP Router ID:192.0.2.4        AS:64496       Local AS:64496
===============================================================================

===============================================================================
BGP VPN Next Hop
===============================================================================
VPN Next Hop                                                    Owner
   Autobind                                          FibProg    Reason
   Labels (User-labels)                              FlexAlgo   Metric
   Admin-tag-policy (strict-tunnel-tagging)                     Last Mod.
-------------------------------------------------------------------------------
192.0.2.1                                                       UDP
   ldp rsvp sr-isis sr-ospf bgp sr-te udp sr-polic Y
   y rib-api mpls-fwd-policy sr-ospf3
   -- (2)                                          --         20
   -- (N)                                                     00h00m25s
192.0.2.2                                                       UDP
   ldp rsvp sr-isis sr-ospf bgp sr-te udp sr-polic Y
   y rib-api mpls-fwd-policy sr-ospf3
   -- (2)                                          --         10
   -- (N)                                                     00h00m25s
192.168.48.2                                                    LOCAL
   ldp rsvp sr-isis sr-ospf bgp sr-te udp sr-polic Y
   y rib-api mpls-fwd-policy sr-ospf3
   -- (2)                                          --         0
   -- (N)                                                     00h00m25s
-------------------------------------------------------------------------------
Next Hops : 3
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
```

Conclusion

VPRN services support the resolution of VPN-IPv4 and VPN-IPv6 BGP next-hops to MPLS over UDP tunnels. MPLS over UDP tunnels are useful in IP-based fabric networks, such as DCs. SR OS supports inter-AS model B for any type of MPLS-based tunnels, including MPLS over UDP.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Inter-AS VPRN Model C

This chapter provides information about virtual private routed network (VPRN) inter-autonomous system (AS) virtual private network (VPN) model C.

Topics in this chapter include:

Applicability

This chapter was initially written for SR OS Release 7.0. The MD-CLI in the current edition corresponds to SR OS Release 22.2.R1. There are no prerequisites for this configuration.

Overview

Introduction

Section 10 of RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*, describes three potential methods for service providers to interconnect their IP-VPN (Internet Protocol — Virtual Private Network) backbones to provide an end-to-end MPLS-VPN where one or more sites of the VPN are connected to different service provider autonomous systems (ASs). The purpose of this chapter is to describe the configuration and troubleshooting for inter-AS VPN model C.

In this architecture, VPN prefixes are neither held, nor re-advertised by the Autonomous System Border Router – Provider Edge (ASBR-PE) routers, which makes Model C more scalable than Model B (where the only prefixes exchanged between ASs are VPN-IPv4). In Model C, the only MPLS data plane resources consumed in the ASBRs are for infrastructure addresses of PEs and RRs rather than VPN prefixes.

In this example, an export policy is configured to ensure that the nodes advertise their system IP addresses (IPv4 /32 addresses) in labeled BGP to all their BGP peers within the AS. Therefore, the ASBR-PE maintains labeled IPv4 /32 BGP routes to other PE routers within its own AS. These BGP routes are inactive, because for each destination within the AS, an IGP route exists which is preferred to BGP routes. The ASBR redistributes these inactive /32 IPv4 prefixes in external Border Gateway Protocol (EBGP) to the ASBR-PE in other service providers ASs, because **advertise-inactive** is configured in EBGP. No export policy is required in EBGP.

At the same time, the ASBR programs a label switch for the received and advertised BGP labels. The receiving ASBR advertises the received IP system prefixes to its IBGP peers (in this case, a Route Reflector (RR)) within their AS, and eventually, all PEs in the AS learn the system IP prefixes of the peer AS. However, there is no need to learn the system IP address of the ASBRs in peer ASs, because they do not exchange customer VPN prefixes.

After the system IP addresses have been learned in the peer AS, it is possible for PE routers in different ASs to establish multi-hop Multi Protocol – external Border Gateway Protocol (MP-EBGP) sessions for address family VPN-IPv4 to each other to exchange customer VPN prefixes over those connections. The

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

multihop sessions can be established between the RR in the ASs, but these RRs should not modify the next-hop attribute of the BGP update across the EBGP session.

A three-level label stack is imposed on the ingress PE. The bottom-level label is assigned by the egress PE (advertised in multi-hop MP-EBGP without next-hop override) and is commonly referred to as the VPN-label. The middle label is assigned by the local ASBR-PE and corresponds to the /32 route of the egress PE (in a different AS) using BGP-LBL (RFC 3107, *Carrying Label Information in BGP-4*). The top level label is the label assigned by the local ASBR-PE /32 loopback address, which is assigned by the IGP next-hop of the ingress PE. This label is referred to as the LDP-LBL.

Figure 308: Inter-AS VPN Model C illustrates this mechanism. The VPN-LBL is assigned by PE-5, the BGP-LBL is assigned by PE-4 and the LDP-LBL is also assigned by PE-4. The BGP-LBL is swapped in both ASBRs. The label stack contains three labels in each AS: VPN-LBL, BGP-LBL, and LDP-LBL) and two labels on the EBGP link between the ASs: VPN-LBL and BGP-LBL.

Note:
This configuration that uses **advertise-inactive** is preferred to a configuration where the BGP routes are not exchanged within their AS and the ASBRs use an export policy with a prefix list for all local system prefixes to be advertised to the peer ASs. The routes for those prefixes are taken from the RTM, where these routes are not known via BGP, but via IS-IS. In that case, IS-IS routes are effectively redistributed into labeled BGP (which most operators do not want) and as a result, the ASBR is not programming a label switch for the BGP label. Furthermore, the label stack is asymmetrical: three labels in the originating AS (VPN-LBL, BGP-LBL, and LDP-LBL) and only two labels in the target AS (VPN-LBL, LDP-LBL), because the local routes are not known via labeled BGP in this scenario. This scenario is not described in this chapter; only the preferred scenario with local labeled BGP routes in each AS is described.

Figure 308: Inter-AS VPN Model C

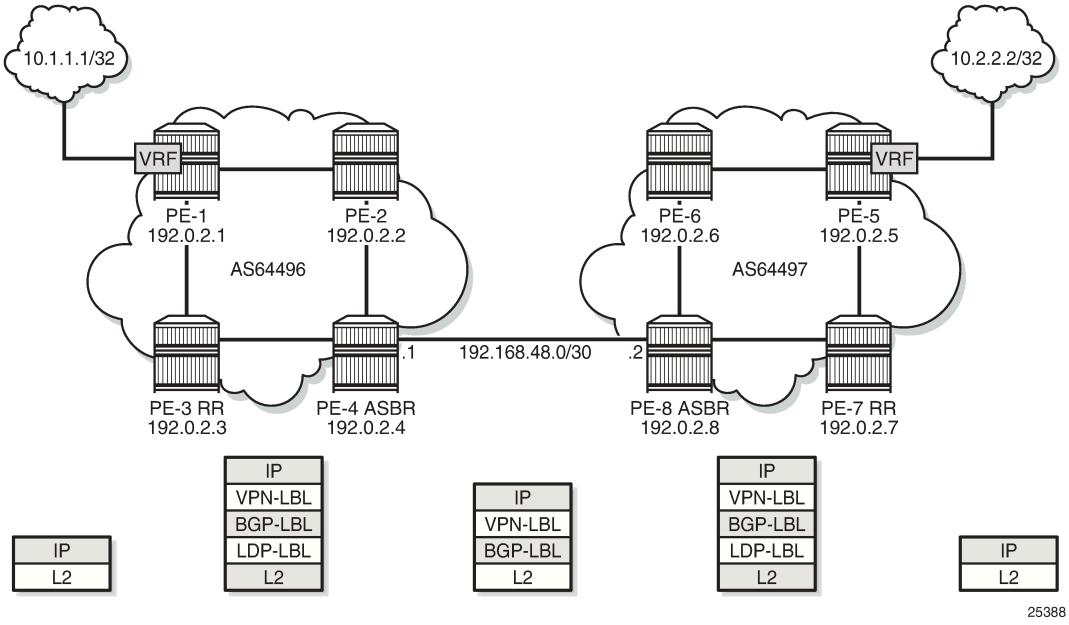

The VPN connectivity is established using Labeled VPN route exchange using MP-EBGP without next-hop override. The PE connectivity is established as follows.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

EBGP PE /32 loopback leaking routing exchange using EBGP LBL (RFC 3107) at the ASBR-PE. The /32 PE routes learned from the other AS through the ASBR-PE are further distributed into the local AS using IBGP and optionally through Route Reflectors (RRs). This model uses a three label stack and is referred to as Model C. Resilience for ASBR-PE failures depends on BGP.

Figure 309: Protocol overview

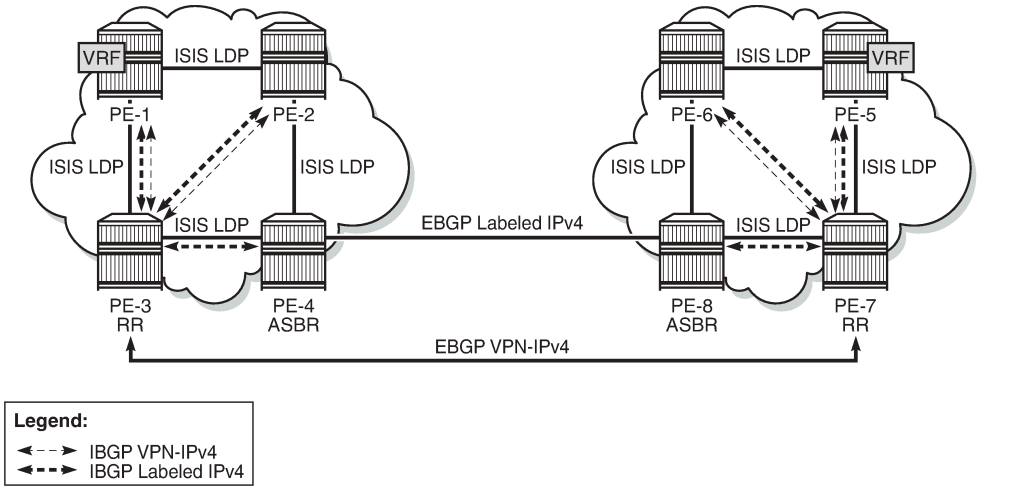

Figure 309: Protocol overview shows the protocols used when implementing Inter-AS Model C. Inside each AS, there is an IS-IS adjacency and a link LDP session between each pair of adjacent nodes. As an alternative, OSPF can be used as IGP. There is also an IBGP session between each PE and the RR. The address family is both VPN-IPv4 for the exchange of customer VPN prefixes and Labeled IPv4 for the exchange of labeled IPv4 prefixes. Between the RR and the ASBR, only Labeled IPv4 is required, because the ASBR does not exchange any customer VPN prefixes. When no RR is used, a full mesh of IBGP sessions can be established in each AS.

Between the ASBRs, there is an EBGP session for the exchange of labeled IPv4 prefixes. The ASBRs override the next-hop for those prefixes. Between the RRs in the different ASs, there is a multihop EBGP session for the exchange of VPN-IPv4 customer prefixes. The RRs do not override the next-hop for those prefixes.

The main advantage of this model is that no VPN routes need to be held on the ASBR-PEs and therefore, it scales the best among all the three Inter-AS IP-VPN models. However, leaking /32 PE addresses between service providers raises some security concerns. Therefore, we see Model C typically deployed within a service provider network.

Figure 308: Inter-AS VPN Model C shows the example topology which consists of four SR OS nodes in AS 64496 and four SR OS nodes in AS 64497. There is an AS interconnection between ASBR PE-4 to ASBR PE-8. PE-3 and PE-7 act as RRs for their AS. An IP-VPN is configured in each AS. The initial configuration includes the following:

- IS-IS or OSPF on all interfaces within each of the ASs.
- LDP on all interfaces within each of the ASs.
- MP-IBGP sessions between the PE routers and the RRs in each of the ASs, as shown in the following section.
- IP-VPN on PE-1 and on PE-5 with identical route targets.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- A loopback interface in the VRF on PE-1 and PE-5.

Configuration

The first step is to configure an MP-IBGP session between the PEs in both ASs. An export policy is configured to export the system prefixes from the PEs in labeled BGP.

PE-3 and PE-7 act as RR in the ASs. In AS 64496, PE-1 and PE-2 are peered with RR PE-3 for the labeled IPv4 and VPN-IPv4 address families; ASBR PE-4 is peered with RR PE-3 for the labeled IPv4 address family only. In AS 64497, PE-5 and PE-6 are peered with RR PE-7 for the labeled IPv4 and VPN-IPv4 address families; ASBR PE-8 is peered with RR PE-7 for the labeled IPv4 address family only.

Address family **label-ipv4** is required to advertise labeled IPv4 routes toward each neighbor PE. Address family **vpn-ipv4** is required to advertise IPv4 customer VPN routes within the AS.

The initial BGP configuration for RR PE-3 is as follows:

```
# on RR PE-3:
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            split-horizon true
            group "IBGP" {
                peer-as 64496
                cluster {
                    cluster-id 192.0.2.3
                }
                export {
                    policy ["export-bgp"]
                }
            }
            neighbor "192.0.2.1" {
                advertise-inactive true
                group "IBGP"
                family {
                    vpn-ipv4 true
                    label-ipv4 true
                }
            }
            neighbor "192.0.2.2" {
                advertise-inactive true
                group "IBGP"
                family {
                    vpn-ipv4 true
                    label-ipv4 true
                }
            }
            neighbor "192.0.2.4" {
                advertise-inactive true
                group "IBGP"
                family {
                    label-ipv4 true
                }
            }
```

The export policy exports the system IP address and is defined as follows:

```
# on PE-1, PE-2, PE-3, PE-5, PE-6, PE-7:
configure {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    policy-options {
        prefix-list "PE-sys" {
            prefix 192.0.2.0/28 type longer {
            }
        }
        policy-statement "export-bgp" {
            entry 10 {
                from {
                    prefix-list ["PE-sys"]
                    protocol {
                        name [direct]
                    }
                }
                action {
                    action-type accept
                }
            }
        }
```

On the ASBRs in both ASs, EBGP and IBGP need to be configured. The EBGP session is configured with **advertise-inactive** and is used to redistribute labeled IPv4 routes for the /32 system IP addresses between the ASs, even if those routes are not the most preferred routes within the system for a specific destination.

The configuration for ASBR PE-4 is as follows. The address family **label-ipv4** is required to enable the advertising of labeled IPv4 routes. This address family is also required on the RR neighbor to propagate the labeled IPv4 routes toward the other PEs in the AS.

```
# on ASBR PE-4:
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            split-horizon true
            group "EBGP" {
            }
            group "IBGP" {
                peer-as 64496
            }
            neighbor "192.0.2.3" {
                group "IBGP"
                family {
                    label-ipv4 true
                }
            }
            neighbor "192.168.48.2" {
                advertise-inactive true
                group "EBGP"
                peer-as 64497
                family {
                    label-ipv4 true
                }
                ebgp-default-reject-policy {
                    import false
                    export false
                }
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

On the remaining PE nodes in AS 64496, PE-1 and PE-2, the address families **label-ipv4** and **vpn-ipv4** must be enabled, as follows:

```
# on PE-1, PE-2:
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            split-horizon true
            group "IBGP" {
                peer-as 64496
                export {
                    policy ["export-bgp"]
                }
            }
            neighbor "192.0.2.3" {
                group "IBGP"
                family {
                    vpn-ipv4 true
                    label-ipv4 true
                }
            }
        }
```

The configuration for the nodes in AS 64497 is similar. The IP addresses can be derived from Figure 308: Inter-AS VPN Model C.

The following command on ASBR PE-4 verifies that the EBGP and IBGP sessions for the labeled IPv4 address family are up:

```
[/]
A:admin@PE-4# show router bgp summary all

===============================================================================
BGP Summary
===============================================================================
Legend : D - Dynamic Neighbor
===============================================================================
Neighbor
Description
ServiceId          AS PktRcvd InQ  Up/Down   State|Rcv/Act/Sent (Addr Family)
                      PktSent OutQ
-------------------------------------------------------------------------------
192.0.2.3
Def. Inst       64496      11    0 00h02m36s 3/0/3 (Lbl-IPv4)
                           9    0
192.168.48.2
Def. Inst       64497       7    0 00h01m40s 3/3/3 (Lbl-IPv4)
                           8    0

-------------------------------------------------------------------------------
```

On ASBR PE-4, three inactive labeled IPv4 routes have been received from the IBGP peers and three active labeled IPv4 routes have been received via EBGP, as follows:

```
[/]
A:admin@PE-4# show router bgp routes label-ipv4
===============================================================================
 BGP Router ID:192.0.2.4        AS:64496       Local AS:64496
===============================================================================
 Legend -
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP LABEL-IPV4 Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop (Router)                                 Path-Id     IGP Cost
      As-Path                                                      Label
-------------------------------------------------------------------------------
*i    192.0.2.1/32                                     100         None
      192.0.2.1                                        None        20
      No As-Path                                                   524283
*i    192.0.2.2/32                                     100         None
      192.0.2.2                                        None        10
      No As-Path                                                   524283
*i    192.0.2.3/32                                     100         None
      192.0.2.3                                        None        10
      No As-Path                                                   524283
u*>i  192.0.2.5/32                                     None        None
      192.168.48.2                                     None        0
      64497                                                        524283
u*>i  192.0.2.6/32                                     None        None
      192.168.48.2                                     None        0
      64497                                                        524282
u*>i  192.0.2.7/32                                     None        None
      192.168.48.2                                     None        0
      64497                                                        524281
-------------------------------------------------------------------------------
Routes : 6
===============================================================================
```

The following three routes have been received from EBGP peer PE-8: one for each system IP address in the remote AS, except for the ASBR itself:

```
[/]
A:admin@PE-4# show router bgp neighbor 192.168.48.2 routes received-routes family label-ipv4
===============================================================================
 BGP Router ID:192.0.2.4        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP LABEL-IPV4 Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop (Router)                                 Path-Id     IGP Cost
      As-Path                                                      Label
-------------------------------------------------------------------------------
u*>i  192.0.2.5/32                                     n/a         None
      192.168.48.2                                     None        0
      64497                                                        524283
u*>i  192.0.2.6/32                                     n/a         None
      192.168.48.2                                     None        0
      64497                                                        524282
u*>i  192.0.2.7/32                                     n/a         None
      192.168.48.2                                     None        0
      64497                                                        524281
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Routes : 3
===============================================================================
```

In this example, the IP prefix for PE-8 itself is not included. The prefix of the ASBR need not be advertised in labeled BGP to the remote AS, because ASBRs do not advertise VPN-IPv4 prefixes.

More detailed information about the advertised route from PE-5 can be seen with following command on PE-4:

```
[/]
A:admin@PE-4# show router bgp routes 192.0.2.5/32 label-ipv4 hunt
===============================================================================
 BGP Router ID:192.0.2.4        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP LABEL-IPV4 Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : 192.0.2.5/32
Nexthop        : 192.168.48.2
Path Id        : None
From           : 192.168.48.2
Res. Nexthop   : 192.168.48.2
Local Pref.    : None                   Interface Name : int-PE-4-PE-8
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : 0
Connector      : None
Community      : No Community Members
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.8
Fwd Class      : None                   Priority       : None
IPv4 Label     : 524283
Flags          : Used Valid Best IGP In-TTM In-RTM
Route Source   : External
AS-Path        : 64497
Route Tag      : 0
Neighbor-AS    : 64497
Orig Validation: NotFound
Source Class   : 0                      Dest Class     : 0
Add Paths Send : Default
RIB Priority   : Normal
Last Modified  : 00h03m58s

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
Network        : 192.0.2.5/32
Nexthop        : 192.0.2.4
Path Id        : None
To             : 192.0.2.3
Res. Nexthop   : n/a
Local Pref.    : 100                    Interface Name : NotAvailable
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : 0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Connector      : None
Community      : No Community Members
Cluster        : No Cluster Members
Originator Id  : None                  Peer Router Id : 192.0.2.3
IPv4 Label     : 524283                Label Type     : SWAP
Lbl Allocation : NEXT-HOP
Origin         : IGP
AS-Path        : 64497
Route Tag      : 0
Neighbor-AS    : 64497
Orig Validation: NotFound
Source Class   : 0                     Dest Class     : 0

-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

In the RIB In entries, the received label from PE-8 can be seen (524283). In the RIB Out entries, the locally assigned (Advertised) label for this prefix can be seen (524283). These labels need not match. The ASBR PE-4 swaps BGP labels, according to the following label mapping:

```
[/]
A:admin@PE-4# show router bgp inter-as-label

===============================================================================
BGP Inter-AS labels
Flags: B - entry has backup, P - entry is promoted
===============================================================================
NextHop                      Received       Advertised     Label
                             Label          Label          Origin
-------------------------------------------------------------------------------
192.0.2.1                    524283         524280         Internal
192.0.2.2                    524283         524279         Internal
192.0.2.3                    524283         524278         Internal
192.168.48.2                 524281         524281         External
192.168.48.2                 524282         524282         External
192.168.48.2                 524283         524283         External
-------------------------------------------------------------------------------
Total Labels allocated:   6
===============================================================================
```

The route from PE-1 toward PE-5 uses received label 524283 and advertised label 524283, as indicated on the sixth row in the table. The BGP label in the label stack sent by PE-1 contains BGP label 524283 toward ASBR PE-4, where it is swapped to BGP label 524283 toward ASBR PE-8.

ASBR PE-8 swaps BGP label 524283 to BGP label 524283 toward PE-5, as follows:

```
[/]
A:admin@PE-8# show router bgp inter-as-label

===============================================================================
BGP Inter-AS labels
Flags: B - entry has backup, P - entry is promoted
===============================================================================
NextHop                      Received       Advertised     Label
                             Label          Label          Origin
-------------------------------------------------------------------------------
192.0.2.5                    524283         524283         Internal
192.0.2.6                    524283         524282         Internal
192.0.2.7                    524283         524281         Internal
192.168.48.1                 524278         524278         External
192.168.48.1                 524279         524279         External
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
192.168.48.1                  524280         524280         External
-------------------------------------------------------------------------------
Total Labels allocated:   6
===============================================================================
```

On ASBR PE-4, the routes toward PE-5, PE-6, and PE-7 in the remote AS have been installed in the route table, as follows:

```
[/]
A:admin@PE-4# show router route-table

===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
192.0.2.1/32                                    Remote  ISIS      00h07m04s  18
       192.168.24.1                                                 20
192.0.2.2/32                                    Remote  ISIS      00h07m04s  18
       192.168.24.1                                                 10
192.0.2.3/32                                    Remote  ISIS      00h07m04s  18
       192.168.34.1                                                 10
192.0.2.4/32                                    Local   Local     00h07m05s  0
       system                                                       0
192.0.2.5/32                                    Remote  BGP_LABEL 00h05m08s  170
       192.168.48.2                                                 0
192.0.2.6/32                                    Remote  BGP_LABEL 00h05m08s  170
       192.168.48.2                                                 0
192.0.2.7/32                                    Remote  BGP_LABEL 00h05m08s  170
       192.168.48.2                                                 0
192.168.24.0/30                                 Local   Local     00h07m05s  0
       int-PE-4-PE-2                                                0
192.168.34.0/30                                 Local   Local     00h07m05s  0
       int-PE-4-PE-3                                                0
192.168.48.0/30                                 Local   Local     00h07m05s  0
       int-PE-4-PE-8                                                0
-------------------------------------------------------------------------------
No. of Routes: 10
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

The BGP labeled routes for the remote PE system prefixes are further advertised toward all the PEs in the AS (through the RR) and are installed in the routing table on all PEs.

At this point, all PEs in one AS have the /32 system IPs of the remote PEs in their routing table, for example for PE-1:

```
[/]
A:admin@PE-1# show router route-table

===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
192.0.2.1/32                                    Local   Local     00h07m25s  0
       system                                                       0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
192.0.2.2/32                                    Remote  ISIS      00h07m18s  18
       192.168.12.2                                                    10
192.0.2.3/32                                    Remote  ISIS      00h07m12s  18
       192.168.13.2                                                    10
192.0.2.4/32                                    Remote  ISIS      00h07m05s  18
       192.168.12.2                                                    20
192.0.2.5/32                                    Remote  BGP_LABEL 00h04m42s  170
       192.0.2.4 (tunneled)                                            20
192.0.2.6/32                                    Remote  BGP_LABEL 00h04m42s  170
       192.0.2.4 (tunneled)                                            20
192.0.2.7/32                                    Remote  BGP_LABEL 00h04m42s  170
       192.0.2.4 (tunneled)                                            20
192.168.12.0/30                                 Local   Local     00h07m25s  0
       int-PE-1-PE-2                                                   0
192.168.13.0/30                                 Local   Local     00h07m25s  0
       int-PE-1-PE-3                                                   0
-------------------------------------------------------------------------------
No. of Routes: 9
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

All PEs in one AS have also received labels for all /32 system IP addresses of the remote PEs. Therefore, an MP-EBGP session can be created between the RRs in the different ASs to exchange VPN-IPv4 routes.

The additional BGP configuration for RR PE-3 is as follows. The configuration for RR PE-7 is similar. The IP addresses can be derived from Figure 309: Protocol overview.

```
# on RR PE-3:
configure {
    router "Base" {
        bgp {
            group "peer-AS-RR" {
                peer-as 64497
                local-address 192.0.2.3
                family {
                    vpn-ipv4 true
                }
            }
            neighbor "192.0.2.7" {
                group "peer-AS-RR"
                multihop 10
                vpn-apply-export true
                vpn-apply-import true
                import {
                    policy ["import-ebgp-vpn"]
                }
                export {
                    policy ["export-ebgp-vpn"]
                }
            }
        }
```

Policies can be applied on the peering session using the **export** or **import** command followed by a policy name, together with the **vpn-apply-export** or **vpn-apply-import** command necessary to enforce base BGP instance policy on VPN-IPv4 prefixes.

On the RRs, the MP-EBGP session is up, as follows:

```
[/]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
A:admin@PE-3# show router bgp neighbor 192.0.2.7

===============================================================================
BGP Neighbor
===============================================================================
-------------------------------------------------------------------------------
Peer                : 192.0.2.7
Description         : (Not Specified)
Group               : peer-AS-RR
-------------------------------------------------------------------------------
Peer AS             : 64497           Peer Port           : 179
Peer Address        : 192.0.2.7
Local AS            : 64496           Local Port          : 51192
Local Address       : 192.0.2.3
Peer Type           : External        Dynamic Peer        : No
State               : Established     Last State          : Active
Last Event          : recvOpen
Last Error          : Unrecognized Error
Local Family        : VPN-IPv4
Remote Family       : VPN-IPv4
---snip---
```

The EBGP session between the two RRs is established.

The VPRNs on PE-1 in AS 64496 and PE-5 in AS 64497 are now interconnected. The route table for VPRN 1 shows that the remote PE can be reached via a BGP tunnel, as follows:

```
[/]
A:admin@PE-1# show router 1 route-table

===============================================================================
Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
10.1.1.1/32                                   Local   Local     00h09m51s  0
       loopback                                                      0
10.2.2.2/32                                   Remote  BGP VPN   00h00m29s  170
       192.0.2.5 (tunneled:BGP)                                      1000
-------------------------------------------------------------------------------
No. of Routes: 2
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

Packets originating in AS 64496 with a destination in AS 64497 have three labels in AS 64496 (and in AS 64497). Originate a VPRN ping on PE-1 toward the VPRN loopback IP address on PE-5:

```
[/]
A:admin@PE-1# ping 10.2.2.2 router-instance "VPRN1"
PING 10.2.2.2 56 data bytes
64 bytes from 10.2.2.2: icmp_seq=1 ttl=64 time=5.80ms.
64 bytes from 10.2.2.2: icmp_seq=2 ttl=64 time=5.89ms.
64 bytes from 10.2.2.2: icmp_seq=3 ttl=64 time=6.00ms.
64 bytes from 10.2.2.2: icmp_seq=4 ttl=64 time=6.54ms.
64 bytes from 10.2.2.2: icmp_seq=5 ttl=64 time=6.11ms.

---- 10.2.2.2 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
round-trip min = 5.80ms, avg = 6.07ms, max = 6.54ms, stddev = 0.257ms
```

The top label is the LDP label to reach the exit point of the AS (PE-4). This label has value 524284, as can be seen with following command on PE-1:

```
[/]
A:admin@PE-1# show router ldp bindings active prefixes prefix 192.0.2.4/32

===============================================================================
LDP Bindings (IPv4 LSR ID 192.0.2.1)
             (IPv6 LSR ID ::)
===============================================================================
Label Status:
        U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
        e - Label ELC
FEC Flags:
        LF - Lower FEC, UF - Upper FEC, M - Community Mismatch,
        BA - ASBR Backup FEC
        (S) - Static           (M) - Multi-homed Secondary Support
        (B) - BGP Next Hop     (BU) - Alternate Next-hop for Fast Re-Route
        (I) - SR-ISIS Next Hop (O) - SR-OSPF Next Hop
        (C) - FEC resolved with class-based-forwarding
===============================================================================
LDP IPv4 Prefix Bindings (Active)
===============================================================================
Prefix                                  Op
IngLbl                                  EgrLbl
EgrNextHop                              EgrIf/LspId
-------------------------------------------------------------------------------
192.0.2.4/32                            Push
  --                                    524284
192.168.12.2                            1/1/1

192.0.2.4/32                            Swap
524284                                  524284
192.168.12.2                            1/1/1

-------------------------------------------------------------------------------
No. of IPv4 Prefix Active Bindings: 2
===============================================================================
```

This LDP label is popped by ASBR PE-4. No LDP label is used between the ASBRs. ASBR PE-8 pushes another LDP label.

To reach a PE in the remote AS, a BGP transport label is required, which is the middle label in the stack. The tunnel table on PE-1 shows a BGP tunnel toward PE-5, as follows:

```
[/]
A:admin@PE-1# show router tunnel-table

===============================================================================
IPv4 Tunnel Table (Router: Base)
===============================================================================
Destination           Owner     Encap TunnelId  Pref   Nexthop        Metric
   Color
-------------------------------------------------------------------------------
192.0.2.2/32          ldp       MPLS  65537     9      192.168.12.2   10
192.0.2.3/32          ldp       MPLS  65538     9      192.168.13.2   10
192.0.2.4/32          ldp       MPLS  65539     9      192.168.12.2   20
192.0.2.5/32          bgp       MPLS  262145    12     192.0.2.4      1000
192.0.2.6/32          bgp       MPLS  262146    12     192.0.2.4      1000
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
192.0.2.7/32          bgp       MPLS  262147    12     192.0.2.4      1000
-------------------------------------------------------------------------------
Flags: B = BGP or MPLS backup hop available
       L = Loop-Free Alternate (LFA) hop available
       E = Inactive best-external BGP route
       k = RIB-API or Forwarding Policy backup hop
===============================================================================
```

The BGP label is assigned by the next hop, in this case by the local ASBR PE-4. This IPv4 label can be seen with following command on PE-1:

```
[/]
A:admin@PE-1# show router bgp routes 192.0.2.5/32 label-ipv4 hunt
===============================================================================
 BGP Router ID:192.0.2.1        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP LABEL-IPV4 Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : 192.0.2.5/32
Nexthop        : 192.0.2.4
Path Id        : None
From           : 192.0.2.3
Res. Nexthop   : 192.0.2.4 (LDP)
Local Pref.    : 100                    Interface Name : NotAvailable
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : 20
Connector      : None
Community      : No Community Members
Cluster        : 192.0.2.3
Originator Id  : 192.0.2.4              Peer Router Id : 192.0.2.3
Fwd Class      : None                   Priority       : None
IPv4 Label     : 524283
Flags          : Used Valid Best IGP In-TTM In-RTM
Route Source   : Internal
AS-Path        : 64497
Route Tag      : 0
Neighbor-AS    : 64497
Orig Validation: NotFound
Source Class   : 0                      Dest Class     : 0
Add Paths Send : Default
RIB Priority   : Normal
Last Modified  : 00h10m56s

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

This BGP label is swapped by ASBR PE-4 in AS 64496 and by ASBR PE-8 in AS 64497.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The bottom label is the VPN label assigned by the remote PE in the remote AS for the destination network. This VPN label is retrieved on PE-1, as follows:

```
[/]
A:admin@PE-1# show router bgp routes 10.2.2.2/32 vpn-ipv4 hunt
===============================================================================
 BGP Router ID:192.0.2.1        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv4 Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : 10.2.2.2/32
Nexthop        : 192.0.2.5
Route Dist.    : 64497:1               VPN Label      : 524282
Path Id        : None
From           : 192.0.2.3
Res. Nexthop   : n/a
Local Pref.    : 100                   Interface Name : NotAvailable
Aggregator AS  : None                  Aggregator     : None
Atomic Aggr.   : Not Atomic            MED            : None
AIGP Metric    : None                  IGP Cost       : 0
Connector      : None
Community      : target:64497:1
Cluster        : No Cluster Members
Originator Id  : None                  Peer Router Id : 192.0.2.3
Fwd Class      : None                  Priority       : None
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : 64497
Route Tag      : 0
Neighbor-AS    : 64497
Orig Validation: N/A
Source Class   : 0                     Dest Class     : 0
Add Paths Send : Default
Last Modified  : 00h02m23s
VPRN Imported  :  1

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

Conclusion

Inter-AS option C allows the delivery of Layer 3 VPN services to customers who have sites connected in different ASs. This example shows the configuration of inter-AS option C (specific to this feature) together with the associated show output which can be used for verification and troubleshooting.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Layer 3 VPN: VPRN Type Spoke

This chapter provides information about Layer 3 VPRN CE hub and spoke architecture.

Topics in this chapter include:

Applicability

This chapter was initially written for SR OS Release 12.0. However, the MD-CLI in the current edition is based on SR OS Release 22.2.R1.

Knowledge of Nokia's Layer 3 VPN concepts is assumed throughout this document.

Overview

This chapter provides a basic technology overview and configuration examples of a network topology used for a Layer 3 VPRN CE hub and spoke architecture.

VPRN type hub

In SR OS releases earlier than 12.0, a CE hub and spoke architecture was partially supported. Internal optimization was available for the hub sites connected to the same PE router only. This feature is known as VPRN type hub. If, on the other hand, multiple spoke sites were connected to the same PE router, separate VPRN instances had to be created to maintain the split horizon forwarding behavior. This approach was complex, hard to maintain and consumed extra VPRN instances.

VPRN type spoke

Release 12.0.R1 added new functionality to overcome these limitations. Introducing the VPRN type spoke feature allows multiple spoke sites to be kept within the same VPRN instance while at the same time maintaining the split horizon approach such that spoke sites cannot send traffic directly to each other.

The primary goal of the feature is to allow multiple spoke sites to be part of a single VPRN instance without allowing direct communication between the spoke CE sites which are part of that VPRN (of type spoke). The packet flow is demonstrated in Figure 310: CE hub and spoke data path.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 310: CE hub and spoke data path

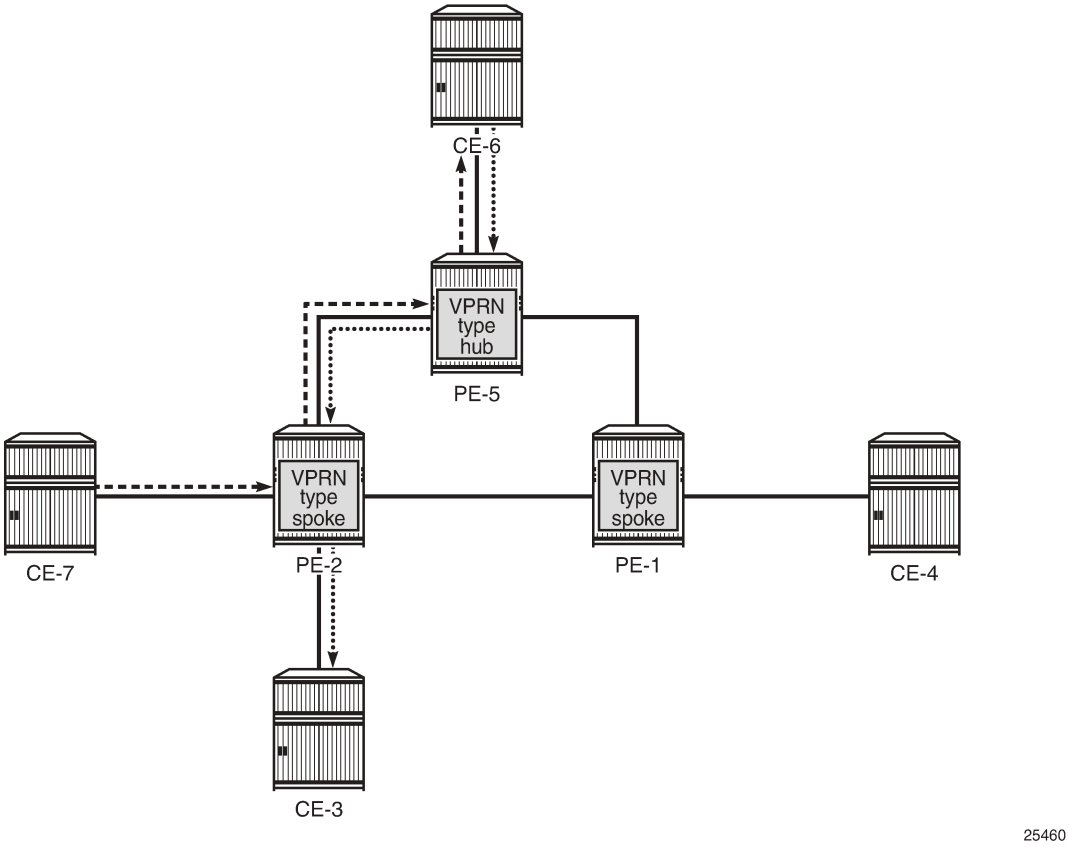

The only way for CE-7 to communicate with CE-3 is via hub site CE-6. The same applies to the communication between CE-7 and CE-4. The VPRN on PE-2 is configured as **vprn-type spoke** and has IP interfaces using SAPs or spoke SDPs that are considered spoke sites only. No direct communication between any of the spoke CE sites in the network is allowed.

Direct communication between the spoke CE sites is blocked using two techniques, as illustrated in Figure 311: CE hub and spoke control plane isolation.

- Using the **vprn-type spoke** command under the **vprn** context as explained later.
- The extended community configuration using route-target policies (this is not covered in detail in this chapter).

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 311: CE hub and spoke control plane isolation

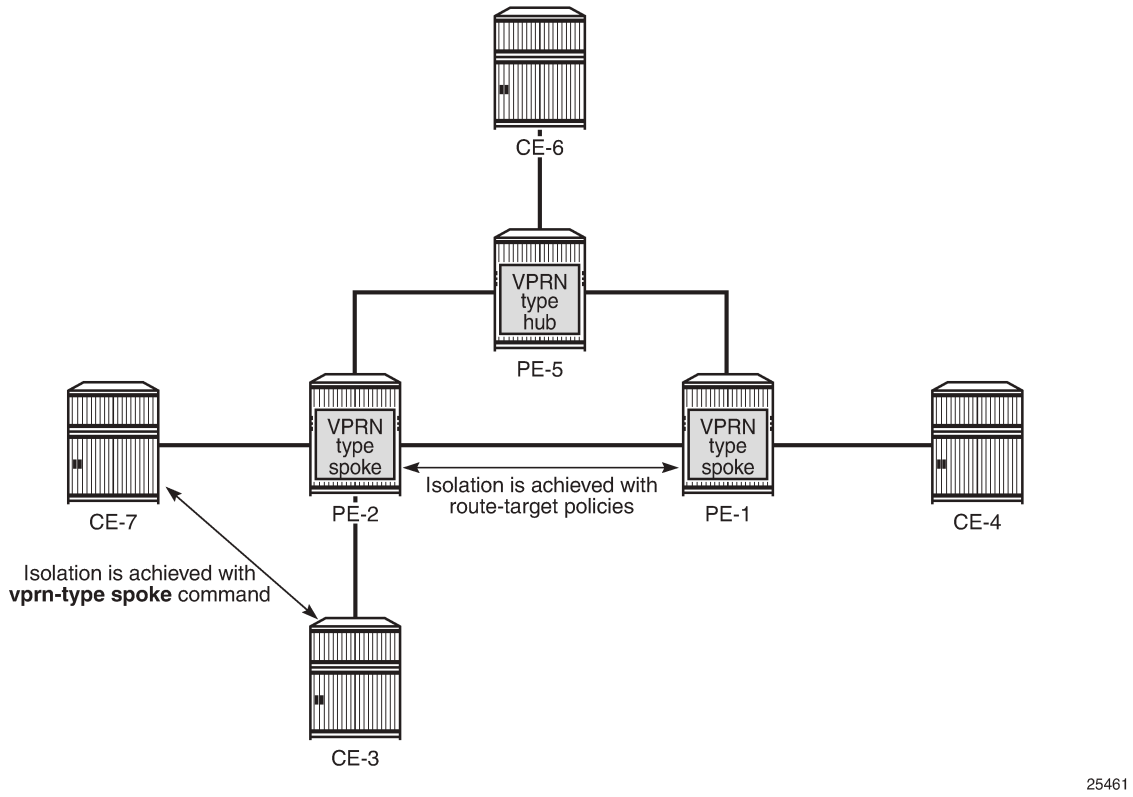

When a VPRN on a PE router is configured as **vprn-type spoke**, then the internal forwarding logic changes as demonstrated in Figure 312: Internal VPRN logic on a PE router.

Figure 312: Internal VPRN logic on a PE router

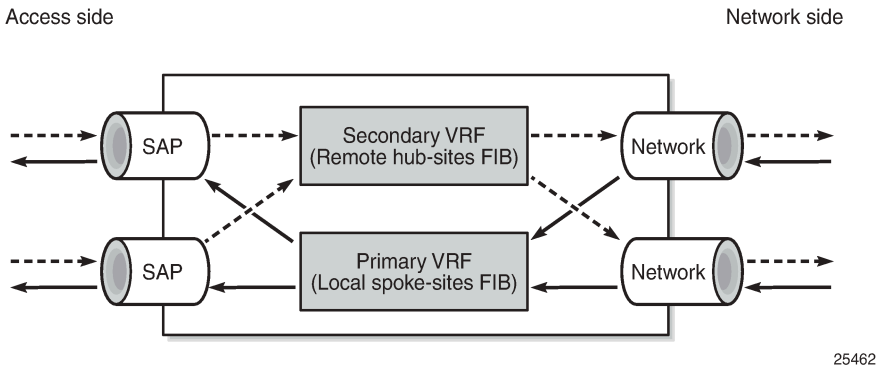

- VPRNs of type spoke create a primary and a secondary VRF internally to the VPRN:
  - The primary VRF is used for forwarding traffic from the network interfaces toward the IP interfaces using SAPs or spoke SDPs. This VRF is populated with routes learned from the spoke CE sites connected to the local PE through IP interfaces using SAPs or spoke SDPs.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

  - The secondary VRF is used for forwarding traffic from the IP interfaces using SAPs or spoke SDPs toward the network interfaces or other VPRN instances. This VRF is populated with routes learned via MP-BGP from hub sites.
- VPRNs of type spoke export routes using a specific extended community (for instance, spoke-ext-comm) via an export policy to identify them as spoke site originated routes.
  - This community is not hard-coded and has to be configured manually (see configuration example later).
- VPRNs of type spoke import routes (using an import policy) received from other PEs or VPRN instances with a hub specific community only (for example, hub-ext-comm). Routes with spoke-ext-comm community are ignored.
  - This community is not hard-coded and has to be configured manually (see configuration example later).
- Multiple VPRNs of type spoke and hub can coexist on the same PE if they use different VPRN instances.
- The configuration of type hub and type spoke is mutually exclusive within one VPRN instance.

Configuration

The physical topology and addressing scheme are presented in Figure 313: CE hub and spoke topology and addressing scheme.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 313: CE hub and spoke topology and addressing scheme

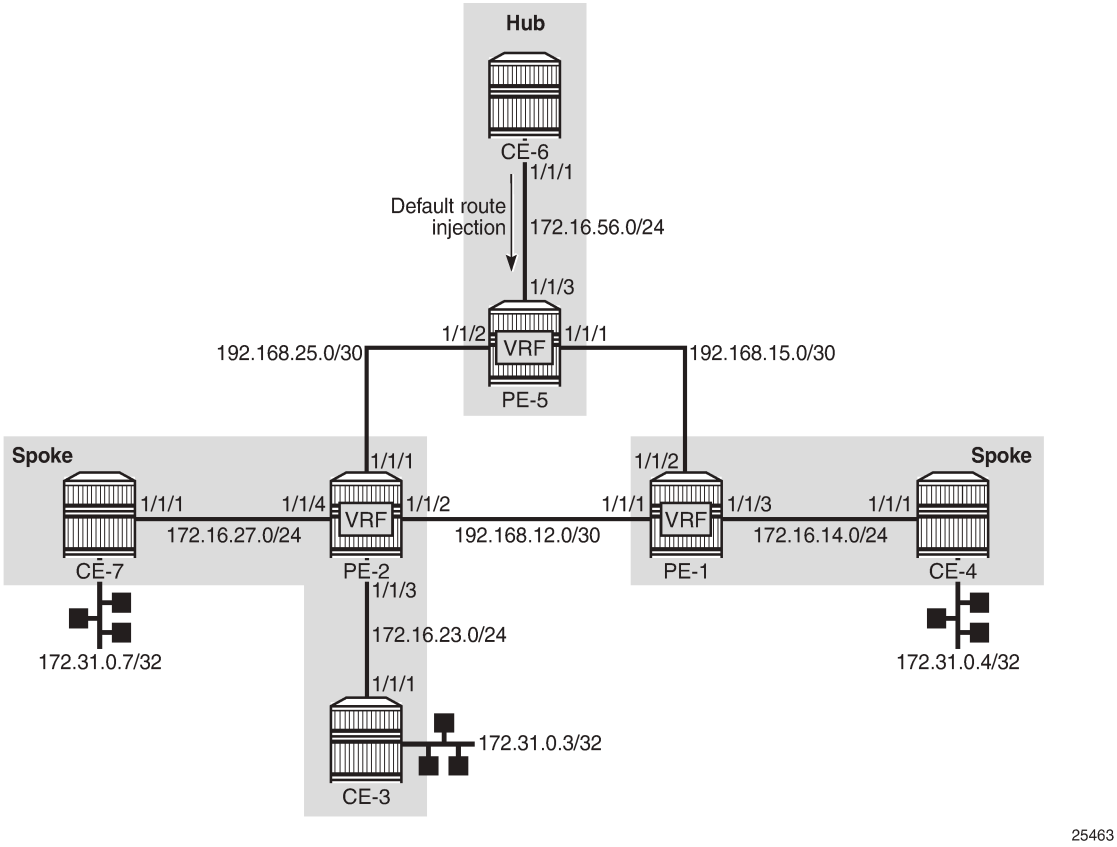

The configuration of PE-2 and PE-5 are the main focus of this example. The configuration of PE-1 is similar to that of PE-2.

Hub site configuration

Only the essential part of the configuration is provided for the hub site.

Vrf-import and vrf-export policies are used to manipulate the vrf-target in order to achieve logical isolation between the spoke sites in the network.

```
# on PE-5:
configure {
    policy-options {
        community "hub-ext-comm" {
            member "target:64500:11" { }
        }
        community "spoke-ext-comm" {
            member "target:64500:12" { }
        }
        policy-statement "export-ospf" {
            entry 10 {
                from {
                    protocol {
                        name [direct]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    }
                }
                action {
                    action-type accept
                }
            }
            default-action {
                action-type accept
            }
        }
        policy-statement "vrf-export" {
            default-action {
                action-type accept
                community {
                    add ["hub-ext-comm"]
                }
            }
        }
        policy-statement "vrf-import" {
            entry 10 {
                from {
                    community {
                        name "spoke-ext-comm"
                    }
                }
                action {
                    action-type accept
                }
            }
            default-action {
                action-type reject
            }
        }
```

PE-5 is configured with VPRN "VPRN1" providing OSPF connectivity to customer CE-6.

```
# on PE-5:
configure {
    service {
        vprn "VPRN1" {
            admin-state enable
            description "VPRN type hub"
            service-id 1
            customer "1"
            vprn-type hub
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64500:15"
                    vrf-import {
                        policy ["vrf-import"]
                    }
                    vrf-export {
                        policy ["vrf-export"]
                    }
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            interface "int-PE-5-CE-6" {
                ipv4 {
                    primary {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        address 172.16.56.1
                        prefix-length 24
                    }
                }
                sap 1/1/3:1 {
                }
            }
            ospf 0 {
                admin-state enable
                export-policy ["export-ospf"]
                area 0.0.0.0 {
                    interface "int-PE-5-CE-6" {
                        interface-type point-to-point
                        mtu 1500
                    }
                }
            }
```

At the same time, CE-6 is configured to advertise a default route which is used by all remote spoke CE sites to forward traffic via CE-6.

```
# on CE-6:
configure {
    policy-options {
        policy-statement "export-ospf-default" {
            entry 10 {
                from {
                    protocol {
                        name [static]
                    }
                }
                action {
                    action-type accept
                }
            }
        }
    }
    service {
        vprn "VPRN1" {
            admin-state enable
            service-id 1
            customer "1"
            interface "int-CE-6-PE-5" {
                ipv4 {
                    primary {
                        address 172.16.56.2
                        prefix-length 24
                    }
                }
                sap 1/1/1:1 {
                }
            }
            static-routes {
                route 0.0.0.0/0 route-type unicast {
                    blackhole {
                        admin-state enable
                    }
                }
            }
            ospf 0 {
                admin-state enable
                router-id 192.0.2.6
                export-policy ["export-ospf-default"]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                ignore-dn-bit true
                suppress-dn-bit true
                area 0.0.0.0 {
                    interface "int-CE-6-PE-5" {
                        interface-type point-to-point
                        mtu 1500
                    }
                }
            }
        }
```

Spoke site configuration

According to the example topology, two spoke VPRNs are present: one VPRN with two CE spoke sites connected is located on PE-2, and another VPRN with one spoke CE site on PE-1. The service configuration for PE-2 is as follows with the one for PE-1 being similar.

Vrf-import and vrf-export policies are used to build a hub-and-spoke topology in order to achieve a logical isolation between spoke sites connected to different PE routers.

```
# on PE-2:
configure {
    policy-options {
        community "hub-ext-comm" {
            member "target:64500:11" { }
        }
        community "spoke-ext-comm" {
            member "target:64500:12" { }
        }
        policy-statement "export-ospf" {
            default-action {
                action-type accept
            }
        }
        policy-statement "vrf-export" {
            default-action {
                action-type accept
                community {
                    add ["spoke-ext-comm"]
                }
            }
        }
        policy-statement "vrf-import" {
            entry 10 {
                from {
                    community {
                        name "hub-ext-comm"
                    }
                }
                action {
                    action-type accept
                }
            }
            default-action {
                action-type reject
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

PE-2 is configured with VPRN 1, which has OSPF connectivity to the customer CE-3 and CE-7. The **vprn-type spoke** command is used to prevent direct CE spoke to CE spoke communications for this VPRN.

```
# on PE-2:
configure {
    service {
        vprn "VPRN1" {
            admin-state enable
            description "VPRN type spoke"
            service-id 1
            customer "1"
            vprn-type spoke
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64500:12"
                    vrf-import {
                        policy ["vrf-import"]
                    }
                    vrf-export {
                        policy ["vrf-export"]
                    }
                    auto-bind-tunnel {
                        resolution any
                    }
                }
            }
            interface "int-PE-2-CE-3" {
                ipv4 {
                    primary {
                        address 172.16.23.1
                        prefix-length 24
                    }
                }
                sap 1/1/3:1 {
                }
            }
            interface "int-PE-2-CE-7" {
                ipv4 {
                    primary {
                        address 172.16.27.1
                        prefix-length 24
                    }
                }
                sap 1/1/4:1 {
                }
            }
            ospf 0 {
                admin-state enable
                export-policy ["export-ospf"]
                area 0.0.0.0 {
                    interface "int-PE-2-CE-3" {
                        interface-type point-to-point
                        mtu 1500
                    }
                    interface "int-PE-2-CE-7" {
                        interface-type point-to-point
                        mtu 1500
                    }
                }
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

For connectivity verification purposes, CE-3, CE-4, and CE-7 are configured to advertise their internal loopback interfaces via OSPF:

- CE-3 advertises 172.31.0.3/32
- CE-4 advertises 172.31.0.4/32
- CE-7 advertises 172.31.0.7/32

```
# on CE-3:
configure {
    service {
        vprn "VPRN1" {
            admin-state enable
            service-id 1
            customer "1"
            interface "int-CE-3-PE-2" {
                ipv4 {
                    primary {
                        address 172.16.23.2
                        prefix-length 24
                    }
                }
                sap 1/1/1:1 {
                }
            }
            interface "lo0" {
                loopback true
                ipv4 {
                    primary {
                        address 172.31.0.3
                        prefix-length 32
                    }
                }
            }
            ospf 0 {
                admin-state enable
                router-id 192.0.2.3
                ignore-dn-bit true
                suppress-dn-bit true
                area 0.0.0.0 {
                    interface "int-CE-3-PE-2" {
                        interface-type point-to-point
                        mtu 1500
                    }
                    interface "lo0" {
                    }
                }
            }
        }
```

Hub site verification

The Routing Information Base (RIB) for VPRN 1 on hub site PE-5 lists all reachable networks:

```
[/]
A:admin@PE-5# show router 1 route-table

===============================================================================
Route Table (Service: 1)
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Dest Prefix[Flags]                                Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
0.0.0.0/0                                         Remote  OSPF      00h02m32s  150
       172.16.56.2                                                     1
172.16.14.0/24                                    Remote  BGP VPN   00h01m07s  170
       192.0.2.1 (tunneled)                                            10
172.16.14.1/32                                    Remote  BGP VPN   00h01m07s  170
       192.0.2.1 (tunneled)                                            10
172.16.23.0/24                                    Remote  BGP VPN   00h01m01s  170
       192.0.2.2 (tunneled)                                            10
172.16.23.1/32                                    Remote  BGP VPN   00h01m01s  170
       192.0.2.2 (tunneled)                                            10
172.16.27.0/24                                    Remote  BGP VPN   00h01m01s  170
       192.0.2.2 (tunneled)                                            10
172.16.27.1/32                                    Remote  BGP VPN   00h01m01s  170
       192.0.2.2 (tunneled)                                            10
172.16.56.0/24                                    Local   Local     00h03m57s  0
       int-PE-5-CE-6                                                   0
172.31.0.3/32                                     Remote  BGP VPN   00h01m01s  170
       192.0.2.2 (tunneled)                                            10
172.31.0.4/32                                     Remote  BGP VPN   00h00m40s  170
       192.0.2.1 (tunneled)                                            10
172.31.0.7/32                                     Remote  BGP VPN   00h00m30s  170
       192.0.2.2 (tunneled)                                            10
-------------------------------------------------------------------------------
No. of Routes: 11
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

The forwarding table (FIB) for the primary VRF of VPRN 1 is displayed using following command. All remote spoke and hub sites are reachable via this VRF.

```
[/]
A:admin@PE-5# show router 1 fib 1

===============================================================================
FIB Display
===============================================================================
Prefix [Flags]                                                  Protocol
  NextHop
-------------------------------------------------------------------------------
0.0.0.0/0                                                       OSPF
  172.16.56.2 (int-PE-5-CE-6)
172.16.14.0/24                                                  BGP_VPN
  192.0.2.1 (VPRN Label:524284 Transport:LDP)
172.16.14.1/32                                                  BGP_VPN
  192.0.2.1 (VPRN Label:524284 Transport:LDP)
172.16.23.0/24                                                  BGP_VPN
  192.0.2.2 (VPRN Label:524284 Transport:LDP)
172.16.23.1/32                                                  BGP_VPN
  192.0.2.2 (VPRN Label:524284 Transport:LDP)
172.16.27.0/24                                                  BGP_VPN
  192.0.2.2 (VPRN Label:524284 Transport:LDP)
172.16.27.1/32                                                  BGP_VPN
  192.0.2.2 (VPRN Label:524284 Transport:LDP)
172.16.56.0/24                                                  LOCAL
  172.16.56.0 (int-PE-5-CE-6)
172.31.0.3/32                                                   BGP_VPN
  192.0.2.2 (VPRN Label:524284 Transport:LDP)
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
172.31.0.4/32                                                   BGP_VPN
  192.0.2.1 (VPRN Label:524284 Transport:LDP)
172.31.0.7/32                                                   BGP_VPN
  192.0.2.2 (VPRN Label:524284 Transport:LDP)
-------------------------------------------------------------------------------
Total Entries : 11
-------------------------------------------------------------------------------
===============================================================================
```

The forwarding table for the secondary VRF of VPRN 1 is displayed using following command, including the **secondary** keyword. All local hub CE sites are reachable via this VRF.

```
[/]
A:admin@PE-5# show router 1 fib 1 secondary

===============================================================================
FIB Display
===============================================================================
Prefix [Flags]                                                  Protocol
  NextHop
-------------------------------------------------------------------------------
0.0.0.0/0                                                       OSPF
  172.16.56.2 (int-PE-5-CE-6)
172.16.56.0/24                                                  LOCAL
  172.16.56.0 (int-PE-5-CE-6)
-------------------------------------------------------------------------------
Total Entries : 2
-------------------------------------------------------------------------------
===============================================================================
```

Spoke site verification

The RIB for VPRN 1 on PE-2 (spoke VPRN) lists all reachable networks.

The other spoke sites connected to the remote PEs are not present in the routing table, in this example, CE-4 with prefixes such as 172.31.0.4/32 and 172.16.14.0/24.

The local interface addresses of PE-2 (172.16.23.1/32 and 172.16.27.1/32) are present in the routing table of VPRN 1, as follows. From a FIB point of view, these are reachable from any spoke VPRN, but the spoke CE's router host addresses are not. This fact does not influence the data plane isolation for the customer networks.

```
[/]
A:admin@PE-2# show router 1 route-table

===============================================================================
Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                          Type    Proto     Age        Pref
      Next Hop[Interface Name]                                Metric
-------------------------------------------------------------------------------
0.0.0.0/0                                   Remote  BGP VPN   00h22m49s  170
       192.0.2.5 (tunneled)                                       10
172.16.23.0/24                              Local   Local     00h22m53s  0
       int-PE-2-CE-3                                              0
172.16.23.1/32                              Local   Host      00h22m53s  0
       int-PE-2-CE-3                                              0
172.16.27.0/24                              Local   Local     00h22m53s  0
       int-PE-2-CE-7                                              0
172.16.27.1/32                              Local   Host      00h22m53s  0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
       int-PE-2-CE-7                                                    0
172.16.56.0/24                          Remote  BGP VPN   00h22m49s  170
       192.0.2.5 (tunneled)                                             10
172.16.56.1/32                          Remote  BGP VPN   00h22m49s  170
       192.0.2.5 (tunneled)                                             10
172.31.0.3/32                           Remote  OSPF      00h22m40s  10
       172.16.23.2                                                      10
172.31.0.7/32                           Remote  OSPF      00h22m24s  10
       172.16.27.2                                                      10
-------------------------------------------------------------------------------
No. of Routes: 9
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

The FIB for the primary VRF of VPRN 1 shows all local spoke sites are reachable via this VRF, as follows:

```
[/]
A:admin@PE-2# show router 1 fib 1

===============================================================================
FIB Display
===============================================================================
Prefix [Flags]                                                  Protocol
  NextHop
-------------------------------------------------------------------------------
172.16.23.0/24                                                  LOCAL
  172.16.23.0 (int-PE-2-CE-3)
172.16.23.1/32                                                  HOST
  Blackhole
172.16.27.0/24                                                  LOCAL
  172.16.27.0 (int-PE-2-CE-7)
172.16.27.1/32                                                  HOST
  Blackhole
172.31.0.3/32                                                   OSPF
  172.16.23.2 (int-PE-2-CE-3)
172.31.0.7/32                                                   OSPF
  172.16.27.2 (int-PE-2-CE-7)
-------------------------------------------------------------------------------
Total Entries : 6
-------------------------------------------------------------------------------
===============================================================================
```

The FIB for the secondary VRF of VPRN 1 shows the remote hub site (address 172.16.56.0/24) is reachable via this VRF, as follows:

```
[/]
A:admin@PE-2# show router 1 fib 1 secondary

===============================================================================
FIB Display
===============================================================================
Prefix [Flags]                                                  Protocol
  NextHop
-------------------------------------------------------------------------------
0.0.0.0/0                                                       BGP_VPN
  192.0.2.5 (VPRN Label:524284 Transport:LDP)
172.16.23.1/32                                                  HOST
  Blackhole
172.16.27.1/32                                                  HOST
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
  Blackhole
172.16.56.0/24                                                  BGP_VPN
  192.0.2.5 (VPRN Label:524284 Transport:LDP)
172.16.56.1/32                                                  BGP_VPN
  192.0.2.5 (VPRN Label:524284 Transport:LDP)
-------------------------------------------------------------------------------
Total Entries : 5
-------------------------------------------------------------------------------
===============================================================================
```

Spoke sites connectivity verification

Without the VPRN spoke type configuration in VPRN 1 on PE-2, CE-3 takes the shortest path to CE-7, which violates the hub-and-spoke design approach explained earlier.

```
# on PE-2:
configure exclusive
    service {
        vprn "VPRN1" {
            delete vprn-type
            commit
```

Note:
In this setup, a VPRN is configured on CE-3, but that is not necessary.

Traffic from CE-3 takes the shortest path to CE-7, because VPRN 1 on PE-2 is not configured with spoke type anymore.

```
[/]
A:admin@CE-3# traceroute 172.31.0.7 router-instance "VPRN1" numeric
traceroute to 172.31.0.7, 30 hops max, 40 byte packets
  1  172.16.23.1    2.54 ms  2.80 ms  2.93 ms
  2  172.31.0.7    3.28 ms  3.23 ms  3.05 ms
```

After enabling the **vprn-type spoke** feature on PE-2, CE-3 takes the longest path via hub CE-6 to reach CE-7, as it should.

```
# on PE-2:
configure exclusive
    service {
        vprn "VPRN1" {
            vprn-type spoke
            commit
```

```
[/]
A:admin@CE-3# traceroute 172.31.0.7 router-instance "VPRN1" numeric
traceroute to 172.31.0.7, 30 hops max, 40 byte packets
  1  172.16.23.1    2.34 ms  2.71 ms  2.58 ms
  2  0.0.0.0  * * *
  3  172.16.56.2    4.64 ms  4.62 ms  3.95 ms
  4  172.16.56.1    4.24 ms  4.44 ms  4.76 ms
  5  172.16.27.1    4.31 ms  4.30 ms  4.20 ms
  6  172.31.0.7    6.32 ms  6.39 ms  6.36 ms
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Similarly, the long path is taken by CE-3 to reach CE-4, as follows. This is unrelated to the VPRN type. It is achieved by policies.

```
[/]
A:admin@CE-3# traceroute 172.31.0.4 router-instance "VPRN1" numeric
traceroute to 172.31.0.4, 30 hops max, 40 byte packets
  1  172.16.23.1    2.39 ms  2.80 ms  2.58 ms
  2  0.0.0.0  * * *
  3  172.16.56.2    4.69 ms  4.46 ms  4.23 ms
  4  172.16.56.1    4.53 ms  4.61 ms  4.12 ms
  5  172.16.14.1    5.83 ms  6.02 ms  6.24 ms
  6  172.31.0.4    7.00 ms  6.54 ms  6.99 ms
```

Conclusion

The VPRN type spoke feature completes the CE hub and spoke solution. It brings an additional level of simplicity, scalability, and flexibility to operators using this VPRN architecture for their customers.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Spoke Termination for IPv6-6VPE

This chapter provides information about spoke termination for IPv6-6VPE.

Topics in this chapter include:

Applicability

This chapter was originally written for SR OS Release 8.0, where Epipe Virtual Leased Line (VLL) is supported for IPv6 spoke termination within a Virtual Private Routed Network (VPRN). The MD-CLI in the current edition corresponds to SR OS Release 21.10.R2.

Overview

RFC 4659, *BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN*, standardized the use of an IPv6 over IPv4 tunneling scheme. SR OS supports the standardized IPv6 over IPv4 tunneling scheme for VPRN services using Multi-Protocol Border Gateway Protocol (MP-BGP), also known as 6VPE.

SR OS supports pseudowire termination by a VPRN from an Epipe Virtual Leased Line (VLL) or VPLS spoke Service Distribution Point (spoke SDP) where the pseudowire can be given IPv6 addresses and run IPv6 protocols. In the example used in this chapter, any advertisements across the Multi-Protocol Labeled Switching (MPLS) network between VPRN Provider Edge (PE) devices will use 6VPE. This chapter describes the configuration for IPv6 spoke termination to a VPRN over an Epipe VLL and transporting IPv6 packets over 6VPE tunnels between PE devices.

This solution can be used where a service provider is providing VPRN services built on a transport network whose Interior Gateway Protocol (IGP) is using IPv4 addressing on the network interfaces. The customer's CE and the service provider's PE must support IPv6 pseudowires, IPv6 interfaces and in addition, the service provider also must be able to support the advertisements of IPv6 prefixes between CE-PE peerings and between the transport PE routers using MP-BGP. The advertisement of IPv6 prefixes across the MPLS network and the transport of IPv6 traffic is tunneled using 6VPE.

The VPRN PE supports spoke termination of Epipe VLL services on access with IPv6 addressing between the CE and VPRN PE. The IPv6 spoke termination on VPRN services has the same functionality as VPRN IPv4 spoke termination.

The example in Figure 314: Spoke termination for IPv6 illustrates a CE device that connects to a VPRN PE on an IPv6 interface addressing using spoke termination.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 314: Spoke termination for IPv6

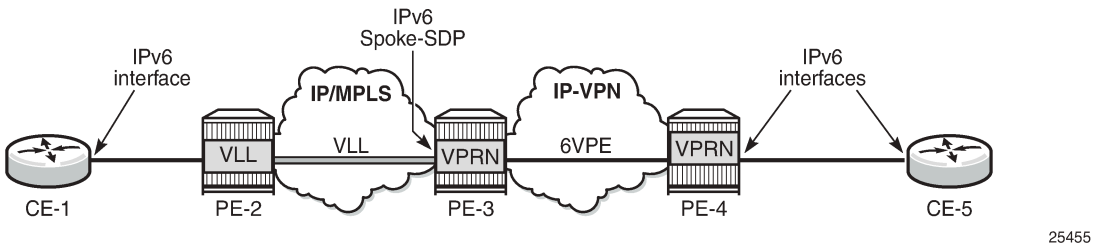

CE-1 is connected to the VPRN service on PE-3, using IPv6 interfaces. CE-1 reaches PE-3 by connecting to PE-2. PE-2 uses an Epipe VLL for transport to the VPRN on the PE-3. The connectivity between the VLL service on the VPRN service on PE-3 is using spoke termination with IPv6 addressing on the spoke SDP interface on PE-3.

Figure 315: IPv6 addressing and IPv6 prefixes

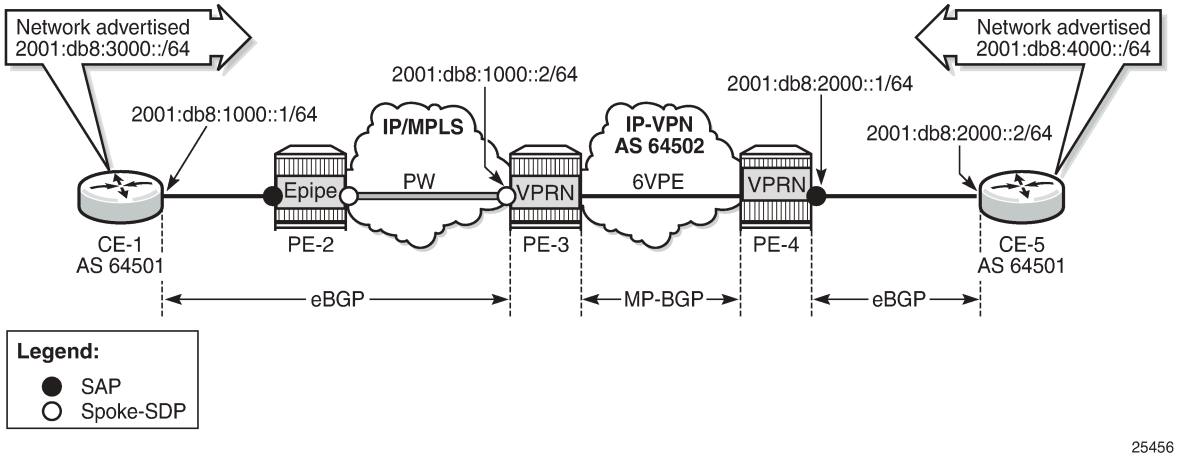

Figure 315: IPv6 addressing and IPv6 prefixes shows the overall IPv6 addressing from interfaces to prefixes advertised from CE-1 and CE-5 across the VPRN network.

- Link between CE-1 and PE-3: 2001:db8:1000::/64
- Link between CE-5 and PE-4: 2001:db8:2000::/64
- Advertised prefix from CE-1: 2001:db8:3000::/64
- Advertised prefix from CE-5: 2001:db8:4000::/64

PE-3 has an MP-eBGP session with CE-1 to receive and advertise IPv6 routes. PE-3 also has an MP-iBGP peering session with PE-4 to use 6VPE to tunnel IPv6 routes and traffic to and from PE-4. PE-4 has an IPv6 SAP interface to CE-5 and uses MP-eBGP to advertise to and receive routes from CE-5 (no spoke termination). The configuration of PE-3 is included to provide examples of the end-to-end VPRN service using a 6VPE model.

This network topology illustrates the use of spoke termination using IPv6 interfaces and the tunneling of IPv6 traffic over a 6VPE MPLS network.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Configuration

First an MPLS network is established where the VPRN service can use 6VPE to tunnel traffic across the IPv4 IGP.

Figure 316: MP-BGP VPN IPv6

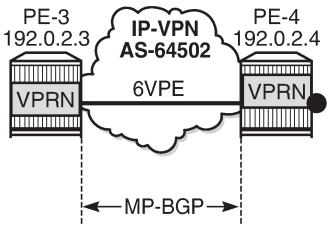

In Figure 316: MP-BGP VPN IPv6, PE-3 and PE-4 are edge routers running VPRN services on access with IPv6 interfaces. The MPLS network is configured using IPv4 link addressing. Interior Border Gateway Protocol (iBGP) peerings need to be established with MP-BGP for the VPN-IPv6 address family between PE-3 and PE-4.

```
# on PE-3:
configure {
    router "Base" {
        autonomous-system 64502
        bgp {
            group "iBGP" {
                description "iBGP peering in AS 64502"
                peer-as 64502
                family {
                    vpn-ipv6 true
                }
            }
            neighbor "192.0.2.4" {
                description "PE-4"
                group "iBGP"
            }
        }
```

```
# on PE-4:
configure {
    router "Base" {
        autonomous-system 64502
        bgp {
            group "iBGP" {
                description "iBGP peering in AS 64502"
                peer-as 64502
                family {
                    vpn-ipv6 true
                }
            }
            neighbor "192.0.2.3" {
                description "PE-3"
                group "iBGP"
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Configuring address family VPN-IPv6 between VPRN PE edge routers in BGP enables MP-BGP for the Layer 3 VPNs supporting the customer's IPv6 addressing (6VPE).

The following commands verify the BGP sessions for the VPN-IPv6 address family between PE-3 and PE-4:

```
[/]
A:admin@PE-3# show router bgp neighbor 192.0.2.4

===============================================================================
BGP Neighbor
===============================================================================
-------------------------------------------------------------------------------
Peer                 : 192.0.2.4
Description          : PE-4
Group                : iBGP
-------------------------------------------------------------------------------
Peer AS              : 64502            Peer Port            : 49727
Peer Address         : 192.0.2.4
Local AS             : 64502            Local Port           : 179
Local Address        : 192.0.2.3
Peer Type            : Internal         Dynamic Peer         : No
State                : Established      Last State           : Established
Last Event           : recvOpen
Last Error           : Cease (Connection Collision Resolution)
Local Family         : VPN-IPv6
Remote Family        : VPN-IPv6
---snip---
```

```
[/]
A:admin@PE-4# show router bgp neighbor 192.0.2.3

===============================================================================
BGP Neighbor
===============================================================================
-------------------------------------------------------------------------------
Peer                 : 192.0.2.3
Description          : PE-3
Group                : iBGP
-------------------------------------------------------------------------------
Peer AS              : 64502            Peer Port            : 179
Peer Address         : 192.0.2.3
Local AS             : 64502            Local Port           : 49727
Local Address        : 192.0.2.4
Peer Type            : Internal         Dynamic Peer         : No
State                : Established      Last State           : Active
Last Event           : recvOpen
Last Error           : Unrecognized Error
Local Family         : VPN-IPv6
Remote Family        : VPN-IPv6
---snip---
```

After the MP-BGP sessions are established for the VPN-IPv6 address-family, 6VPE tunnel support is provided between PE-3 and PE-4.

Figure 317: Spoke termination for IPv6 addressing illustrates the model for spoke termination for IPv6 using VPRN services.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 317: Spoke termination for IPv6 addressing

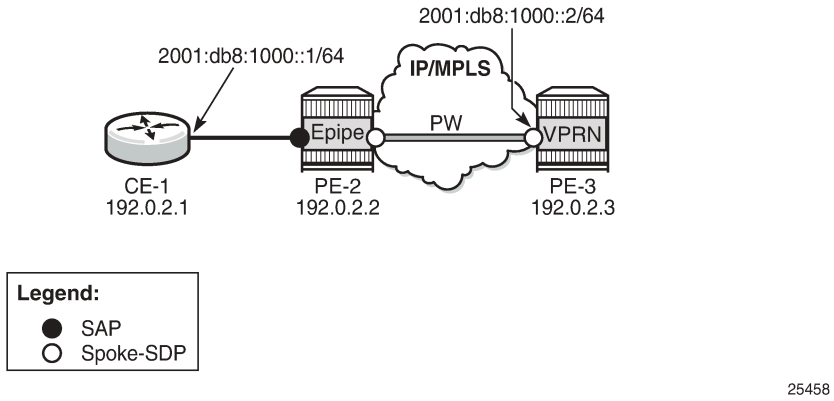

CE-1 is configured with IPv6 addressing on the access interface facing the VPRN service. CE-1's access is backhauled to the VPRN service on PE-3 using Epipe VLL with spoke termination. The configuration of the Epipe VLL on PE-2 is as follows:

```
# on PE-2:
configure {
    service {
        epipe "Epipe 1" {
            admin-state enable
            service-id 1
            customer "1"
            service-mtu 9190
            spoke-sdp 231:1 {
            }
            sap 1/1/2 {
            }
        }
        sdp 231 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.3
            }
        }
```

Epipe 1 on PE-2 is configured with a SAP interface facing the customer and a spoke SDP facing PE-3. The spoke SDP is terminated into the customer's VPRN service on PE-3.

The possible IPv6 options for spoke SDP interfaces on the CLI for VPRN services are as follows (compliant with RFC 4213, *Basic Transition Mechanisms for IPv6 Hosts and Routers <draft-ietf-v6ops-mech-v2-07.txt>*):

- Interface spoke SDP (IPv6 options only)

```
[ex:/configure service vprn "VPRN 1" interface "int-PE-3-PE-2"]
A:admin@PE-3# ?
---snip---
 ipv6                 + Enter the ipv6 context
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
---snip---
```

```
[ex:/configure service vprn "VPRN 1" interface "int-PE-3-PE-2" ipv6]
A:admin@PE-3# ?

 address               + Enter the address list instance
 bfd                   + Enter the bfd context
 dhcp6                 + Enter the dhcp6 context
 duplicate-address-    - Enable/disable Duplicate Address Detection
                       detection
 forward-ipv4-packets  - Forward unencapsulated IPv4 packets
 icmp6                 + Enter the icmp6 context
 link-local-address    + Enter the link-local-address context
 local-dhcp-server     - DHCP server for the interface
 neighbor-discovery    + Enter the neighbor-discovery context
 qos-route-lookup      - QoS Route lookup
 tcp-mss               - TCP maximum segment size for the interface
 urpf-check            + Enable the urpf-check context
 vrrp                  + Enter the vrrp list instance
```

- IPv6 address

```
[ex:/configure service vprn "VPRN 1" interface "int-PE-3-PE-2" ipv6]
A:admin@PE-3# address ?

 [ipv6-address] <ipv6-address>
 <ipv6-address>  - (<x:x:x:x:x:x:x:x>|<x:x:x:x:x:x:d.d.d.d>)

    IPv6 address assigned to the interface
```

- DHCPv6 relay parameters for the VPRN service

```
[ex:/configure service vprn "VPRN 1" interface "int-PE-3-PE-2" ipv6 dhcp6]
A:admin@PE-3# ?

 apply-groups          - Apply a configuration group at this level
 apply-groups-exclude  - Exclude a configuration group at this level
 relay                 + Enter the relay context
 server                + Enter the server context
```

```
[ex:/configure service vprn "VPRN 1" interface "int-PE-3-PE-2" ipv6 dhcp6 relay]
A:admin@PE-3# ?

 admin-state           - Administrative state of DHCPv6 Relay
 apply-groups          - Apply a configuration group at this level
 apply-groups-exclude  - Exclude a configuration group at this level
 description           - Text description
 lease-populate        + Enter the lease-populate context
 link-address          - Link address of the DHCPv6 relay messages
 neighbor-resolution   - Enable neighbor resolution via DHCPv6 relay
 option                + Enter the option context
 python-policy         - Python policy name
 server                - DHCPv6 server to which the DHCPv6 requests are forwarded
 source-address        - Source IPv6 address of the DHCPv6 relay messages
 user-db               - Local user database for authentication
```

- DHCPv6 server parameters for the VPRN service

```
[ex:/configure service vprn "VPRN 1" interface "int-PE-3-PE-2" ipv6 dhcp6 server]
A:admin@PE-3# ?
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
 apply-groups          - Apply a configuration group at this level
 apply-groups-exclude  - Exclude a configuration group at this level
 max-nbr-of-leases     - DHCPv6 leases allowed
 prefix-delegation     + Enter the prefix-delegation context
```

- ICMPv6

```
[ex:/configure service vprn "VPRN 1" interface "int-PE-3-PE-2" ipv6 icmp6]
A:admin@PE-3# ?

 packet-too-big        + Enter the packet-too-big context
 param-problem         + Enter the param-problem context
 redirects             + Enter the redirects context
 time-exceeded         + Enter the time-exceeded context
 unreachables          + Enter the unreachables context
```

- Link-local-addressing, for the VPRN interface. By default, link-local addressing is assigned dynamically. Use this command if you want to add a static link-local-address.

```
[ex:/configure service vprn "VPRN 1" interface "int-PE-3-PE-2" ipv6 link-local-address]
A:admin@PE-3# ?

 address               - IPv6 link local address
 duplicate-address-    - Enable Duplicate Address Detection
  detection
```

- Enabling local proxy neighbor discovery

```
*[ex:/configure service vprn "VPRN 1" interface "int-PE-3-PE-2" ipv6 neighbor-discovery]
A:admin@PE-3# local-proxy-nd ?

 local-proxy-nd <boolean>
 <boolean>  - ([true]|false)
 Default    - false

    Enable neighbor discovery on interface

 host-route            + Enter the host-route context
 learn-unsolicited     - Type of entries learned from unsolicited NA messages
 limit                 + Enter the limit context
 proactive-refresh     - Neighbor entries to be refreshed proactively
 proxy-nd-policy       - Name of the proxy Neighbor Discovery policies for the interface
 reachable-time        - Timer for neighbor reachability detection
 secure-nd             + Enter the secure-nd context
 stale-time            - Time during which a neighbor discovery cache entry remains stale
 static-neighbor       + Enter the static-neighbor list instance
```

- VRRP

```
*[ex:/configure service vprn "VPRN 1" interface "int-PE-3-PE-2" ipv6 vrrp 1]
A:admin@PE-3# ?

 Immutable fields      - owner, passive

 admin-state           - Administrative state of VRRP
 apply-groups          - Apply a configuration group at this level
 apply-groups-exclude  - Exclude a configuration group at this level
 backup                - Virtual router IP addresses for the interface
 bfd-liveness          + Enter the bfd-liveness context
 init-delay            - VRRP initialization delay timer
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
  mac                   - MAC address used by virtual router instance overriding the
                          VRRP default derived from the VRID
  master-int-inherit    - Allow the master instance to dictate the master down timer
  message-interval      - Interval for sending VRRP Advertisement messages
  ntp-reply             - Allow processing of NTP Requests
  oper-group            - Operational group name associated with VRRP
  owner                 - Designate virtual router instance as owning the virtual
                          router IP addresses
  passive               - Suppress the transmission and reception of VRRP
                          advertisement messages
  ping-reply            - Allow non-owner master to reply to ICMP echo requests
  policy                - VRRP priority control policy associated with the virtual
                          router instance
  preempt               - Allow the VRRP to override an existing non-owner master
  priority              - Base priority for the VRRP
  standby-forwarding    - Allow the forwarding of packets by a standby router
  telnet-reply          - Allow non-owner master to reply to Telnet requests
  traceroute-reply      - Allow non-owner master to reply to traceroute requests
```

The VPRN on PE-3 exports IPv6 routes (IPv6 route on CE-5) to CE-1 using the following route policy.

```
# on PE-3
configure {
    policy-options {
        prefix-list "PE-3-CE-1" {
            prefix 2001:db8:4000::/64 type exact {
            }
        }
        policy-statement "PE-3-BGP-CE-1" {
            entry 10 {
                from {
                    prefix-list ["PE-3-CE-1"]
                }
                action {
                    action-type accept
                    origin igp
                }
            }
            default-action {
                action-type reject
            }
        }
```

EBGP routes are discarded by default, so an import policy is required:

```
# on PE-3, PE-4, CE-1, CE-5:
configure {
    policy-options {
        prefix-list "2001:db8::/32" {
            prefix 2001:db8::/32 type longer {
            }
        }
        policy-statement "import-2001:db8::/32" {
            entry 10 {
                from {
                    prefix-list ["2001:db8::/32"]
                }
                action {
                    action-type accept
                    origin igp
                }
            }
            default-action {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                action-type reject
            }
        }
```

The configuration for the VPRN service on PE-3 with IPv6 interface (spoke SDP) as shown in Figure 317: Spoke termination for IPv6 addressing:

```
# on PE-3:
configure {
    service {
        sdp 321 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.2
            }
        }
        vprn "VPRN 1" {
            admin-state enable
            service-id 1
            customer "1"
            autonomous-system 64502
            router-id 192.0.2.31
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64502:1"
                    vrf-target {
                        community "target:64502:1"
                    }
                    auto-bind-tunnel {
                        resolution filter
                        resolution-filter {
                            ldp true
                        }
                    }
                }
            }
            bgp {
                router-id 192.0.2.31
                group "Spoke-CE-1-PE-3" {
                    peer-as 64501
                    local-address 2001:db8:1000::2
                    family {
                        ipv6 true
                    }
                }
                neighbor "2001:db8:1000::1" {
                    group "Spoke-CE-1-PE-3"
                    as-override true
                    type external
                    import {
                        policy ["import-2001:db8::/32"]
                    }
                    export {
                        policy ["PE-3-BGP-CE-1"]
                    }
                }
            }
            interface "int-PE-3-PE-2" {
                description "Spoke SDP"
                spoke-sdp 321:1 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
                ipv6 {
                    address 2001:db8:1000::2 {
                        prefix-length 64
                    }
                }
            }
            interface "loopback" {
                loopback true
                ipv4 {
                    primary {
                        address 192.0.2.31
                        prefix-length 32
                    }
                }
            }
        }
```

In the preceding configuration example, PE-3 has been configured with an IPv6 spoke SDP (spoke termination) with interface int-PE-3-PE-2. The VPRN configuration has also been set up for MP-eBGP peering to CE-1 through the IPv6 spoke interface. The MP-eBGP peering receives and advertises IPv6 prefixes from and to CE-1. The included route policy configuration shows how IPv6 routes are advertised to CE-1 from PE-3 (policy-statement PE-3-BGP-CE-1).

The configuration on PE-4 is similar, but with a SAP interface to CE-5 instead of a spoke-SDP.

Figure 318: PE-4 VPRN with SAP to CE-5

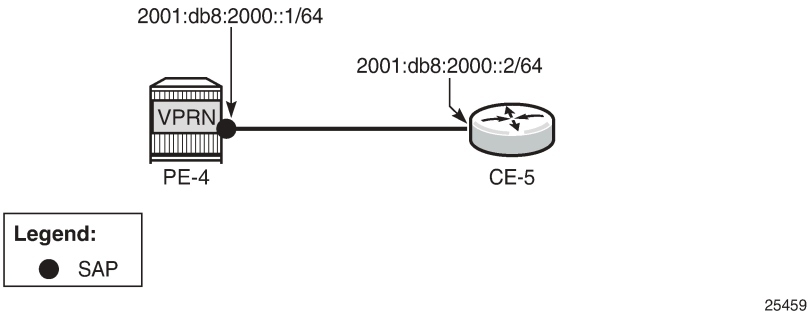

The IPv6 configuration options for the SAP interface (int-PE-4-CE-5) are similar to those in the preceding example for the spoke SDP on PE-3. The PE-4 BGP export policy (PE-4-BGP-CE-5) is also similar to the example for PE-3 in advertising the learned IPv6 route to CE-5.

```
# on PE-4:
configure {
    policy-options {
        prefix-list "PE-4-CE-5" {
            prefix 2001:db8:3000::/64 type exact {
            }
        }
        policy-statement "PE-4-BGP-CE-5" {
            entry 10 {
                from {
                    prefix-list ["PE-4-CE-5"]
                }
                action {
                    action-type accept
                    origin igp
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
            }
            default-action {
                action-type reject
            }
        }
```

```
configure {
    service {
        vprn "VPRN 1" {
            admin-state enable
            service-id 1
            customer "1"
            autonomous-system 64502
            router-id 192.0.2.41
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64502:1"
                    vrf-target {
                        community "target:64502:1"
                    }
                    auto-bind-tunnel {
                        resolution filter
                        resolution-filter {
                            ldp true
                        }
                    }
                }
            }
            bgp {
                router-id 192.0.2.41
                group "CE-5-PE-4" {
                    peer-as 64501
                    local-address 2001:db8:2000::1
                    family {
                        ipv6 true
                    }
                }
                neighbor "2001:db8:2000::2" {
                    group "CE-5-PE-4"
                    as-override true
                    type external
                    import {
                        policy ["import-2001:db8::/32"]
                    }
                    export {
                        policy ["PE-4-BGP-CE-5"]
                    }
                }
            }
            interface "int-PE-4-CE-5" {
                sap 1/1/1 {
                }
                ipv6 {
                    address 2001:db8:2000::1 {
                        prefix-length 64
                    }
                }
            }
            interface "loopback" {
                loopback true
                ipv4 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    primary {
                        address 192.0.2.41
                        prefix-length 32
                    }
                }
            }
        }
```

In this setup, the configuration on CE-1 is as follows.

```
# on CE-1:
configure {
    policy-options {
        prefix-list "2001:db8::/32" {
            prefix 2001:db8::/32 type longer {
            }
        }
        prefix-list "CE-1-192.0.2.1" {
            prefix 2001:db8:3000::/64 type exact {
            }
        }
        policy-statement "CE-1-sys-to-eBGP" {
            entry 10 {
                from {
                    prefix-list ["CE-1-192.0.2.1"]
                }
                action {
                    action-type accept
                    origin igp
                }
            }
            default-action {
                action-type reject
            }
        }
        policy-statement "import-2001:db8::/32" {
            entry 10 {
                from {
                    prefix-list ["2001:db8::/32"]
                }
                action {
                    action-type accept
                    origin igp
                }
            }
            default-action {
                action-type reject
            }
        }
    }
    router "Base" {
        autonomous-system 64501
        bgp {
            router-id 192.0.2.1
            group "eBGP_to_64502" {
                description "eBGP_to_PE-3_AS64502"
                type external
                peer-as 64502
                local-address 2001:db8:1000::1
                family {
                    ipv6 true
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            neighbor "2001:db8:1000::2" {
                group "eBGP_to_64502"
                import {
                    policy ["import-2001:db8::/32"]
                }
                export {
                    policy ["CE-1-sys-to-eBGP"]
                }
            }
        }
        static-routes {
            route 2001:db8:3000::/64 route-type unicast {
                blackhole {
                    admin-state enable
                }
            }
        }
    }
    service
        ies "IES 1" {
            admin-state enable
            service-id 1
            customer "1"
            interface "int-CE-1-PE-2" {
                description "SAP_toward_VPRN_Service"
                sap 1/1/1 {
                }
                ipv6 {
                    address 2001:db8:1000::1 {
                        prefix-length 64
                    }
                }
            }
        }
```

The configuration on CE-5 is similar.

The following command on PE-2 shows that the Epipe VLL is established with the SAP facing CE-1 and spoke SDP facing VPRN 1 on PE-3.

```
[/]
A:admin@PE-2# show service id 1 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 1                   Vpn Id            : 0
Service Type      : Epipe
MACSec enabled    : no
Name              : Epipe 1
Description       : (Not Specified)
Customer Id       : 1                   Creation Origin   : manual
Last Status Change: 01/25/2022 10:53:53
Last Mgmt Change  : 01/25/2022 10:53:39
Test Service      : No
Admin State       : Up                  Oper State        : Up
MTU               : 9190
Vc Switching      : False
SAP Count         : 1                   SDP Bind Count    : 1
Per Svc Hashing   : Disabled
Vxlan Src Tep Ip  : N/A
Force QTag Fwd    : Disabled
Lcl Switch Svc St : sap
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Oper Group        : <none>

------------------------------------------------------------------------------
Service Access & Destination Points
------------------------------------------------------------------------------
Identifier                               Type         AdmMTU  OprMTU  Adm  Opr
------------------------------------------------------------------------------
sap:1/1/2                                null         9212    9212    Up   Up
sdp:231:1 S(192.0.2.3)                   Spok         0       9190    Up   Up
==============================================================================
```

The same command can be launched on PE-3 to verify that the VPRN service is up and that the spoke SDP is up (admin state up/oper state up).

```
[/]
A:admin@PE-3# show service id 1 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 1                    Vpn Id            : 0
Service Type      : VPRN
MACSec enabled    : no
Name              : VPRN 1
Description       : (Not Specified)
Customer Id       : 1                    Creation Origin   : manual
Last Status Change: 01/25/2022 10:53:48
Last Mgmt Change  : 01/25/2022 10:53:48
Admin State       : Up                   Oper State        : Up

Router Oper State : Up
Route Dist.       : 64502:1              VPRN Type         : regular
Oper Route Dist   : 64502:1
Oper RD Type      : configured
AS Number         : 64502                Router Id         : 192.0.2.31
ECMP              : Enabled              ECMP Max Routes   : 1
Max IPv4 Routes   : No Limit
Local Rt Domain-Id: None                 D-Path Lng Ignore : Disabled

Auto Bind Tunnel
Allow Flex-Alg-Fb : Disabled
Resolution        : filter
Filter Protocol   : ldp
Weighted ECMP     : Disabled             ECMP Max Routes   : 1
Strict Tnl Tag    : Disabled

Max IPv6 Routes   : No Limit
Ignore NH Metric  : Disabled
Hash Label        : Disabled
Entropy Label     : Disabled
Vrf Target        : target:64502:1
Vrf Import        : None
Vrf Export        : None
MVPN Vrf Target   : None
MVPN Vrf Import   : None
MVPN Vrf Export   : None
Car. Sup C-VPN    : Disabled
Label mode        : vrf
BGP VPN Backup    : Disabled
BGP Export Inactv : Disabled
LOG all events    : Disabled

SAP Count         : 0                    SDP Bind Count    : 1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
VSD Domain        : <none>

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                              Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sdp:321:1 S(192.0.2.2)                  TLDP         0       9190    Up   Up
===============================================================================
```

The following command shows that the IPv6 interface is established and its IPv6 address is preferred (2001:db8:1000::2/64). The IPv6 link local address (fe80::17:ffff:fe00:0/64) has been dynamically assigned and is in the preferred state.

```
[/]
A:admin@PE-3# show service id 1 interface

===============================================================================
Interface Table
===============================================================================
Interface-Name                   Adm       Opr(v4/v6)  Type    Port/SapId
   IP-Address                                                   PfxState
-------------------------------------------------------------------------------
int-PE-3-PE-2                    Up        Down/Up     VPRN    spoke-321:1
   2001:db8:1000::2/64                                        PREFERRED
   fe80::fe:1856:5892:3d76/64                                 PREFERRED
loopback                         Up        Up/Down     VPRN    loopback
   192.0.2.31/32                                              n/a
-------------------------------------------------------------------------------
Interfaces : 2
===============================================================================
```

With the **show service id 1 all** command, an extensive list of parameters is displayed, including IPv6-related fields that can be checked if configured: DHCP6-relay, DHCP6-server, and so on. It is possible to use filters to reduce the output.

After verification of the services (Epipe, VPRN), the MP-eBGP peering connectivity (through IPv6 interfaces) on the VPRN between PE-3 and CE-1 can be verified as follows:

```
[/]
A:admin@PE-3# show router 1 bgp neighbor

===============================================================================
BGP Neighbor
===============================================================================
-------------------------------------------------------------------------------
Peer                 : 2001:db8:1000::1
Description          : (Not Specified)
Group                : Spoke-CE-1-PE-3
-------------------------------------------------------------------------------
Peer AS              : 64501            Peer Port            : 179
Peer Address         : 2001:db8:1000::1
Local AS             : 64502            Local Port           : 50788
Local Address        : 2001:db8:1000::2
Peer Type            : External         Dynamic Peer         : No
State                : Established      Last State           : Active
Last Event           : recvOpen
Last Error           : Unrecognized Error
Local Family         : IPv6
Remote Family        : IPv6
---snip---
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Local Capability     : RtRefresh MPBGP 4byte ASN
Remote Capability    : RtRefresh MPBGP 4byte ASN
Local AddPath Capabi*: Disabled
Remote AddPath Capab*: Send - None
                     : Receive - None
Import Policy        : import-2001:db8::/32
                     : Default Reject
Export Policy        : PE-3-BGP-CE-1
                     : Default Reject
---snip---

-------------------------------------------------------------------------------
Ingress prefix counters per family.
IPv4 received        : 0               IPv6 received        : 1
IPv4 active          : 0               IPv6 active          : 1
IPv4 suppressed      : 0               IPv6 suppressed      : 0
IPv4 rejected        : 0               IPv6 rejected        : 0
VPN-IPv4 received    : 0               VPN-IPv6 received    : 0
VPN-IPv4 active      : 0               VPN-IPv6 active      : 0
VPN-IPv4 suppressed  : 0               VPN-IPv6 suppressed  : 0
VPN-IPv4 rejected    : 0               VPN-IPv6 rejected    : 0
---snip---
```

Not only is the MP-eBGP session on the VPRN established, but the MP-BGP capabilities are also supported (locally and remotely). PE-3 and its BGP peer CE-1 have advertised and received an IPv6 prefix.

The same command can be launched on PE-4. The status of the VPRN service on PE-4 and its interface to CE-5 can be verified as follows:

```
[/]
A:admin@PE-4# show service id 1 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 1                   Vpn Id            : 0
Service Type      : VPRN
MACSec enabled    : no
Name              : VPRN 1
Description       : (Not Specified)
Customer Id       : 1                   Creation Origin   : manual
Last Status Change: 01/25/2022 10:53:56
Last Mgmt Change  : 01/25/2022 10:53:56
Admin State       : Up                  Oper State        : Up

Router Oper State : Up
Route Dist.       : 64502:1             VPRN Type         : regular
Oper Route Dist   : 64502:1
Oper RD Type      : configured
AS Number         : 64502               Router Id         : 192.0.2.41
ECMP              : Enabled             ECMP Max Routes   : 1
Max IPv4 Routes   : No Limit
Local Rt Domain-Id: None                D-Path Lng Ignore : Disabled

Auto Bind Tunnel
Allow Flex-Alg-Fb : Disabled
Resolution        : filter
Filter Protocol   : ldp
Weighted ECMP     : Disabled            ECMP Max Routes   : 1
Strict Tnl Tag    : Disabled

Max IPv6 Routes   : No Limit
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Ignore NH Metric  : Disabled
Hash Label        : Disabled
Entropy Label     : Disabled
Vrf Target        : target:64502:1
Vrf Import        : None
Vrf Export        : None
MVPN Vrf Target   : None
MVPN Vrf Import   : None
MVPN Vrf Export   : None
Car. Sup C-VPN    : Disabled
Label mode        : vrf
BGP VPN Backup    : Disabled
BGP Export Inactv : Disabled
LOG all events    : Disabled

SAP Count         : 1                   SDP Bind Count    : 0
VSD Domain        : <none>

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                               Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/1                                null         9212    9212    Up   Up
===============================================================================
```

The VPRN service is up and the SAP is up.

The following command shows that the IPv6 interface is established and its IPv6 address is in the preferred state.

```
[/]
A:admin@PE-4# show service id 1 interface

===============================================================================
Interface Table
===============================================================================
Interface-Name                   Adm       Opr(v4/v6)  Type    Port/SapId
   IP-Address                                                   PfxState
-------------------------------------------------------------------------------
int-PE-4-CE-5                    Up        Down/Up     VPRN    1/1/1
   2001:db8:2000::1/64                                        PREFERRED
   fe80::13:6a5e:9959:90aa/64                                 PREFERRED
loopback                         Up        Up/Down     VPRN    loopback
   192.0.2.41/32                                              n/a
-------------------------------------------------------------------------------
Interfaces : 2
===============================================================================
```

MP-iBGP, providing 6VPE is configured and built between PE-3 and PE-4 across the MPLS network. IPv6 prefixes are received on PE-3 from CE-1 (2001:db8:3000::/64) and on PE-4 from CE-5 (2001:db8:4000::/64) across the MPLS network using MP-iBGP (6VPE).

CE-1 advertises IPv6 prefix 2001:db8:3000::/64 and CE-5 advertises IPv6 prefix 2001:db8:4000::/64.

The following command on PE-3 shows whether VPN-IPv6 routes were received from and advertised to its iBGP peer PE-4:

```
[/]
A:admin@PE-3# show router bgp summary
===============================================================================
 BGP Router ID:192.0.2.3        AS:64502       Local AS:64502
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
BGP Admin State         : Up          BGP Oper State              : Up
Total Peer Groups       : 1           Total Peers                 : 1
Total VPN Peer Groups   : 1           Total VPN Peers             : 1
Current Internal Groups : 1           Max Internal Groups         : 1
Total BGP Paths         : 23          Total Path Memory           : 8168
---snip---

Total VPN-IPv4 Rem. Rts : 0           Total VPN-IPv4 Rem. Act. Rts: 0
Total VPN-IPv6 Rem. Rts : 2           Total VPN-IPv6 Rem. Act. Rts: 2
Total VPN-IPv4 Bkup Rts : 0           Total VPN-IPv6 Bkup Rts     : 0
Total VPN Local Rts     : 5           Total VPN Supp. Rts         : 0
Total VPN Hist. Rts     : 0           Total VPN Decay Rts         : 0
---snip---

===============================================================================
BGP Summary
===============================================================================
Legend : D - Dynamic Neighbor
===============================================================================
Neighbor
Description
                   AS PktRcvd InQ  Up/Down   State|Rcv/Act/Sent (Addr Family)
                      PktSent OutQ
-------------------------------------------------------------------------------
192.0.2.4
PE-4
                64502      17    0 00h05m58s 2/2/2 (VpnIPv6)
                           17    0
-------------------------------------------------------------------------------
```

PE-3 has received and learned a valid and best IPv6 route for prefix 2001:db8:3000::/64 with a BGP next hop of 2001:db8:1000:: (CE-1), as follows:

```
[/]
A:admin@PE-3# show router 1 bgp routes ipv6
===============================================================================
 BGP Router ID:192.0.2.31       AS:64502       Local AS:64502
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP IPv6 Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop (Router)                                 Path-Id     IGP Cost
      As-Path                                                      Label
-------------------------------------------------------------------------------
u*>i  2001:db8:3000::/64                               None        None
      2001:db8:1000::1                                 None        0
      64501                                                        -
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

The following output shows the 2001:db8:4000::/64 prefix as BGP IPv6 route advertised by PE-3 to eBGP peer CE-1.

```
[/]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
A:admin@PE-3# show router 1 bgp neighbor 2001:db8:1000::1 advertised-routes ipv6
===============================================================================
 BGP Router ID:192.0.2.31       AS:64502       Local AS:64502
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP IPv6 Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop (Router)                                 Path-Id     IGP Cost
      As-Path                                                      Label
-------------------------------------------------------------------------------
i     2001:db8:4000::/64                               n/a         None
      2001:db8:1000::2                                 None        n/a
      64502 64502                                                  -
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

The IPv6 route 2001:db8:4000::/64 originates from CE-5 and was advertised from CE-5 to its eBGP peer PE-4, then from PE-4 as VPN-IPv6 route to its iBGP peer PE-3 with next-hop PE-4, as follows:

```
[/]
A:admin@PE-3# show router bgp routes vpn-ipv6
===============================================================================
 BGP Router ID:192.0.2.3        AS:64502       Local AS:64502
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv6 Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop (Router)                                 Path-Id     IGP Cost
      As-Path                                                      Label
-------------------------------------------------------------------------------
u*>i  64502:1:2001:db8:2000::/64                       100         None
      ::ffff:192.0.2.4                                 None        10
      No As-Path                                                   524283
u*>i  64502:1:2001:db8:4000::/64                       100         None
      ::ffff:192.0.2.4                                 None        10
      64501                                                        524283
-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

PE-3 is advertising the VPN-IPv6 route of 2001:db8:3000::/64 to its MP-iBGP peer PE-4, as follows. The IPv6 prefix 2001:db8:3000::/64 was learned from CE-1 in an MP-eBGP session:

```
[/]
A:admin@PE-3# show router bgp neighbor 192.0.2.4 advertised-routes vpn-ipv6
===============================================================================
 BGP Router ID:192.0.2.3        AS:64502       Local AS:64502
===============================================================================
 Legend -
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv6 Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop (Router)                                 Path-Id     IGP Cost
      As-Path                                                      Label
-------------------------------------------------------------------------------
i     64502:1:2001:db8:1000::/64                       100         None
      ::ffff:192.0.2.3                                 None        n/a
      No As-Path                                                   524283
i     64502:1:2001:db8:3000::/64                       100         None
      ::ffff:192.0.2.3                                 None        n/a
      64501                                                        524283
-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

The list of VPN-IPv6 routes on PE-4 includes the VPN-IPv6 route that was learned from PE-3: 2001:db8:3000::/64, as follows:

```
[/]
A:admin@PE-4# show router bgp routes vpn-ipv6
===============================================================================
 BGP Router ID:192.0.2.4        AS:64502       Local AS:64502
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv6 Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop (Router)                                 Path-Id     IGP Cost
      As-Path                                                      Label
-------------------------------------------------------------------------------
u*>i  64502:1:2001:db8:1000::/64                       100         None
      ::ffff:192.0.2.3                                 None        10
      No As-Path                                                   524283
u*>i  64502:1:2001:db8:3000::/64                       100         None
      ::ffff:192.0.2.3                                 None        10
      64501                                                        524283
-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

The following output shows the advertised VPN-IPv6 route of 2001:db8:4000::/64 from PE-4 to PE-3.

```
[/]
A:admin@PE-4# show router bgp neighbor 192.0.2.3 advertised-routes vpn-ipv6
===============================================================================
 BGP Router ID:192.0.2.4        AS:64502       Local AS:64502
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
BGP VPN-IPv6 Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop (Router)                                 Path-Id     IGP Cost
      As-Path                                                      Label
-------------------------------------------------------------------------------
i     64502:1:2001:db8:2000::/64                       100         None
      ::ffff:192.0.2.4                                 None        n/a
      No As-Path                                                   524283
i     64502:1:2001:db8:4000::/64                       100         None
      ::ffff:192.0.2.4                                 None        n/a
      64501                                                        524283
-------------------------------------------------------------------------------
Routes : 2
===============================================================================
```

The following output from PE-4 shows the IPv6 prefix 2001:db8:4000::/64 learned from CE-5.

```
[/]
A:admin@PE-4# show router 1 bgp routes ipv6
===============================================================================
 BGP Router ID:192.0.2.41       AS:64502       Local AS:64502
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP IPv6 Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop (Router)                                 Path-Id     IGP Cost
      As-Path                                                      Label
-------------------------------------------------------------------------------
u*>i  2001:db8:4000::/64                               None        None
      2001:db8:2000::2                                 None        0
      64501                                                        -
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

The following command on PE-4 confirms that IPv6 prefix 2001:db8:3000::/64 is advertised to CE-5.

```
[/]
A:admin@PE-4# show router 1 bgp neighbor 2001:db8:2000::2 advertised-routes ipv6
===============================================================================
 BGP Router ID:192.0.2.41       AS:64502       Local AS:64502
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP IPv6 Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop (Router)                                 Path-Id     IGP Cost
      As-Path                                                      Label
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
i     2001:db8:3000::/64                                    n/a         None
      2001:db8:2000::1                                      None        n/a
      64502 64502                                                       -
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

The final verification of CE-1 and CE-5 shows that IPv6 routes for AS 64501 have been received and are valid across the VPRN service, as follows:

```
[/]
A:admin@CE-1# show router bgp routes ipv6
===============================================================================
 BGP Router ID:192.0.2.1        AS:64501       Local AS:64501
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP IPv6 Routes
===============================================================================
Flag  Network                                               LocalPref   MED
      Nexthop (Router)                                      Path-Id     IGP Cost
      As-Path                                                           Label
-------------------------------------------------------------------------------
u*>i  2001:db8:4000::/64                                    None        None
      2001:db8:1000::2                                      None        0
      64502 64502                                                       -
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```

```
[/]
A:admin@CE-5# show router bgp routes ipv6
===============================================================================
 BGP Router ID:192.0.2.5        AS:64501       Local AS:64501
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP IPv6 Routes
===============================================================================
Flag  Network                                               LocalPref   MED
      Nexthop (Router)                                      Path-Id     IGP Cost
      As-Path                                                           Label
-------------------------------------------------------------------------------
u*>i  2001:db8:3000::/64                                    None        None
      2001:db8:2000::1                                      None        0
      64502 64502                                                       -
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Conclusion

Spoke termination for IPv6-6VPE extends the use of spoke terminated interfaces from an Epipe VLL into a VPRN service using IPv6 interfaces on the access. Supporting the requirement of IPv6 interfaces, routing of IPv6 prefixes and the use of 6VPE for IPv6 tunneling over an IPv4 network allows SR OS to provide capabilities supporting the growth of IPv6 architectures.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Traffic Leaking from VPRN to GRT

This chapter provides information about Traffic Leaking from VPRN to GRT.

Topics in this chapter include:

Applicability

The information and configuration in this chapter were originally based on SR OS Release 14.0 R4. The MD-CLI in the current edition corresponds to SR OS Release 22.2.R2.

Overview

RFC 4364 *BGP/MPLS IP Virtual Private Networks (VPNs)* describes a method of distributing routing information using BGP and MPLS forwarding data to provide a Layer 3 VPN service to end users. Each Virtual Private Routed Network (VPRN) consists of a set of customer sites connected to one or more PE routers. Each associated PE router maintains a separate IP forwarding table for each VPRN. Additionally, the PE routers exchange the routing information configured or learned from all customer sites via Multi-Protocol Border Gateway Protocol (MP-BGP) peering. Each route exchanged via the MP-BGP protocol includes a route distinguisher (RD), which identifies the VPRN association and resolves any IP address overlap.

It has always been possible to exchange traffic from one VPRN to another, using scenarios such as "extranet", "hub and spoke" and so on, using the vrf-import and vrf-export policies for BGP VPN-IPv4 route distribution.

Traffic leaking to the Global Route Table (GRT) allows service providers to offer VPRN and Internet services over a single virtual routing and forwarding VRF interface. Packets entering a VRF interface can have route processing results derived from the VRF or the GRT. The leaking and preferred lookup settings are configured on a per-VPRN basis.

To allow data flowing from a VPRN to the base router, routing information from the base router must be made available for lookup by the VPRN. The GRT lookup can be general (for example, any lookup miss in the Virtual Routing and Forwarding (VRF) table can be resolved in the GRT), or specific (for example, specific routes should only be searched for in the GRT and ignored by the VPRN).

To enable the GRT lookup from the VPRN, the **grt-leaking>grt-lookup true** command is used. This only provides part of the solution, because packets can now be forwarded from the VPRN to the GRT, but not in the opposite direction. The GRT needs to learn specific destination prefixes from the VPRN and this is achieved by route leaking from the VPRN to the GRT, using policies (**grt-leaking>export-grt** command). The maximum number of routes leaked from a VPRN to the GRT is five by default, but this maximum can be modified or even removed. Prefixes should be globally unique within the service provider network and if these are propagated outside the provider's network, they must be from the public IP space and globally unique.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 319: VPRN to GRT leak process

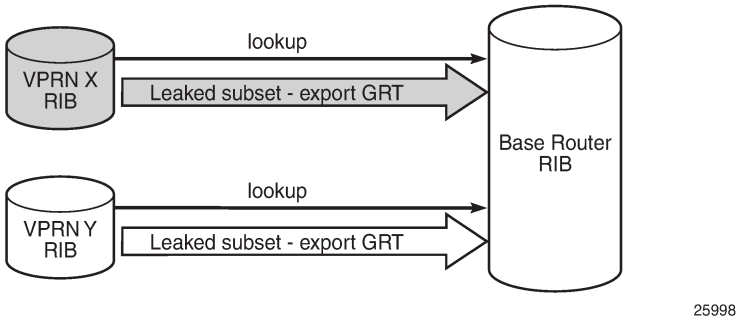

The method described in this chapter allows the network administrator to leak specific or all routes that are inside a VPRN to the GRT. Route leaking from VPRN to GRT is protocol-independent and can be applied for BGP, OSPF(v3), IS-IS, static, local routes, and so on. For BGP routes, there is an improved route leaking mechanism that allows leaking routes preserving all BGP attributes; see chapter the "BGP Route Leaking" in the Unicast Routing Protocols volume in the *7450 ESS, 7750 SR, and 7950 XRS MD-CLI Advanced Configuration Guide - Part I*.

Configuration

Figure 320: Example topology with IPv4 addresses shows the example topology used in this chapter, including the IPv4 addresses. The interfaces also have IPv6 addresses, which will be shown in Figure 322: Example topology with IPv6 addresses.

Figure 320: Example topology with IPv4 addresses

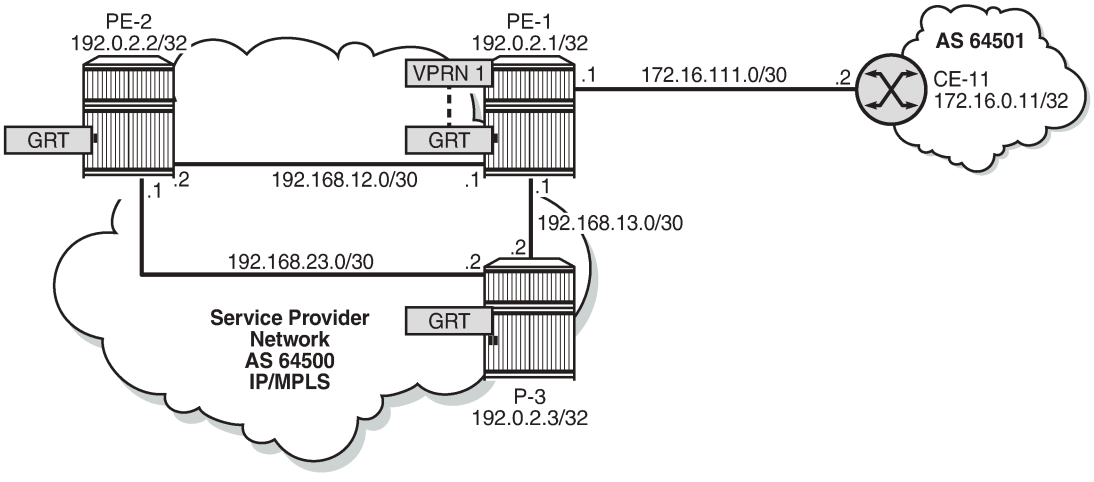

Initial configuration

The nodes in the example topology have the following initial configuration:

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- Cards, MDAs, ports
- Router interfaces
- IGP (IS-IS or OSPF) between the PEs
- LDP between the PEs
- VPRN "VPRN 1" on PE-1
- BGP (IBGP between the PEs; EBGP between PE-1 and CE-11)
  - On PE-1, BGP is configured in the base router and in VPRN 1.
- Loopback addresses on CE-11, such as 192.168.110.2/32.
- Export policies on CE-11 to export routes from direct with certain prefixes.

Protocol-independent IPv4 route leaking from VPRN to GRT

Figure 321: IPv4 VPRN to GRT route leaking for IS-IS shows the topology with the IP addresses for this example. Route leaking from VPRN to GRT is protocol independent and in this example, VPRN "VPRN 1" on PE-1 will leak local routes, static routes, and imported BGP routes to the GRT. IS-IS or OSPF routes can also be leaked, but that is not shown here.

Figure 321: IPv4 VPRN to GRT route leaking for IS-IS

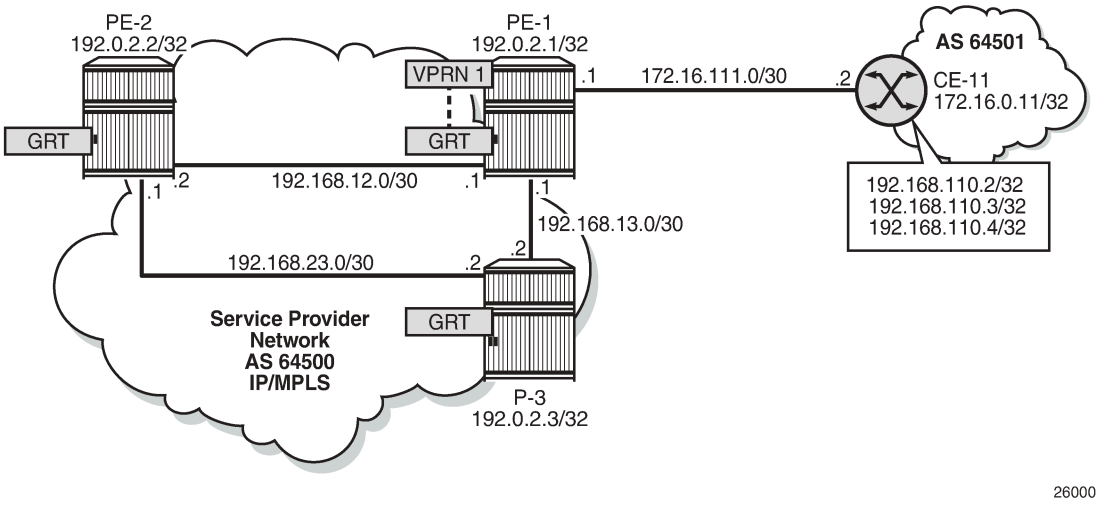

GRT-leak is by default disabled. The routing table for VPRN 1 on PE-1 contains local routes, static routes, and BGP routes that are learned from CE-11, as follows:

```
[/]
A:admin@PE-1# show router 1 route-table

===============================================================================
Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
172.16.1.1/32                                 Local   Local     00h01m14s  0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
       system                                                       0
172.16.111.0/30                          Local   Local     00h01m14s  0
       int-PE-1-CE-11                                               0
192.168.110.2/32                         Remote  BGP       00h00m12s  170
       172.16.111.2                                                 0
192.168.110.3/32                         Remote  BGP       00h00m12s  170
       172.16.111.2                                                 0
192.168.110.4/32                         Remote  BGP       00h00m12s  170
       172.16.111.2                                                 0
192.168.120.0/24                         Remote  Static    00h01m14s  5
       172.16.111.2                                                 1
-------------------------------------------------------------------------------
No. of Routes: 6
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

By default, the GRT is not learning the VPRN routes, as follows:

```
[/]
A:admin@PE-1# show router route-table

===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                       Type    Proto     Age        Pref
      Next Hop[Interface Name]                                Metric
-------------------------------------------------------------------------------
192.0.2.1/32                             Local   Local     00h01m14s  0
       system                                                       0
192.0.2.2/32                             Remote  ISIS      00h01m03s  15
       192.168.12.2                                                 10
192.0.2.3/32                             Remote  ISIS      00h00m51s  15
       192.168.13.2                                                 10
192.168.12.0/30                          Local   Local     00h01m14s  0
       int-PE-1-PE-2                                                0
192.168.13.0/30                          Local   Local     00h01m14s  0
       int-PE-1-P-3                                                 0
192.168.23.0/30                          Remote  ISIS      00h01m03s  15
       192.168.12.2                                                 20
-------------------------------------------------------------------------------
No. of Routes: 6
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

To enable VPRN to GRT leaking, the following route policy is configured on PE-1 and applied in VPRN 1:

```
# on PE-1:
configure {
    policy-options {
        policy-statement "LeakVPRNtoGRT_pref8" {
            entry 10 {
                action {
                    action-type accept
                    preference 8
                }
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    }
    service {
        vprn "VPRN 1" {
            grt-leaking {
                grt-lookup true
                export-grt {
                    policy-name ["LeakVPRNtoGRT_pref8"]
                }
            }
```

This policy allows leaking all routes from a VPRN to the base router, without any match criteria. However, when routes are leaked from VPRNs to the GRT, they need to be unique and only routes that need to be known in the GRT should be leaked. By default, the preference for a leaked route is 180. The preference can be manually configured to a lower value, such as 8, to avoid network inconsistencies between the IGP and the RT on the router where the routes are leaked.

When **grt-leaking>grt-lookup true** is configured, any lookup miss in the VRF table will be resolved in the GRT, if available. This only works from VPRN to GRT and does not require route leaking. However, the base router needs to be able to route packets back to the VPRN and it cannot perform a lookup in the routing table of the VPRN. Therefore, route leaking from VPRN to GRT is required, and **grt-leaking>export-grt** is configured. Prefixes in the VPRN must be leaked to the GRT through a policy. Prefixes leaked from any VPRN should never conflict with prefixes leaked from any other VPRN or existing prefixes in the GRT.

This configuration is protocol-independent. Route leaking from VPRN to GRT is applicable for all kinds of learned routes, such as static routes, local routes, IS-IS, OSPF, BGP, and so on.

After routes are leaked from the VPRN to the GRT, the routing table of the base router includes the leaked routes, with protocol "VPN Leak". For PE-1, the routing table contains the following routes:

```
[/]
A:admin@PE-1# show router route-table

===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
172.16.1.1/32                                 Remote  VPN Leak  00h00m20s  8
       system                                                        0
172.16.111.0/30                               Remote  VPN Leak  00h00m20s  8
       int-PE-1-CE-11                                                0
192.0.2.1/32                                  Local   Local     00h03m06s  0
       system                                                        0
192.0.2.2/32                                  Remote  ISIS      00h02m49s  15
       192.168.12.2                                                  10
192.0.2.3/32                                  Remote  ISIS      00h02m41s  15
       192.168.13.2                                                  10
192.168.12.0/30                               Local   Local     00h03m06s  0
       int-PE-1-PE-2                                                 0
192.168.13.0/30                               Local   Local     00h03m06s  0
       int-PE-1-P-3                                                  0
192.168.23.0/30                               Remote  ISIS      00h02m49s  15
       192.168.12.2                                                  20
192.168.110.3/32                              Remote  VPN Leak  00h00m20s  8
       172.16.111.2                                                  0
192.168.110.4/32                              Remote  VPN Leak  00h00m20s  8
       172.16.111.2                                                  0
192.168.120.0/24                              Remote  VPN Leak  00h00m20s  8
       172.16.111.2                                                  0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
No. of Routes: 11
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

Regardless the preference of the original routes in VPRN 1, all the leaked routes in the GRT have preference 8, as configured. By default, a maximum of five routes are leaked. This export limit can be overruled, as follows:

```
# on PE-1:
configure {
    service {
        vprn "VPRN 1" {
            grt-leaking {
                export-limit 10
                grt-lookup true
                export-grt {
                    policy-name ["LeakVPRNtoGRT_pref8"]
                }
            }
```

The following command shows only the routes leaked from any VPRN to GRT on PE-1:

```
[/]
A:admin@PE-1# show router route-table protocol vpn-leak all

===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                         Active     Metric
-------------------------------------------------------------------------------
172.16.1.1/32                                 Remote  VPN Leak  00h00m05s  8
       system                                          Y            0
172.16.111.0/30                               Remote  VPN Leak  00h00m05s  8
       int-PE-1-CE-11                                  Y            0
192.168.110.2/32                              Remote  VPN Leak  00h00m05s  8
       172.16.111.2                                    Y            0
192.168.110.3/32                              Remote  VPN Leak  00h00m05s  8
       172.16.111.2                                    Y            0
192.168.110.4/32                              Remote  VPN Leak  00h00m05s  8
       172.16.111.2                                    Y            0
192.168.120.0/24                              Remote  VPN Leak  00h00m05s  8
       172.16.111.2                                    Y            0
-------------------------------------------------------------------------------
No. of Routes: 6
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
       E = Inactive best-external BGP route
===============================================================================
```

Different types of routes are leaked to the GRT with protocol type "VPN Leak" and all of them get the same preference, configured or default. The detailed output for any leaked route in the preceding list for PE-1 shows protocol VPN_LEAK and preference 8, as follows:

```
[/]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
A:admin@PE-1# show router route-table protocol vpn-leak 192.168.110.2/32 extensive

===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix             : 192.168.110.2/32
  Protocol              : VPN_LEAK
  Age                   : 00h00m24s
  Preference            : 8
  Next-Hop              : 172.16.111.2
    Interface           : int-PE-1-CE-11 (VPRN 1)
    QoS                 : Priority=n/c, FC=n/c
    Source-Class        : 0
    Dest-Class          : 0
    Metric              : 0
    ECMP-Weight         : N/A
-------------------------------------------------------------------------------
No. of Destinations: 1
===============================================================================
```

Export IPv4 VPN-leak routes to routing protocols

Until now, the VPN-leak routes are leaked locally to the GRT, but they are not advertised in IS-IS, OSPF, or BGP. Router P-3 has not learned any of the leaked routes, as follows:

```
[/]
A:admin@P-3# show router route-table

===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
192.0.2.1/32                                    Remote  ISIS      00h03m44s  15
       192.168.13.1                                                 10
192.0.2.2/32                                    Remote  ISIS      00h03m44s  15
       192.168.23.1                                                 10
192.0.2.3/32                                    Local   Local     00h03m51s  0
       system                                                       0
192.168.12.0/30                                 Remote  ISIS      00h03m44s  15
       192.168.13.1                                                 20
192.168.13.0/30                                 Local   Local     00h03m51s  0
       int-P-3-PE-1                                                 0
192.168.23.0/30                                 Local   Local     00h03m51s  0
       int-P-3-PE-2                                                 0
-------------------------------------------------------------------------------
No. of Routes: 6
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

To reduce the number of routes to be exported on PE-1, a match criterion is added for the routes to be leaked, as follows:

```
# on PE-1:
configure {
    policy-options {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        prefix-list "192.168.110.0" {
            prefix 192.168.110.0/24 type longer {
            }
        }
        policy-statement "LeakVPRNtoGRT_pref8_110" {
            entry 10 {
                from {
                    prefix-list ["192.168.110.0"]
                }
                action {
                    action-type accept
                    preference 8
                }
            }
        }
    }
    service {
        vprn "VPRN 1" {
            grt-leaking {
                export-limit 10
                grt-lookup true
                export-grt {
                    delete policy-name ["LeakVPRNtoGRT_pref8"]
                    policy-name ["LeakVPRNtoGRT_pref8_110"]
                }
```

VPN-leak routes can be exported to any routing protocol. Prefix lists can be used to filter routes, but that is not configured in this example. The following export policy is configured on PE-1 to export the VPN-leak routes:

```
# on PE-1:
configure {
    policy-options
        policy-statement "export-vpn-leak" {
            entry 10 {
                from {
                    protocol {
                        name [vpn-leak]
                    }
                }
                action {
                    action-type accept
                }
            }
```

The same export policy will be used for export to IS-IS, OSPF, and BGP.

Export IPv4 VPN-leak routes to IS-IS

The export policy is applied in the IS-IS context on PE-1, as follows:

```
# on PE-1:
configure {
    router "Base" {
        isis 0 {
            export-policy ["export-vpn-leak"]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The leaked routes are now advertised via IS-IS and appear as IS-IS routes with default preference for IS-IS routes on PE-2 and P-3. The route table on P-3 looks as follows:

```
[/]
A:admin@P-3# show router route-table

===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
192.0.2.1/32                                    Remote  ISIS      00h04m39s  15
       192.168.13.1                                               10
192.0.2.2/32                                    Remote  ISIS      00h04m39s  15
       192.168.23.1                                               10
192.0.2.3/32                                    Local   Local     00h04m46s  0
       system                                                     0
192.168.12.0/30                                 Remote  ISIS      00h04m39s  15
       192.168.13.1                                               20
192.168.13.0/30                                 Local   Local     00h04m46s  0
       int-P-3-PE-1                                               0
192.168.23.0/30                                 Local   Local     00h04m46s  0
       int-P-3-PE-2                                               0
192.168.110.2/32                                Remote  ISIS      00h00m21s  15
       192.168.13.1                                               10
192.168.110.3/32                                Remote  ISIS      00h00m21s  15
       192.168.13.1                                               10
192.168.110.4/32                                Remote  ISIS      00h00m21s  15
       192.168.13.1                                               10
-------------------------------------------------------------------------------
No. of Routes: 9
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

The export policy is removed from the IS-IS context on PE-1, as follows:

```
# on PE-1:
configure {
    router "Base" {
        isis 0 {
            delete export-policy ["export-vpn-leak"]
```

Export IPv4 VPN-leak routes to OSPF

When OSPF is used instead of IS-IS, the behavior is similar. The export policy is applied in the OSPF context on PE-1, as follows:

```
# on PE-1:
configure {
    router "Base"
        ospf 0 {
            export-policy ["export-vpn-leak"]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

To export routes into OSPF using a policy, the router must be configured as ASBR, as follows:

```
# on PE-1:
configure {
    router "Base"
        ospf 0 {
            asbr {
            }
```

The routes with protocol VPN-leak on PE-1 are now exported in OSPF to PE-2 and P-3. The default preference for external OSPF routes is 150. On P-3, the routing table contains the following OSPF routes:

```
[/]
A:admin@P-3# show router route-table protocol ospf

===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
192.0.2.1/32                                    Remote  OSPF      00h00m44s  10
       192.168.13.1                                                 10
192.0.2.2/32                                    Remote  OSPF      00h00m44s  10
       192.168.23.1                                                 10
192.168.12.0/30                                 Remote  OSPF      00h00m44s  10
       192.168.13.1                                                 20
192.168.110.2/32                                Remote  OSPF      00h00m14s  150
       192.168.13.1                                                 1
192.168.110.3/32                                Remote  OSPF      00h00m14s  150
       192.168.13.1                                                 1
192.168.110.4/32                                Remote  OSPF      00h00m14s  150
       192.168.13.1                                                 1
-------------------------------------------------------------------------------
No. of Routes: 6
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

The export policy is removed from the OSPF context on PE-1 as follows:

```
# on PE-1:
configure {
    router "Base" {
        ospf 0 {
            delete export-policy ["export-vpn-leak"]
```

Export IPv4 VPN-leak routes to BGP

The export policy is applied in the general **bgp** context of PE-1, as follows:

```
# on PE-1:
configure {
    router "Base" {
        bgp {
            export {
                policy ["export-vpn-leak"]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
```

The VPN-leak routes from PE-1 will be advertised as BGP routes to BGP neighbors PE-2 and P-3, and the routing tables will contain BGP routes with preference 170. P-3 has the following BGP routes:

```
[/]
A:admin@P-3# show router route-table protocol bgp

===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                     Metric
-------------------------------------------------------------------------------
192.168.110.2/32                                Remote  BGP       00h00m16s  170
       192.168.13.1                                                 10
192.168.110.3/32                                Remote  BGP       00h00m16s  170
       192.168.13.1                                                 10
192.168.110.4/32                                Remote  BGP       00h00m16s  170
       192.168.13.1                                                 10
-------------------------------------------------------------------------------
No. of Routes: 3
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

Note:
If it is required to preserve the BGP path attributes in the leaking process, you must use the BGP Route Leaking process described in chapter *BGP Route Leaking*. However, with this protocol-independent route leaking mechanism, it is possible to leak non-BGP routes to the GRT that will be advertised as BGP routes.

The export policy is removed from the **bgp** context, as follows:

```
# on PE-1:
configure {
    router "Base" {
        bgp {
            delete export
```

Protocol-independent IPv6 route leaking from VPRN to GRT

Figure 322: Example topology with IPv6 addresses shows the topology and the IP addresses used for IPv6. CE-11 exports routes such as 2001:db8:110::2/128 to VPRN 1on PE-1. On PE-1, local routes, static routes, and BGP routes will be leaked to the GRT.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 322: Example topology with IPv6 addresses

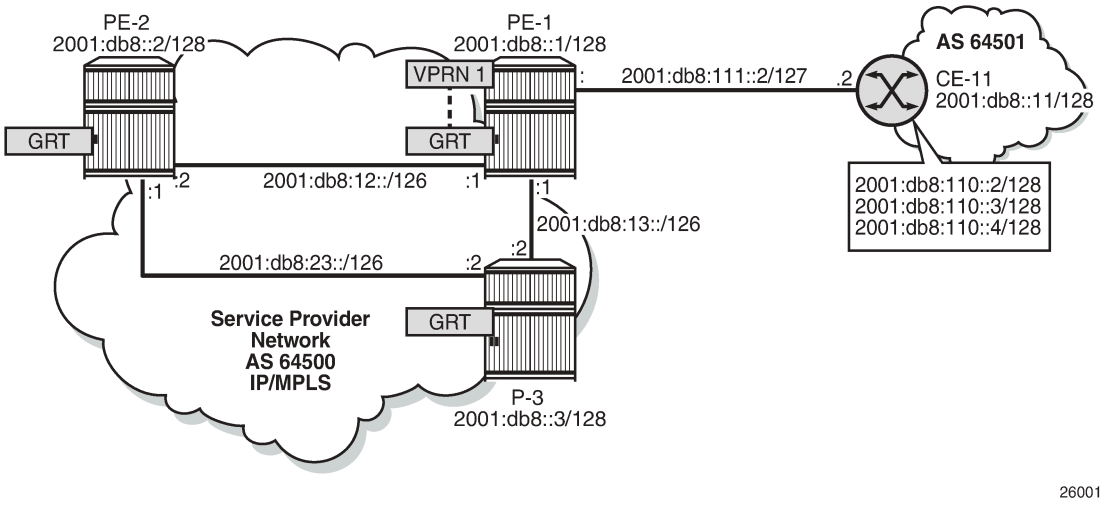

The IPv6 routing table for VPRN 1 on PE-1 includes local addresses, a static route, and three BGP routes exported by CE-11, as follows:

```
[/]
A:admin@PE-1# show router 1 route-table ipv6

===============================================================================
IPv6 Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                     Metric
-------------------------------------------------------------------------------
2001:db8::1:1/128                               Local   Local     00h08m49s  0
       system                                                        0
2001:db8:110::2/128                             Remote  BGP       00h07m49s  170
       2001:db8:111::1                                               0
2001:db8:110::3/128                             Remote  BGP       00h07m49s  170
       2001:db8:111::1                                               0
2001:db8:110::4/128                             Remote  BGP       00h07m49s  170
       2001:db8:111::1                                               0
2001:db8:111::/127                              Local   Local     00h08m48s  0
       int-PE-1-CE-11                                                0
2001:db8:120::/120                              Remote  Static    00h08m48s  5
       2001:db8:111::1                                               1
-------------------------------------------------------------------------------
No. of Routes: 6
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

By default, route leaking is disabled and the IPv6 GRT on PE-1 does not contain any of the IPv6 routes in VPRN 1, as follows:

```
[/]
A:admin@PE-1# show router route-table ipv6
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
IPv6 Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
2001:db8::1/128                               Local   Local     00h08m50s  0
       system                                                         0
2001:db8::2/128                               Remote  OSPF3     00h08m38s  10
       fe80::14:1ff:fe01:1-"int-PE-1-PE-2"                            10
2001:db8::3/128                               Remote  OSPF3     00h08m26s  10
       fe80::18:1ff:fe01:2-"int-PE-1-P-3"                             10
2001:db8:12::/126                             Local   Local     00h08m49s  0
       int-PE-1-PE-2                                                  0
2001:db8:13::/126                             Local   Local     00h08m49s  0
       int-PE-1-P-3                                                   0
2001:db8:23::/126                             Remote  OSPF3     00h08m22s  10
       fe80::18:1ff:fe01:2-"int-PE-1-P-3"                             20
-------------------------------------------------------------------------------
No. of Routes: 6
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

The VPN-leak route policy is the same as for IPv4 routes, and is applied in the **vprn** context in the same way as for IPv4 routes, as follows:

```
# on PE-1:
configure {
    policy-options
        policy-statement "LeakVPRNtoGRT_pref8" {
            entry 10 {
                action {
                    action-type accept
                    preference 8
                }
            }
        }
    }
    service {
        vprn "VPRN 1" {
            grt-leaking {
                grt-lookup true
                export-grt {
                    policy-name ["LeakVPRNtoGRT_pref8"]
                }
            }
```

On PE-1, the IPv6 routing table for VPRN 1 contains six routes, but by default, a maximum of five routes are leaked, as follows:

```
[/]
A:admin@PE-1# show router route-table ipv6 protocol vpn-leak all

===============================================================================
IPv6 Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                           Active     Metric
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
2001:db8::1:1/128                               Remote  VPN Leak  00h00m12s  8
       system                                            Y           0
2001:db8:110::2/128                             Remote  VPN Leak  00h00m12s  8
       2001:db8:111::1                                   Y           0
2001:db8:110::4/128                             Remote  VPN Leak  00h00m12s  8
       2001:db8:111::1                                   Y           0
2001:db8:111::/127                              Remote  VPN Leak  00h00m12s  8
       int-PE-1-CE-11                                    Y           0
2001:db8:120::/120                              Remote  VPN Leak  00h00m12s  8
       2001:db8:111::1                                   Y           0
-------------------------------------------------------------------------------
No. of Routes: 5
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
       E = Inactive best-external BGP route
===============================================================================
```

The export limit for IPv6 routes is removed, as follows:

```
# on PE-1:
configure {
    service {
        vprn "VPRN 1" {
            grt-leaking {
                export-v6-limit 0
                grt-lookup true
                export-grt {
                    policy-name ["LeakVPRNtoGRT_pref8"]
                }
            }
```

As a result, there is no limit to the number of leaked IPv6 routes, and all six IPv6 routes are leaked from VPRN 1 to the GRT with the configured preference 8, as follows:

```
[/]
A:admin@PE-1# show router route-table ipv6 protocol vpn-leak all

===============================================================================
IPv6 Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                           Active     Metric
-------------------------------------------------------------------------------
2001:db8::1:1/128                               Remote  VPN Leak  00h00m05s  8
       system                                            Y           0
2001:db8:110::2/128                             Remote  VPN Leak  00h00m05s  8
       2001:db8:111::1                                   Y           0
2001:db8:110::3/128                             Remote  VPN Leak  00h00m05s  8
       2001:db8:111::1                                   Y           0
2001:db8:110::4/128                             Remote  VPN Leak  00h00m05s  8
       2001:db8:111::1                                   Y           0
2001:db8:111::/127                              Remote  VPN Leak  00h00m05s  8
       int-PE-1-CE-11                                    Y           0
2001:db8:120::/120                              Remote  VPN Leak  00h00m05s  8
       2001:db8:111::1                                   Y           0
-------------------------------------------------------------------------------
No. of Routes: 6
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
       S = Sticky ECMP requested
       E = Inactive best-external BGP route
===============================================================================
```

The details for any of the routes shows that the protocol is VPN-leak and the preference is 8, as follows:

```
[/]
A:admin@PE-1# show router route-table protocol vpn-leak 2001:db8:110::2/128 extensive

===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix             : 2001:db8:110::2/128
  Protocol              : VPN_LEAK
  Age                   : 00h00m26s
  Preference            : 8
  Next-Hop              : 2001:db8:111::1
    Interface           : int-PE-1-CE-11 (VPRN 1)
    QoS                 : Priority=n/c, FC=n/c
    Source-Class        : 0
    Dest-Class          : 0
    Metric              : 0
    ECMP-Weight         : N/A
-------------------------------------------------------------------------------
No. of Destinations: 1
===============================================================================
```

Export IPv6 VPN-leak routes to routing protocols

Until now, the IPv6 VPN-leak routes are leaked locally to the GRT, but they are not advertised in IS-IS, OSPFv3, or BGP. Router P-3 has not learned any of the leaked IPv6 routes, as follows:

```
[/]
A:admin@P-3# show router route-table ipv6

===============================================================================
IPv6 Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
2001:db8::1/128                                 Remote  OSPF3     00h10m51s  10
       fe80::10:1ff:fe01:1-"int-P-3-PE-1"                              10
2001:db8::2/128                                 Remote  OSPF3     00h10m46s  10
       fe80::14:1ff:fe01:2-"int-P-3-PE-2"                              10
2001:db8::3/128                                 Local   Local     00h10m52s  0
       system                                                          0
2001:db8:12::/126                               Remote  OSPF3     00h10m51s  10
       fe80::10:1ff:fe01:1-"int-P-3-PE-1"                              20
2001:db8:13::/126                               Local   Local     00h10m51s  0
       int-P-3-PE-1                                                    0
2001:db8:23::/126                               Local   Local     00h10m51s  0
       int-P-3-PE-2                                                    0
-------------------------------------------------------------------------------
No. of Routes: 6
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

To reduce the number of VPN-leak routes, a match criterion is added to the **export-grt** policy on PE-1, as follows:

```
# on PE-1:
configure {
    policy-options {
        prefix-list "2001:db8:110::" {
            prefix 2001:db8:110::/125 type longer {
            }
        }
        policy-statement "LeakVPRNtoGRT_pref8_110" {
            entry 20 {
                from {
                    prefix-list ["2001:db8:110::"]
                }
                action {
                    action-type accept
                    preference 8
                }
            }
        }
    }
    service {
        vprn "VPRN 1" {
            grt-leaking {
                export-v6-limit 0
                grt-lookup true
                export-grt {
                    delete policy-name ["LeakVPRNtoGRT_pref8"]
                    policy-name ["LeakVPRNtoGRT_pref8_110"]
                }
            }
```

The following IPv6 routes are leaked from VPRN 1 to GRT on PE-1:

```
[/]
A:admin@PE-1# show router route-table ipv6 protocol vpn-leak

===============================================================================
IPv6 Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
2001:db8:110::2/128                           Remote  VPN Leak  00h00m21s  8
       2001:db8:111::1                                            0
2001:db8:110::3/128                           Remote  VPN Leak  00h00m21s  8
       2001:db8:111::1                                            0
2001:db8:110::4/128                           Remote  VPN Leak  00h00m21s  8
       2001:db8:111::1                                            0
-------------------------------------------------------------------------------
No. of Routes: 3
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

IPv6 VPN-leak routes can be exported to routing protocols IS-IS, OSPFv3, and BGP.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The export policy on PE-1 is the same as in all the preceding examples for IPv4, as follows:

```
# on PE-1:
configure {
    policy-options {
        policy-statement "export-vpn-leak" {
            entry 10 {
                from {
                    protocol {
                        name [vpn-leak]
                    }
                }
                action {
                    action-type accept
                }
            }
```

Export IPv6 VPN-leak routes to IS-IS

The export policy for IPv6 routes of protocol VPN-leak is applied for IS-IS, as follows:

```
# on PE-1:
configure {
    router "Base" {
        isis 0 {
            export-policy ["export-vpn-leak"]
```

The three IPv6 VPN-leak routes from PE-1 are now advertised by IS-IS to PE-2 and P-3. The routing table on P-3 contains the following IPv6 IS-IS routes:

```
[/]
A:admin@P-3# show router route-table ipv6 protocol isis

===============================================================================
IPv6 Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
2001:db8::1/128                               Remote  ISIS      00h11m54s  15
       fe80::10:1ff:fe01:1-"int-P-3-PE-1"                           10
2001:db8::2/128                               Remote  ISIS      00h11m54s  15
       fe80::14:1ff:fe01:2-"int-P-3-PE-2"                           10
2001:db8:12::/126                             Remote  ISIS      00h11m54s  15
       fe80::10:1ff:fe01:1-"int-P-3-PE-1"                           20
2001:db8:110::2/128                           Remote  ISIS      00h00m15s  15
       fe80::10:1ff:fe01:1-"int-P-3-PE-1"                           10
2001:db8:110::3/128                           Remote  ISIS      00h00m15s  15
       fe80::10:1ff:fe01:1-"int-P-3-PE-1"                           10
2001:db8:110::4/128                           Remote  ISIS      00h00m15s  15
       fe80::10:1ff:fe01:1-"int-P-3-PE-1"                           10
-------------------------------------------------------------------------------
No. of Routes: 6
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The export policy is removed for IS-IS, as follows:

```
# on PE-1:
configure {
    router "Base" {
        isis 0 {
            delete export-policy ["export-vpn-leak"]
```

Export IPv6 VPN-leak routes to OSPFv3

The export policy for IPv6 routes of protocol VPN-leak is applied for OSPFv3, as follows:

```
# on PE-1:
configure {
    router "Base" {
        ospf3 0 {
            export-policy ["export-vpn-leak"]
```

Routes can only be exported to OSPFv3 if the router is configured as ASBR, as follows:

```
# on PE-1:
configure {
    router "Base" {
        ospf3 0 {
            asbr { }
```

The IPv6 VPN-leak routes from PE-1 are now advertised by OSPFv3 to PE-2 and P-3. The preference for remote OSPFv3 routes is by default 150. The routing table on P-3 contains the following IPv6 OSPFv3 routes:

```
[/]
A:admin@P-3# show router route-table ipv6 protocol ospf3

===============================================================================
IPv6 Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
2001:db8::1/128                               Remote  OSPF3     00h00m48s  10
       fe80::10:1ff:fe01:1-"int-P-3-PE-1"                           10
2001:db8::2/128                               Remote  OSPF3     00h00m48s  10
       fe80::14:1ff:fe01:2-"int-P-3-PE-2"                           10
2001:db8:12::/126                             Remote  OSPF3     00h00m48s  10
       fe80::10:1ff:fe01:1-"int-P-3-PE-1"                           20
2001:db8:110::2/128                           Remote  OSPF3     00h00m22s  150
       fe80::10:1ff:fe01:1-"int-P-3-PE-1"                           1
2001:db8:110::3/128                           Remote  OSPF3     00h00m22s  150
       fe80::10:1ff:fe01:1-"int-P-3-PE-1"                           1
2001:db8:110::4/128                           Remote  OSPF3     00h00m22s  150
       fe80::10:1ff:fe01:1-"int-P-3-PE-1"                           1
-------------------------------------------------------------------------------
No. of Routes: 6
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The export policy is removed for OSPFv3, as follows:

```
# on PE-1:
configure {
    router "Base" {
        ospf3 0 {
            delete export-policy ["export-vpn-leak"]
```

Export IPv6 VPN-leak routes to BGP

The export policy for IPv6 routes of protocol VPN-leak is applied for BGP, as follows:

```
# on PE-1:
configure {
    router "Base" {
        bgp {
            export {
                policy ["export-vpn-leak"]
```

The three IPv6 VPN-leak routes from PE-1 are now advertised by BGP to PE-2 and P-3. The routing table on P-3 contains the following IPv6 BGP routes:

```
[/]
A:admin@P-3# show router route-table ipv6 protocol bgp

===============================================================================
IPv6 Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
2001:db8:110::2/128                           Remote  BGP       00h00m15s  170
       fe80::10:1ff:fe01:1-"int-P-3-PE-1"                           10
2001:db8:110::3/128                           Remote  BGP       00h00m15s  170
       fe80::10:1ff:fe01:1-"int-P-3-PE-1"                           10
2001:db8:110::4/128                           Remote  BGP       00h00m15s  170
       fe80::10:1ff:fe01:1-"int-P-3-PE-1"                           10
-------------------------------------------------------------------------------
No. of Routes: 3
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

The export policy is removed for BGP, as follows:

```
# on PE-1:
configure {
    router "Base" {
        bgp {
            delete export
```

In this example, BGP leaked IPv6 routes are advertised by BGP. For scenarios with only BGP routes, a dedicated BGP route leaking mechanism that preserves all attributes is preferred, as described in chapter "BGP Route Leaking" in the Unicast Routing Protocols volume in the *7450 ESS, 7750 SR, and 7950 XRS*

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

MD-CLI Advanced Configuration Guide - Part I. However, with the same configuration as in this chapter, it is possible to leak non-BGP routes and advertise them using BGP.

Conclusion

Routes learned in a VPRN can be leaked to the base router and advertised using routing protocols. The mechanism described in this chapter is protocol-independent: all kinds of routes can be leaked from a VRF to the GRT: local, static, IS-IS, OSPF, BGP routes, and so on. In some cases, it might be useful to leak the routes from a VPRN to the entire network using the routing protocol, in order to access the resources defined inside the VRF. Routes that are leaked from VPRNs to the GRT must be unique in the network where they will be advertised.

For BGP routes, the protocol-independent route leaking mechanism described here does not preserve the attributes, unlike the dedicated BGP route leaking feature.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Weighted ECMP for VPRN over RSVP-TE or SR-TE LSPs

This chapter provides information about Weighted Equal Cost Multipath (ECMP) for Virtual Private Routed Network (VPRN) over Resource Reservation Protocol with Traffic Engineering (RSVP-TE) or Segment Routing with Traffic Engineering (SR-TE) Label Switched Paths (LSPs).

Topics in this chapter include:

Applicability

The information and configuration in this chapter are based on SR OS Release 23.3.R2. Weighted load balancing over Multi Protocol Label Switching (MPLS) LSPs -as described in chapter *BGP Weighted ECMP*- is supported in SR OS Release 13.0.R1, and later. Weighted load balancing for VPRN with auto-bind-tunnel over RSVP-TE LSPs is supported in SR OS Release 15.0.R2, and later. Weighted load balancing for VPRN with auto-bind-tunnel over SR-TE LSPs is supported in SR OS Release 15.0.R4, and later.

Overview

Equal Load Balancing

In this chapter, ECMP refers to spraying traffic flows over multiple RSVP-TE or SR-TE LSPs within an ECMP set. ECMP spraying consists of hashing the relevant fields in the packet header and selecting the tunnel next-hop based on the modulo operation of the output of the hash and the number of LSPs present in the ECMP set. The maximum number of LSPs in the ECMP set is defined by the **ecmp** command.

Only LSPs with the same lowest LSP metric can be part of the ECMP set. If the number of such LSPs exceeds the maximum number of LSPs allowed in the ECMP set as defined by the **ecmp** command, the LSPs with the lowest tunnel IDs are selected first. By default, all LSPs in the ECMP set have the same weight, and traffic flows are spread evenly over all LSPs in the ECMP set, regardless of the bandwidth of the active path in the LSPs. By default, ECMP is enabled and set to 1.

Figure 323: Regular ECMP in AS 64496 shows that PE-1 sprays the traffic flows equally over two LSPs between PE-1 and PE-3. If three or more LSPs with the same lowest LSP metric were available from PE-1 to PE-3, only two of those would be used, because an ECMP value of 2 allows the traffic to be sprayed over two LSPs.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 323: Regular ECMP in AS 64496

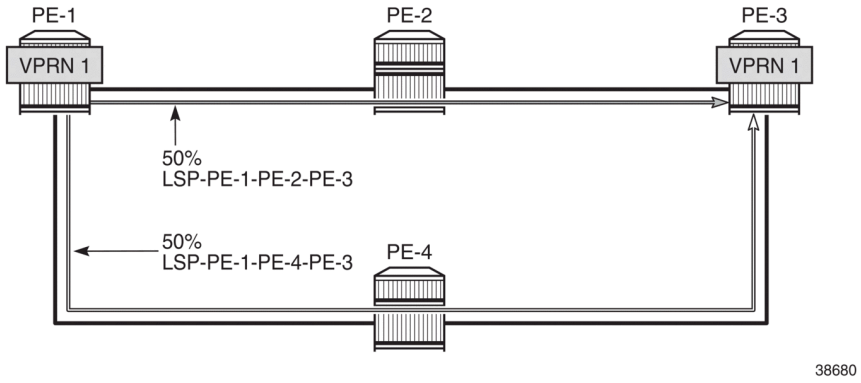

Unequal Load Balancing

Weighted ECMP sprays traffic flows over MPLS LSPs proportionally to the **load-balancing-weight** <*weight*> value configured on each MPLS LSP in the ECMP set. Figure 324: Weighted ECMP in AS 64496 shows that PE-1 forwards two thirds of the traffic flows on LSP-PE-1-PE-2-PE-3 with weight 2 and one third on LSP-PE-1-PE-4-PE-3 with weight 1. Each of the links can be link aggregation group (LAG) ports. For instance, when LSP-PE-1-PE-2-PE-3 uses LAG ports, 67% of the traffic is sprayed evenly over all ports belonging to the LAG.

Figure 324: Weighted ECMP in AS 64496

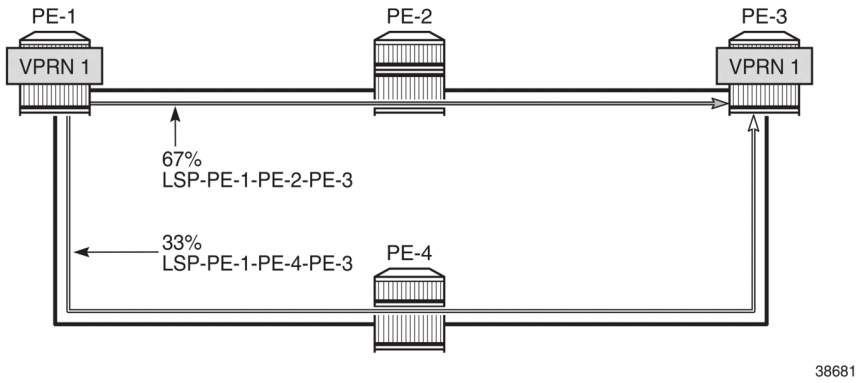

The LSP load balancing weight can be configured in an LSP template or on an LSP. By default, the load balancing weight equals zero, in which case regular ECMP applies.

The following command is used to configure the weight in an LSP template:

```
A:admin@PE-1# configure {
    router "Base" {
        mpls {
            lsp-template "LSPtemplate1" {
                load-balancing-weight ?

 load-balancing-weight <number>
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
 <number> - <1..4294967295>

    Load balancing weight for an MPLS LSP template

    Warning: Modifying this element toggles
    'configure router "Base" mpls lsp-template "LSPtemplate1" admin-state'
    automatically for the new value to take effect.
```

The following command is used to configure the weight on an LSP (for example on LSP "LSP-PE-1-PE-2-PE-3"):

```
A:admin@PE-1# configure {
   router "Base" {
        mpls {
            lsp "LSP-PE-1-PE-2-PE-3" {
                load-balancing-weight ?

 load-balancing-weight <number>
 <number> - <1..4294967295>

    Load balancing weight for an MPLS LSP
```

The LSP load balancing weight on LSP-PE-1-PE-2-PE-3 is configured with a value of 2, as follows:

```
configure {
   router "Base" {
        mpls {
            path "path-PE-1-PE-2-PE-3" {
                admin-state enable
                hop 10 {
                    ip-address 192.168.12.2
                    type strict
                }
                hop 20 {
                    ip-address 192.168.23.2
                    type strict
                }
            }
            lsp "LSP-PE-1-PE-2-PE-3" {
                admin-state enable
                type p2p-rsvp
                to 192.0.2.3
                path-computation-method local-cspf
                metric 100
                load-balancing-weight 2
                primary "path-PE-1-PE-2-PE-3" {
                }
            }
        }
```

Weighted ECMP is enabled in the **vprn "1" bgp-ipvpn mpls auto-bind-tunnel** context as follows:

```
configure {
    service {
        vprn "1" {
            customer "1"
            bgp-ipvpn {
                mpls {
                    auto-bind-tunnel {
                        ecmp 2
                        weighted-ecmp true
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    }
```

Weighted load balancing within a **vprn** context can be performed only when the next-hops are associated with the same neighbor and all LSPs in the ECMP set are configured with non-zero load balancing weights. If one or more LSPs in the ECMP set toward a specific next-hop do not have a load balancing weight configured, regular ECMP spraying is used. The weighted ECMP support for ECMP routes applies to both IPv4 and IPv6.

Additionally, it is possible to enable ECMP in the **vprn** context, with: **configure service vprn "1" ecmp** *<max-ecmp-routes>*, to control load balancing to a different next-hop. The **weighted-ecmp true** option in the **vprn "1" bgp-ipvpn mpls auto-bind-tunnel** context controls load balancing to the same next-hop.

Configuration

Figure 325: Example Topology shows the example topology with four PEs. VPRN 1 is configured on PE-1 and PE-3. A traffic generator is connected to VPRN 1 SAP 2/1/c36/2 on PE-1 and VPRN 1 SAP 2/1/c36/2 on PE-3. The traffic generator generates multiple traffic flows with random IP addresses and TCP/UDP port numbers. As a result, these flows are sprayed over different MPLS LSPs between PE-1 and PE-3.

Figure 325: Example Topology

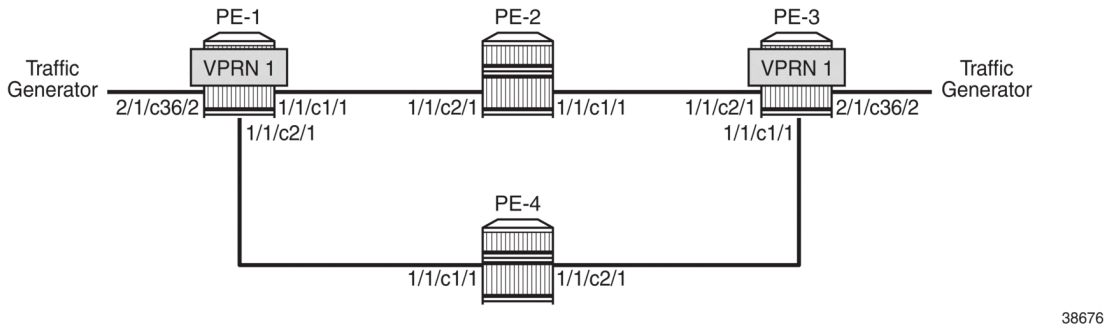

The initial configuration on the PEs includes the following:

- Cards, MDAs, ports
- Router interfaces
- IS-IS as IGP (alternatively, OSPF can be used) with traffic engineering enabled
- MPLS and RSVP enabled on all router interfaces

The initial configuration on PE-1 is as follows:

```
configure {
   router "Base" {
        interface "int-PE-1-PE-2" {
            port 1/1/c1/1
            ipv4 {
                primary {
                    address 192.168.12.1
                    prefix-length 30
                }
            }
        }
        interface "int-PE-1-PE-4" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            port 1/1/c2/1
            ipv4 {
                primary {
                    address 192.168.14.1
                    prefix-length 30
                }
            }
        }
        interface "system" {
            ipv4 {
                primary {
                    address 192.0.2.1
                    prefix-length 32
                }
            }
        }
        isis 0 {
            admin-state enable
            area-address [49.0001]
            traffic-engineering true
            interface "system" {
            }
            interface "int-PE-1-PE-2" {
                interface-type point-to-point
            }
            interface "int-PE-1-PE-4" {
                interface-type point-to-point
            }
        }
        mpls {
            admin-state enable
            interface "int-PE-1-PE-2" {
            }
            interface "int-PE-1-PE-4" {
            }
        }
        rsvp {
            admin-state enable
            interface "int-PE-1-PE-2" {
            }
            interface "int-PE-1-PE-4" {
            }
        }
    }
```

The initial configuration on the other PEs is similar.

With the preceding configuration, MPLS and RSVP are enabled on all interfaces, including the system interface, which is added automatically.

In the next sections, the following use cases are described:

- Weighted ECMP for VPRN with Auto-bind-tunnel RSVP-TE
- Weighted ECMP for an SDP used as a spoke SDP in a VPRN
- Weighted ECMP for VPRN with Auto-bind-tunnel SR-TE

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Weighted ECMP for VPRN with Auto-bind-tunnel RSVP-TE

On PE-1, the following paths and LSPs are configured. LSP-PE-1-PE-2-PE-3 is configured with a load balancing weight of 2; LSP-PE-1-PE-4-PE-3 is configured with a load balancing weight of 1.

```
configure {
   router "Base" {
        mpls {
            path "path-PE-1-PE-2-PE-3" {
                admin-state enable
                hop 10 {
                    ip-address 192.168.12.2
                    type strict
                }
                hop 20 {
                    ip-address 192.168.23.2
                    type strict
                }
            }
            path "path-PE-1-PE-4-PE-3" {
                admin-state enable
                hop 10 {
                    ip-address 192.168.14.2
                    type strict
                }
                hop 20 {
                    ip-address 192.168.34.1
                    type strict
                }
            }
            lsp "LSP-PE-1-PE-2-PE-3" {
                admin-state enable
                type p2p-rsvp
                to 192.0.2.3
                path-computation-method local-cspf
                metric 100
                load-balancing-weight 2
                primary "path-PE-1-PE-2-PE-3" {
                }
            }
            lsp "LSP-PE-1-PE-4-PE-3" {
                admin-state enable
                type p2p-rsvp
                to 192.0.2.3
                path-computation-method local-cspf
                metric 100
                load-balancing-weight 1
                primary "path-PE-1-PE-4-PE-3" {
                }
            }
        }
```

On PE-1, VPRN 1 is configured as follows. ECMP and weighted ECMP can be configured in the **vprn** context, for example, **configure service vprn "1" ecmp 2** and **configure service vprn "1" weighted-ecmp true** but it is not required when the next-hop for the MPLS LSPs is the same. In this example, ECMP and weighted ECMP are only configured in the **vprn "1" bgp-ipvpn mpls auto-bind-tunnel** context. The resolution filter only allows RSVP-TE tunnels, no other MPLS LSPs, such as LDP, BGP, or segment routing (SR) tunnels.

```
configure {
    service {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        vprn "1" {
            admin-state enable
            customer "1"
            description "CE-1"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:1"
                    vrf-target {
                        community "target:64496:1"
                    }
                    auto-bind-tunnel {
                        ecmp 2
                        weighted-ecmp true
                        resolution filter
                        resolution-filter {
                            rsvp true
                        }
                    }
                }
            }
            interface "loopback1" {
                loopback true
                ipv4 {
                    primary {
                        address 172.16.0.1
                        prefix-length 32
                    }
                }
                ipv6 {
                    address 2001:db8::1 {
                        prefix-length 128
                    }
                }
            }
            interface "int-CE-1-STC" {
                ipv4 {
                    primary {
                        address 192.168.11.1
                        prefix-length 24
                    }
                }
                ipv6 {
                    address 2001:db8::11:1 {
                        prefix-length 120
                    }
                }
                sap 2/1/c36/2 {
                }
            }
        }
```

The service configuration on PE-3 is similar.

VPRN 1 is dual stacked. Weighted ECMP applies to both IPv4 and IPv6 traffic streams. BGP is configured for the VPN-IPv4 and VPN-IPv6 address family to exchange the routes used in VPRN 1 between PE-1 and PE-3. The BGP configuration on PE-1 is as follows:

```
configure {
   router "Base" {
        autonomous-system 64496
        bgp {
            group "iBGP" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
            neighbor 192.0.2.3 {
                group "iBGP"
                peer-as 64496
                family {
                    vpn-ipv4 true
                    vpn-ipv6 true
                }
                export {
                    policy ["export-vpn-ipv4" "export-vpn-ipv6"]
                }
            }
        }
```

The BGP configuration on PE-3 is similar.

The export policies on PE-1 are defined as follows:

```
configure {
    policy-options {
        prefix-list "vpn-ipv4" {
            prefix 172.16.0.0/16 type longer {
            }
            prefix 192.168.11.0/24 type exact {
            }
        }
        prefix-list "vpn-ipv6" {
            prefix 2001:db8::/120 type longer {
            }
            prefix 2001:db8::11:0/120 type exact {
            }
        }
        policy-statement "export-vpn-ipv4" {
            entry 10 {
                from {
                    prefix-list ["vpn-ipv4"]
                }
                action {
                    action-type accept
                }
            }
        }
        policy-statement "export-vpn-ipv6" {
            entry 10 {
                from {
                    prefix-list ["vpn-ipv6"]
                }
                action {
                    action-type accept
                }
            }
        }
    }
```

The export policies on PE-3 are similar.

With ECMP enabled for MPLS LSPs with the same next-hop and two RSVP-TE LSPs available with equal metric, the route table of VPRN 1 on PE-1 shows two routes for each prefix with the same next-hop 192.0.2.3: one via RSVP LSP 1 and the other via RSVP LSP 3, as follows:

```
[/]
A:admin@PE-1# show router 1 route-table
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
172.16.0.1/32                                   Local   Local     00h00m58s  0
       loopback1                                                        0
172.16.0.3/32 [2]                               Remote  BGP VPN   00h00m10s  170
       192.0.2.3 (tunneled:RSVP:1)                                      100
172.16.0.3/32 [2]                               Remote  BGP VPN   00h00m10s  170
       192.0.2.3 (tunneled:RSVP:3)                                      100
192.168.11.0/24                                 Local   Local     00h00m58s  0
       int-CE-1-STC                                                     0
192.168.33.0/24 [2]                             Remote  BGP VPN   00h00m10s  170
       192.0.2.3 (tunneled:RSVP:1)                                      100
192.168.33.0/24 [2]                             Remote  BGP VPN   00h00m10s  170
       192.0.2.3 (tunneled:RSVP:3)                                      100
-------------------------------------------------------------------------------
No. of Routes: 6
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

The flag [2] indicates that next-hop 192.0.2.3 occurs twice for the prefix 172.16.0.3/32; next-hop 192.0.2.3 also occurs twice for the prefix 192.168.33.0/24.

The following IPv6 route table is similar, with next-hop 192.0.2.3 occurring twice for prefix 2001:db8::3/128 and twice for prefix 2001:db8::33:0/120.

```
[/]
A:admin@PE-1# show router 1 route-table ipv6

===============================================================================
IPv6 Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
2001:db8::1/128                                 Local   Local     00h00m57s  0
       loopback1                                                        0
2001:db8::3/128 [2]                             Remote  BGP VPN   00h00m10s  170
       192.0.2.3 (tunneled:RSVP:1)                                      100
2001:db8::3/128 [2]                             Remote  BGP VPN   00h00m10s  170
       192.0.2.3 (tunneled:RSVP:3)                                      100
2001:db8::11:0/120                              Local   Local     00h00m57s  0
       int-CE-1-STC                                                     0
2001:db8::33:0/120 [2]                          Remote  BGP VPN   00h00m10s  170
       192.0.2.3 (tunneled:RSVP:1)                                      100
2001:db8::33:0/120 [2]                          Remote  BGP VPN   00h00m10s  170
       192.0.2.3 (tunneled:RSVP:3)                                      100
-------------------------------------------------------------------------------
No. of Routes: 6
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The following tunnel table output on PE-1 shows that RSVP-TE LSP 1 goes via PE-2 (next-hop 192.168.12.2) and RSVP-TE LSP 3 via PE-4 (next-hop 192.168.14.2):

```
[/]
A:admin@PE-1# show router tunnel-table

===============================================================================
IPv4 Tunnel Table (Router: Base)
===============================================================================
Destination           Owner     Encap TunnelId  Pref   Nexthop        Metric
   Color
-------------------------------------------------------------------------------
192.0.2.3/32          rsvp      MPLS  1         7      192.168.12.2   100
192.0.2.3/32          rsvp      MPLS  3         7      192.168.14.2   100
-------------------------------------------------------------------------------
Flags: B = BGP or MPLS backup hop available
       L = Loop-Free Alternate (LFA) hop available
       E = Inactive best-external BGP route
       k = RIB-API or Forwarding Policy backup hop
===============================================================================
```

To verify the weighted load balancing between the two RSVP-TE LSPs, the traffic generator sends multiple IPv4 and IPv6 traffic flows with random IP addresses and TCP/UDP port numbers via PE-1 to PE-3. The traffic enters PE-1 through port 2/1/c36/2. When LSP-PE-1-PE-2-PE-3 is configured with weight 2 and LSP-PE-1-PE-4-PE-3 with weight 1, PE-1 forwards two thirds of the traffic via port 1/1/c1/1 toward PE-2 and one third of the traffic via port 1/1/c2/1 toward PE-3, as follows:

```
[/]
A:admin@PE-1# monitor port 1/1/c1/1 rate interval 3 repeat 3

===============================================================================
Monitor statistics for Port 1/1/c1/1
===============================================================================
                                                  Input                 Output
-------------------------------------------------------------------------------
---snip---
-------------------------------------------------------------------------------
At time t = 3 sec (Mode: Rate)
-------------------------------------------------------------------------------
Octets                                              107                 432147
Packets                                               2                    419
Errors                                                0                      0
Bits                                                856                3457176
Utilization (% of port capacity)                  ~0.00                   0.03
---snip---
===============================================================================


[/]
A:admin@PE-1# monitor port 1/1/c2/1 rate interval 3 repeat 3

===============================================================================
Monitor statistics for Port 1/1/c2/1
===============================================================================
                                                  Input                 Output
-------------------------------------------------------------------------------
---snip---
-------------------------------------------------------------------------------
At time t = 3 sec (Mode: Rate)
-------------------------------------------------------------------------------
Octets                                               95                 178293
Packets                                               1                    174
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
---snip---
===============================================================================


[/]
A:admin@PE-1# monitor port 2/1/c36/2 rate interval 3 repeat 3

===============================================================================
Monitor statistics for Port 2/1/c36/2
===============================================================================
                                                    Input                 Output
-------------------------------------------------------------------------------
---snip---
-------------------------------------------------------------------------------
At time t = 3 sec (Mode: Rate)
-------------------------------------------------------------------------------
Octets                                           622933                      0
Packets                                             608                      0
---snip---
===============================================================================
```

This can also be verified as follows:

```
[/]
A:admin@PE-1# show port 1/1/c1/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c1/1                                 49                           4532
                                      15056                       15493446
===============================================================================


[/]
A:admin@PE-1# show port 1/1/c2/1 statistics

===============================================================================
Port Statistics on Slot 1
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
1/1/c2/1                                 40                           3798
                                       6280                        6443384
===============================================================================


[/]
A:admin@PE-1# show port 2/1/c36/2 statistics

===============================================================================
Port Statistics on Slot 2
===============================================================================
Port                        Ingress Packets                 Ingress Octets
Id                           Egress Packets                  Egress Octets
-------------------------------------------------------------------------------
2/1/c36/2                             21249                       21758976
                                          0                              0
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Restrictions

All RSVP-TE LSPs in the ECMP set must have a load balancing weight configured. When at least one RSVP-TE LSP in the ECMP set is configured without weight, regular ECMP is applied.

RSVP-TE LSP without Weight in ECMP Set

If one of the RSVP-TE LSPs in the ECMP set does not have a load balancing weight configured, the traffic flows are sprayed equally between all RSVP-TE LSPs, regardless of the configured weight of the other RSVP-TE LSPs in the ECMP set.

On PE-1, LSP-PE-1-PE-2-PE-3 is configured without a load balancing weight, as follows:

```
configure {
   router "Base" {
        mpls {
            lsp "LSP-PE-1-PE-2-PE-3" {
                delete load-balancing-weight
            }
```

LSP-PE-1-PE-4-PE-3 is still configured with a load balancing weight of 1, but it is impossible to calculate its relative weight, because the sum of the weight values is not defined. Therefore, PE-1 reverts to regular ECMP for the load balancing between the two RSVP-TE LSPs to PE-3. When the traffic generator sends multiple traffic flows via PE-1 to PE-3, the load is spread equally over both RSVP-TE LSPs, as shown in the following monitor output. Port 1/1/c1/1 is used for traffic sent via LSP-PE-1-PE-2-PE-3 and port 1/1/c2/1 for traffic sent via LSP-PE-1-PE-4-PE-3.

```
[/]
A:admin@PE-1# monitor port 1/1/c1/1 rate interval 3 repeat 3

===============================================================================
Monitor statistics for Port 1/1/c1/1
===============================================================================
                                                    Input                 Output
-------------------------------------------------------------------------------
---snip---
-------------------------------------------------------------------------------
At time t = 3 sec (Mode: Rate)
-------------------------------------------------------------------------------
Octets                                                101                303429
Packets                                                 1                   294
Errors                                                  0                     0
Bits                                                  808               2427432
Utilization (% of port capacity)                    ~0.00                  0.02
---snip---
===============================================================================


[/]
A:admin@PE-1# monitor port 1/1/c2/1 rate interval 3 repeat 3

===============================================================================
Monitor statistics for Port 1/1/c2/1
===============================================================================
                                                    Input                 Output
-------------------------------------------------------------------------------
---snip---
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
At time t = 3 sec (Mode: Rate)
-------------------------------------------------------------------------------
Octets                                             61                315509
Packets                                             0                   306
---snip---
===============================================================================


[/]
A:admin@PE-1# monitor port 2/1/c36/2 rate interval 3 repeat 3

===============================================================================
Monitor statistics for Port 2/1/c36/2
===============================================================================
                                                  Input                Output
-------------------------------------------------------------------------------
---snip---
-------------------------------------------------------------------------------
At time t = 3 sec (Mode: Rate)
-------------------------------------------------------------------------------
Octets                                          602112                     0
Packets                                            588                     0
---snip---
===============================================================================
```

The configuration is restored as follows:

```
configure {
   router "Base" {
        mpls {
            lsp "LSP-PE-1-PE-2-PE-3" {
                load-balancing-weight 2
            }
```

Weighted ECMP for an SDP used as a spoke SDP in a VPRN

The following LSPs are configured on PE-1. The LSP load balancing weight values are 4 and 1 and the metric is 101 for both LSPs.

```
configure {
   router "Base" {
        mpls {
            admin-state enable
            lsp "LSP-PE-1-PE-2-PE-3-spoke" {
                admin-state enable
                type p2p-rsvp
                to 192.0.2.3
                path-computation-method local-cspf
                metric 101
                load-balancing-weight 4
                primary "path-PE-1-PE-2-PE-3" {
                }
            }
            lsp "LSP-PE-1-PE-4-PE-3-spoke" {
                admin-state enable
                type p2p-rsvp
                to 192.0.2.3
                path-computation-method local-cspf
                metric 101
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                load-balancing-weight 1
                primary "path-PE-1-PE-4-PE-3" {
                }
            }
        }
        rsvp {
            admin-state enable
            interface "int-PE-1-PE-2" {
            }
            interface "int-PE-1-PE-4" {
            }
        }
    }
```

Similar LSPs are configured on PE-3.

On PE-1, an SDP is configured, as follows:

```
configure {
    service {
        sdp 13 {
            admin-state enable
            delivery-type mpls
            far-end {
                ip-address 192.0.2.3
            }
            lsp "LSP-PE-1-PE-2-PE-3-spoke" { }
            lsp "LSP-PE-1-PE-4-PE-3-spoke" { }
        }
```

A similar SDP is configured on PE-3.

These SDPs are configured as spoke SDPs in a VPRN, as follows:

```
configure {
    service {
        vprn "1" {
            spoke-sdp 13:2 {
            }


configure {
    router "Base" {
        ldp {
            admin-state enable
        }
```

On PE-1, weighted ECMP is enabled on an SDP, as follows:

```
configure {
    service {
        sdp 13 {
            weighted-ecmp true
        }
```

The ECMP configuration on PE-3 is similar.

With ECMP enabled for MPLS LSPs with the same next-hop and two RSVP-TE LSPs available with equal metric, the route table of VPRN 1 on PE-1 shows one route for each prefix via the SDP tunnel, as follows:

```
[/]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
A:admin@PE-1# show router 1 route-table

===============================================================================
Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
172.16.0.1/32                                 Local   Local     00h41m48s  0
       loopback1                                                    0
172.16.0.3/32                                 Remote  BGP VPN   00h06m15s  170
       192.0.2.3 (tunneled)                                         0
192.168.11.0/24                               Local   Local     00h41m48s  0
       int-CE-1-STC                                                 0
192.168.33.0/24                               Remote  BGP VPN   00h06m15s  170
       192.0.2.3 (tunneled)                                         0
-------------------------------------------------------------------------------
No. of Routes: 4
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

The following tunnel table output on PE-1 shows that the preferred route to PE-3 is via the SDP tunnel. It uses RSVP-TE LSP 2 that goes via PE-2 (next-hop 192.168.12.2) and RSVP-TE LSP 4 that goes via PE-4 (next-hop 192.168.14.2):

```
[/]
A:admin@PE-1# show router tunnel-table

===============================================================================
IPv4 Tunnel Table (Router: Base)
===============================================================================
Destination           Owner     Encap TunnelId  Pref   Nexthop        Metric
   Color
-------------------------------------------------------------------------------
192.0.2.3/32          sdp       MPLS  13        5      192.0.2.3      0
---snip---
192.0.2.3/32          rsvp      MPLS  2         7      192.168.12.2   101
192.0.2.3/32          rsvp      MPLS  4         7      192.168.14.2   101
-------------------------------------------------------------------------------
Flags: B = BGP or MPLS backup hop available
       L = Loop-Free Alternate (LFA) hop available
       E = Inactive best-external BGP route
       k = RIB-API or Forwarding Policy backup hop
===============================================================================
```

To verify the weighted load balancing between the two RSVP-TE LSPs, the traffic generator sends multiple IPv4 traffic flows with random IP addresses and TCP/UDP port numbers via PE-1 to PE-3. The traffic enters PE-1 through port 2/1/c36/2. When LSP-PE-1-PE-2-PE-3-spoke is configured with weight 4 and LSP-PE-1-PE-4-PE-3-spoke with weight 1, PE-1 forwards four fifths of the traffic via port 1/1/c1/1 toward PE-2 and one fifth of the traffic via port 1/1/c2/1 toward PE-3, as follows:

```
[/]
A:admin@PE-1# monitor port 1/1/c1/1 rate interval 3 repeat 3

===============================================================================
Monitor statistics for Port 1/1/c1/1
===============================================================================
                                                 Input                Output
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
---snip---
-------------------------------------------------------------------------------
At time t = 3 sec (Mode: Rate)
-------------------------------------------------------------------------------
Octets                                              21                488922
Packets                                              0                   474
Errors                                               0                     0
Bits                                               168               3911376
Utilization (% of port capacity)                 ~0.00                  0.03
---snip---
===============================================================================


[/]
A:admin@PE-1# monitor port 1/1/c2/1 rate interval 3 repeat 3

===============================================================================
Monitor statistics for Port 1/1/c2/1
===============================================================================
                                                    Input                Output
-------------------------------------------------------------------------------
---snip---
-------------------------------------------------------------------------------
At time t = 3 sec (Mode: Rate)
-------------------------------------------------------------------------------
Octets                                            107                118204
Packets                                             1                   115
---snip---
===============================================================================


[/]
A:admin@PE-1# monitor port 2/1/c36/2 rate interval 3 repeat 3

===============================================================================
Monitor statistics for Port 2/1/c36/2
===============================================================================
                                                    Input                Output
-------------------------------------------------------------------------------
---snip---
-------------------------------------------------------------------------------
At time t = 3 sec (Mode: Rate)
-------------------------------------------------------------------------------
Octets                                          602112                     0
Packets                                            588                     0
---snip---
===============================================================================
```

Figure 326: Weighted ECMP over RSVP LSPs used in a spoke SDP shows how the traffic flows are sprayed over the two RSVP LSPs:

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 326: Weighted ECMP over RSVP LSPs used in a spoke SDP

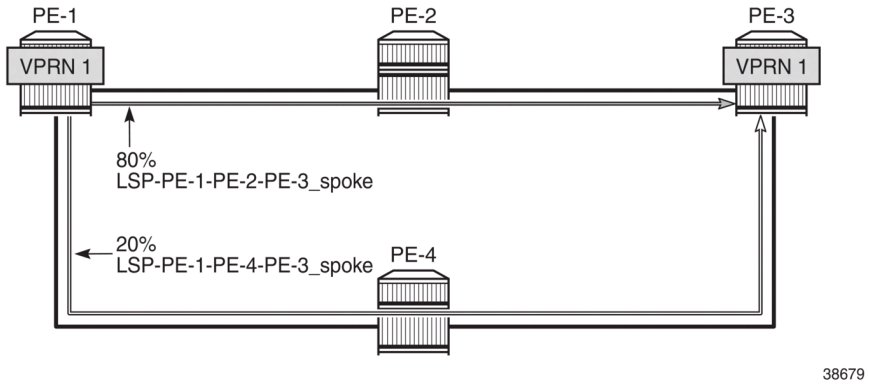

Weighted ECMP for VPRN with Auto-bind-tunnel SR-TE

The following configuration is added to enable SR-ISIS on PE-1.

```
configure {
   router "Base" {
        mpls-labels {
            sr-labels {
                start 20000
                end 20099
            }
        }
        isis 0 {
            advertise-router-capability as
            interface "system" {
                ipv4-node-sid {
                    label 20001
                }
            }
            segment-routing {
                admin-state enable
                prefix-sid-range {
                    start-label 20000
                    max-index 99
                }
            }
        }
    }
```

The configuration on the other PEs is identical, except for the **ipv4-node-sid label** value.

The following SR-TE LSPs are configured on PE-1. For more information about SR-TE LSPs, see the *Segment Routing Traffic Engineered Tunnels* chapter. The load balancing weight values are 75 and 25. The values 3 and 1, which have the same ratio, can be used instead.

```
configure {
   router "Base" {
        mpls {
            lsp "LSP-PE-1-PE-2-PE-3_SR-TE" {
                admin-state enable
                type p2p-sr-te
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                to 192.0.2.3
                load-balancing-weight 75
                primary "path-PE-1-PE-2-PE-3" {
                }
            }
            lsp "LSP-PE-1-PE-4-PE-3_SR-TE" {
                admin-state enable
                type p2p-sr-te
                to 192.0.2.3
                load-balancing-weight 25
                primary "path-PE-1-PE-4-PE-3" {
                }
            }
        }
```

The configuration on PE-3 is similar.

The following tunnel table on PE-1 shows two SR-TE tunnels with equal metrics: SR-TE tunnel 655362 has PE-2 as next-hop (192.168.12.2) and SR-TE tunnel 655363 has PE-4 as next-hop (192.168.14.2).

```
[/]
A:admin@PE-1# show router tunnel-table protocol sr-te

===============================================================================
IPv4 Tunnel Table (Router: Base)
===============================================================================
Destination           Owner     Encap TunnelId  Pref   Nexthop        Metric
   Color
-------------------------------------------------------------------------------
192.0.2.3/32          sr-te     MPLS  655362    8      192.168.12.2   16777215
192.0.2.3/32          sr-te     MPLS  655363    8      192.168.14.2   16777215
-------------------------------------------------------------------------------
Flags: B = BGP or MPLS backup hop available
       L = Loop-Free Alternate (LFA) hop available
       E = Inactive best-external BGP route
       k = RIB-API or Forwarding Policy backup hop
===============================================================================
```

The resolution filter for VPRN 1 is configured on PE-1 and PE-3 to only allow SR-TE tunnels. ECMP and weighted ECMP are enabled in the **vprn bgp-ipvpn mpls auto-bind-tunnel** context.

```
configure {
    service {
        vprn "1" {
            admin-state enable
            customer "1"
            description "CE-1"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:1"
                    vrf-target {
                        community "target:64496:1"
                    }
                    auto-bind-tunnel {
                        ecmp 2
                        weighted-ecmp true
                        resolution filter
                        resolution-filter {
                            sr-te true
                        }
                    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
        }
        interface "loopback1" {
            loopback true
            ipv4 {
                primary {
                    address 172.16.0.1
                    prefix-length 32
                }
            }
            ipv6 {
                address 2001:db8::1 {
                    prefix-length 128
                }
            }
        }
        interface "int-CE-1-STC" {
            ipv4 {
                primary {
                    address 192.168.11.1
                    prefix-length 24
                }
            }
            ipv6 {
                address 2001:db8::11:1 {
                    prefix-length 120
                }
            }
            sap 2/1/c36/2 {
            }
        }
    }
```

The following route table for VPRN 1 on PE-1 shows two entries for each remote prefix with the same next-hop 192.0.2.3 and a different SR-TE LSP: SR-TE tunnel 655362 and SR-TE tunnel 655363.

```
[/]
A:admin@PE-1# show router 1 route-table

===============================================================================
Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
172.16.0.1/32                                 Local   Local     01h01m17s  0
       loopback1                                                    0
172.16.0.3/32 [2]                             Remote  BGP VPN   00h02m24s  170
       192.0.2.3 (tunneled:SR-TE:655362)                            16777215
172.16.0.3/32 [2]                             Remote  BGP VPN   00h02m24s  170
       192.0.2.3 (tunneled:SR-TE:655363)                            16777215
192.168.11.0/24                               Local   Local     01h01m17s  0
       int-CE-1-STC                                                 0
192.168.33.0/24 [2]                           Remote  BGP VPN   00h02m24s  170
       192.0.2.3 (tunneled:SR-TE:655362)                            16777215
192.168.33.0/24 [2]                           Remote  BGP VPN   00h02m24s  170
       192.0.2.3 (tunneled:SR-TE:655363)                            16777215
-------------------------------------------------------------------------------
No. of Routes: 6
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
```

The following IPv6 route table for VPRN 1 is similar:

```
[/]
A:admin@PE-1# show router 1 route-table ipv6

===============================================================================
IPv6 Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
2001:db8::1/128                               Local   Local     01h01m16s  0
       loopback1                                                    0
2001:db8::3/128 [2]                           Remote  BGP VPN   00h02m24s  170
       192.0.2.3 (tunneled:SR-TE:655362)                            16777215
2001:db8::3/128 [2]                           Remote  BGP VPN   00h02m24s  170
       192.0.2.3 (tunneled:SR-TE:655363)                            16777215
2001:db8::11:0/120                            Local   Local     01h01m16s  0
       int-CE-1-STC                                                 0
2001:db8::33:0/120 [2]                        Remote  BGP VPN   00h02m24s  170
       192.0.2.3 (tunneled:SR-TE:655362)                            16777215
2001:db8::33:0/120 [2]                        Remote  BGP VPN   00h02m24s  170
       192.0.2.3 (tunneled:SR-TE:655363)                            16777215
-------------------------------------------------------------------------------
No. of Routes: 6
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

When multiple IPv4 and IPv6 traffic flows are sent from PE-1 to PE-3, the load balancing is weighted: 75% is sent via port 1/1/c1/1 toward PE-2 (LSP-PE-1-PE-2-PE-3_SR-TE) and 25% is sent via port 1/1/c2/1 toward PE-4 (LSP-PE-1-PE-4-PE-3_SR-TE), as follows:

```
[/]
A:admin@PE-1# monitor port 1/1/c1/1 rate interval 3 repeat 3

===============================================================================
Monitor statistics for Port 1/1/c1/1
===============================================================================
                                                  Input                 Output
-------------------------------------------------------------------------------
---snip---
-------------------------------------------------------------------------------
At time t = 3 sec (Mode: Rate)
-------------------------------------------------------------------------------
Octets                                              95                452765
Packets                                              1                   439
Errors                                               0                     0
Bits                                               760               3622120
Utilization (% of port capacity)                 ~0.00                  0.03
---snip---
===============================================================================


[/]
A:admin@PE-1# monitor port 1/1/c2/1 rate interval 3 repeat 3

===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Monitor statistics for Port 1/1/c2/1
================================================================================
                                                    Input                 Output
--------------------------------------------------------------------------------
---snip---
--------------------------------------------------------------------------------
At time t = 3 sec (Mode: Rate)
--------------------------------------------------------------------------------
Octets                                                 86                153866
Packets                                                 1                   150
---snip---
================================================================================


[/]
A:admin@PE-1# monitor port 2/1/c36/2 rate interval 3 repeat 3

================================================================================
Monitor statistics for Port 2/1/c36/2
================================================================================
                                                    Input                 Output
--------------------------------------------------------------------------------
---snip---
--------------------------------------------------------------------------------
At time t = 3 sec (Mode: Rate)
--------------------------------------------------------------------------------
Octets                                              600747                    0
Packets                                                587                    0
---snip---
================================================================================
```

Figure 327: Weighted ECMP over SR-TE LSPs in AS 64496 shows how the traffic flows are sprayed over the two SR-TE LSPs:

Figure 327: Weighted ECMP over SR-TE LSPs in AS 64496

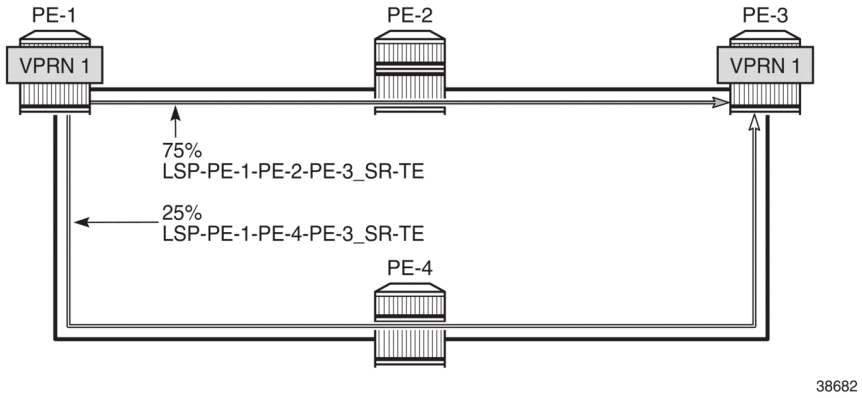

Conclusion

Operators can control how traffic in a VPRN is load balanced unequally over multiple transport tunnels by defining a load balancing weight factor on each LSP and enabling weighted ECMP in the VPRN. In this chapter, weighted ECMP for VPRN over transport LSPs is enabled for RSVP-TE tunnels and for SR-TE tunnels.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Multi-Service Integrated Service Adapter and Extended Services Appliance

This section provides MS-ISA and ESA configuration information for the following topics:

- Multi-Chassis IPSec Redundancy
- N:M MC-IPsec Redundancy

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Multi-Chassis IPSec Redundancy

This chapter provides information about multi-chassis IPSec redundancy configurations.

Topics in this chapter include:

Applicability

This initial version of this chapter was based on SR OS Release 10.0.R8, but the MD-CLI in the current edition corresponds to SR OS Release 22.10.R2.

Overview

Multi-Chassis IPSec redundancy (MC-IPSec) is a stateful inter-chassis IPSec failover mechanism. IPSec tunnel states are synchronized between the primary and standby chassis. A tunnel group failure on the primary chassis or a primary chassis failure could trigger MC-IPSec failover to the standby chassis.

The following are some highlights of this feature:

- Internet Key Exchange version 2 (IKEv2) only
- Multi-active tunnel group only
- The granularity of failover is tunnel group, which means a specific tunnel group could failover to the standby chassis independent of other tunnel groups on the primary chassis
- Both static and dynamic LAN-to-LAN tunnels are supported

This feature has the following building blocks:

- Primary chassis election: MC-IPSec mastership protocol (MIMP) runs between the chassis to elect a primary chassis with independent MIMP runs for each tunnel group
- Synchronization: multi-chassis synchronization (MCS) synchronizes the IPSec states between chassis
- Routing:
  - MC-IPSec-aware routing attracts traffic to the primary chassis
  - Shunting support
  - MC-IPSec-aware virtual router redundancy protocol (VRRP)

The figure Figure 328: MC-IPSec architecture shows two redundant IPSec chassis in the middle: a primary chassis and a standby chassis.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 328: MC-IPSec architecture

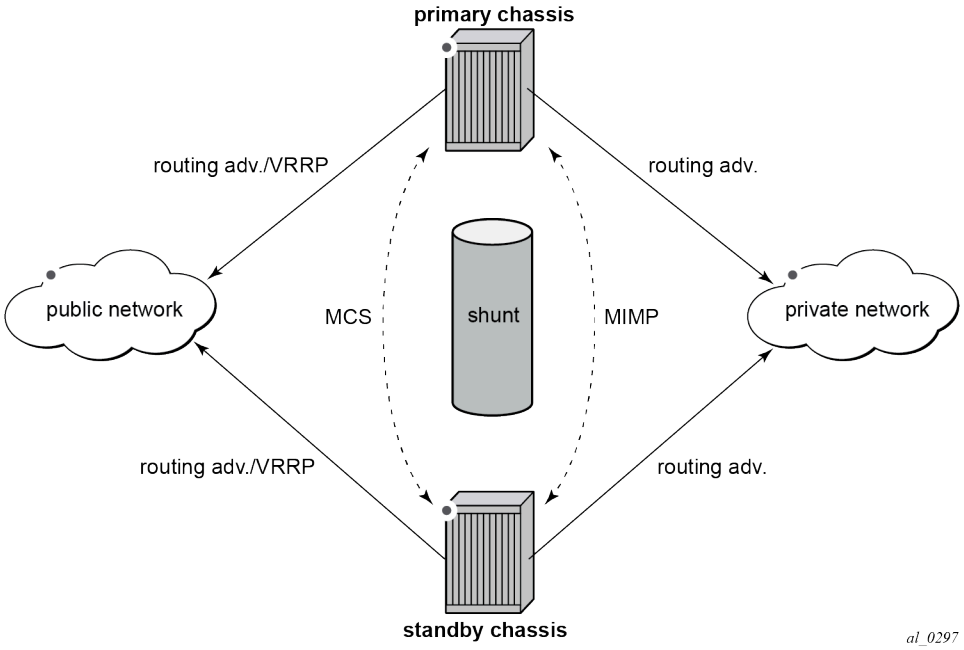

The fundamentals of MC-IPSec are:

- Only the primary chassis processes encapsulating security payload (ESP) and IKE traffic. If the standby chassis receives traffic, it shunts it to the primary chassis, if possible. The traffic is discarded if the standby chassis fails to shunt the traffic.
- The same local gateway address must be provisioned on both chassis.
- MC-IPSec does not synchronize configurations.
- MC-IPSec-aware routing attracts traffic to the primary chassis for both public and private services, which is achieved by exporting the corresponding IPSec routes to the routing protocol using a route policy and setting a different routing metric according to the MC-IPSec state.
- In case of a Layer 2 public network, MC-IPSec-aware VRRP can be used to trigger VRRP switchover upon MC-IPSec switchover.
- MCS synchronizes IPSec states between chassis so that existing IPSec tunnels do not need to be re-established upon switchover.
- MIMP elects mastership between two chassis, and it can also detect chassis failure and tunnel group failure; a central BFD session can be associated with MIMP to achieve fast chassis failure detection.

Configuration

The example topology is shown in the figure Figure 329: Example topology.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 329: Example topology

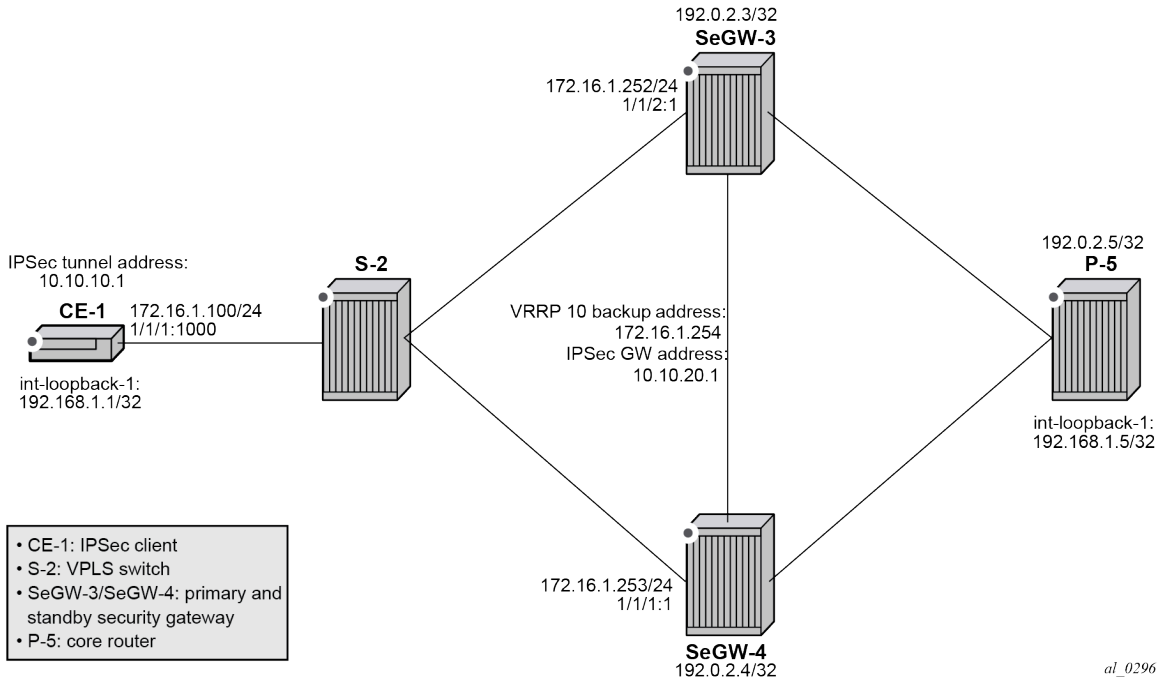

The example setup includes:

- an IPSec tunnel initiated by CE-1 and terminated on the primary chassis of the two SeGWs.
- a public IES service "IES-1" and a private VPRN service "VPRN-2" configured on CE-1, SeGW-3, and SeGW-4.
- VPRN "VPRN-2" (also) configured on P-5.
- a static LAN-to-LAN tunnel with pre-shared key.
- a local VPLS service "VPLS-3" on S-2 to simulate a Layer 2 switch.
- VRRP 10 between SeGW-3 and SeGW-4 to provide a backup address 192.168.1.254, which is the default next hop for CE-1.
- VRRP policy 1 bound to VRRP 10 on the primary chassis SeGW-3 to change the in-use priority upon MC-IPSec switchover.
- OSPF as IGP running in the base routing instance between SeGW-3, SeGW-4, and P-5.
- MP-BGP running between SeGW-3, SeGW-4, and P-5 for the VPN-IPv4 address family.

A ping in VPRN "VPRN-2" between loopback interface address 192.168.1.1 on CE-1 and 192.168.1.5 on P-5 is used to verify the connectivity over the IPSec tunnel.

The MC-IPSec configuration commands are shown below.

```
configure
    redundancy
        multi-chassis
            peer <ip-address>
                sync
                    ipsec
                    tunnel-group <1..64>
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        sync-tag <string>

                mc-ipsec
                    bfd-liveness <boolean>
                    discovery-interval
                        interval-secs <1..1800>
                        boot <1..1800>
                    hold-on-neighbor-failure <2..25>
                    keep-alive-interval <5..500>     # deciseconds
                    tunnel-group <1..64>
                        admin-state <boolean>
                        peer-group <1..64>
                        priority <0..255>
```

```
configure
    policy-options
        policy-statement <string>
            entry <1..4294967295>
                from
                    state ipsec-master-with-peer|ipsec-non-master|ipsec-master-without-peer
                    protocol
                        name ipsec
```

```
configure
    service
        ies <string>
            interface <string>
                dynamic-tunnel-redundant-nexthop <unicast-ipv4-address>
                static-tunnel-redundant-nexthop <unicast-ipv4-address>


configure
    service
        vprn <string>
            interface <string>
                dynamic-tunnel-redundant-nexthop <unicast-ipv4-address>
                static-tunnel-redundant-nexthop <unicast-ipv4-address>
```

```
configure
    isa
        tunnel-group <1..64>
            ipsec-responder-only <boolean>
```

```
configure
    vrrp
        policy <1..9999>
            priority-event
                mc-ipsec-non-forwarding <tunnel-grp-id>
                    hold-clear <1..86400 seconds>
                    hold-set <1..86400 seconds>
                    priority
                        priority-level <1..254>
                        event-type (delta|explicit)
```

The parameters are the following:

- in the **configure redundancy multi-chassis** context:

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- **peer** <*ip-address*> — This command creates or enters a multi-chassis peer. The peer address is by default the system address. This can be changed on the peer using the **configure redundancy multi-chassis peer source-address** command.
  - **sync** — This command enters the sync configuration context.
    - **ipsec** <*boolean*> — This command enables MCS to synchronize IPSec states.
    - **tunnel-group** <*tunnel-group-id*> **sync-tag** <*tag-name*> — This command enables MCS to synchronize the IPSec states of the specified tunnel group. The **sync-tag** parameter is used to match the tunnel group of the peer. The tunnel group states with the same **sync-tag** on both chassis will be synchronized.
  - **mc-ipsec** — This command enters the multi-chassis IPSec configuration context.
    - **bfd-liveness** <*boolean*> — The command **bfd-liveness** *true* enables tracking a central BFD session; if the BFD session goes down, then the system considers the peer as down and changes the MC-IPSec status of the configured tunnel group accordingly.

      The BFD session uses the source address of MCS as its source address and the MCS peer address as the destination address. Other BFD parameters are configured in the **bfd** context on the interface that the MCS source address resides on.

      The configuration of BFD is optional for MC-IPSec.
    - **discovery-interval interval-secs** <*interval-1*> **[boot** <*interval-2*>**]** — This command specifies the time interval that the tunnel group stays in **discovery** state. Interval 1 is used as discovery interval when a new tunnel group is added to multi-chassis redundancy (**mp-ipsec**); interval 2 is used as discovery interval after system boot-up. Interval 2 is optional, and when it is not specified, the value for interval 1 is used. Both intervals have a default value of 300 seconds.
    - **hold-on-neighbor-failure** <*2..25*> — This command specifies the number of keep-alive failures before considering the peer to be down. The default value is 3.
    - **keep-alive-interval** <*5..500*> — This command specifies the time interval of the mastership election protocol keep-alive packets in deciseconds. The default value is 10 deciseconds (1 s).
    - **tunnel-group** <*tunnel-group-id*> — This command enables multi-chassis redundancy for the specified tunnel group, or enters an already configured tunnel group context. The configured tunnel groups can failover independently.

      — **peer-group** <*tunnel-group-id*> — This command specifies the corresponding tunnel group ID on the peer node. The peer tunnel group ID is not necessarily equal to local tunnel group ID.

      — **priority** <*priority*> — This command specifies the local priority of the tunnel group, this is used to elect a primary chassis, where the higher number prevails. If the priorities are the same, then the peer which has more active ISAs wins; if the priority and the number of active ISAs are same, then the peer with higher IP address wins. The range is from 0 to 255 and the default value is 100.

- in a **from** statement of a route policy entry:
  - **state ipsec-master-with-peer | ipsec-non-master | ipsec-master-without-peer** — These commands specify the MC-IPSec state in a **from** statement of a route policy entry:
    - **ipsec-master-with-peer**: The tunnel group is the primary chassis with a peer reachable.
    - **ipsec-master-without-peer**: The tunnel group is the primary chassis with peer unreachable.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

    - **ipsec-non-master**: The tunnel group is not the primary chassis.
  - **protocol name ipsec** — This command specifies IPSec as protocol in a **from** statement of a route policy entry. **protocol name ipsec** refers to the /32 local gateway routes (of both static and dynamic tunnels) and reverse route of dynamic tunnel.
- on a public or private IPSec interface in an IES or VPRN service:
  - **static-tunnel-redundant-nexthop** <*ip-address*> and **dynamic-tunnel-redundant-nexthop** <*ip-address*> — These commands specify the redundant next hop address on a public or private IPSec interface (with public or private tunnel SAP) for a static and dynamic IPSec tunnel respectively. The specified next hop address is used by the standby chassis to shunt traffic to the primary chassis in case it receives any traffic. The next hop address is resolved in the routing table of the corresponding service.

    Note:
    - Shunting is supported over:
      - directly connected SAPs
      - spoke SDP terminated IP interfaces
    - Shunting over auto-bind tunnel is not supported.
    - Shunting does not work if the tunnel group is down.

- in the **isa tunnel-group <id>** context:
  - **ipsec-responder-only** <*boolean*> — With the command **ipsec-responder-only** *true* , the system only acts as IKE responder except for the automatic CHILD_SA rekey upon MC-IPSec switchover.

    This command is required for MC-IPSec support of static LAN-to-LAN tunnels.
- in the **vrrp policy <id> priority-event** context:
  - **mc-ipsec-non-forwarding** <*tunnel-grp-id*> — This command creates a VRRP policy priority event: *mc-ipsec-non-forwarding*, which is triggered whenever the specified tunnel group enters the non-forwarding state.
    - **hold-clear** <*seconds*> — This command configures the hold time before clearing the event. The range is from 0 to 86400 seconds and the default value is 0 s.
    - **hold-set** <*seconds*> — This command configures the hold time before setting the event. The range is from 0 to 86400 seconds and the default value is 0 s.
    - **priority** <*priority-level*> **explicit** — This command sets the VRRP in-use priority to the configured value upon the event. The range is from 0 to 254 and the default value is 0.

The initial configuration must include the following:

- The system time of SeGW-3 and SeGW-4 must be the same for the feature to work. Nokia recommends to use a time synchronization protocol such as NTP or SNTP.
- SeGW-3 and SeGW-4 must be IP reachable in the base routing instance because both MCS and MIMP run in the base routing instance.

Configuration of MC-IPSec

In this section, the following steps are described:

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- configure CE-1
- configure S-2
- configure P-5
- configure IPSec tunnel on SeGW-3
- enable MC-IPSec for tunnel group on SeGW-3
- configure MC-IPSec-aware routing on SeGW-3
- configure MC-IPSec-aware VRRP on SeGW-3
- configure SeGW-4

Configure CE-1

On CE-1, the following is configured:

- a public IES service "IES-1" and a private VPRN service "VPRN-2".
- a static default route pointing to the VRRP backup address 172.16.1.254.
- a static IPSec tunnel "tunnel-1" with local address 10.10.10.1 and remote address 10.10.20.1.
- a loopback interface in VPRN "VPRN-2" with address 192.168.1.1/32 to be used as source address for the ping command to verify the connectivity between CE-1 and P-5 over the IPSec tunnel.

The following base router configuration on CE-1 includes a static route with next hop 172.16.1.254, which is the VRRP backup address.

```
# on CE-1:
configure {
    router "Base" {
        interface "int-CE-1-S-2" {
            port 1/1/1:1000
            ipv4 {
                primary {
                    address 172.16.1.100
                    prefix-length 24
                }
            }
        }
        interface "system" {
            ipv4 {
                primary {
                    address 172.31.2.1
                    prefix-length 32
                }
            }
        }
        static-routes {
            route 0.0.0.0/0 route-type unicast {
                next-hop "172.16.1.254" {       # VRRP backup address
                    admin-state enable
                }
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

IPSec is configured as follows:

```
configure {
    ipsec {
        ike-policy 1 {
            ike-transform [1]
            ike-version-2 {
            }
            dpd {    # dead peer detection (on peer side; not on MC-IPSec chassis)
            }
        }
        ike-transform 1 {
        }
        ipsec-transform 1 {
        }
```

Tunnel group 1 is configured as follows:

```
configure {
    isa {
        tunnel-group 1 {
            admin-state enable
            isa-scale-mode tunnel-limit-2k
            primary 1/2
        }
```

The public IES service is configured as follows:

```
configure {
    service {
        ies "IES-1" {
            admin-state enable
            service-id 1
            customer "1"
            interface "int-IPsec-Public-1" {
                sap tunnel-1.public:1 {
                }
                ipv4 {
                    primary {
                        address 10.10.10.254
                        prefix-length 24
                    }
                }
            }
        }
```

The private VPRN service on CE-1 is configured as follows:

```
configure {
    service {
        vprn "VPRN-2" {
            admin-state enable
            service-id 2
            customer "1"
            ipsec {
                security-policy 1 {
                    entry 10 {
                        local-ip {
                            address 192.168.1.1/32
                        }
                        remote-ip {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        address 192.168.1.5/32
                    }
                }
            }
        }
        interface "int-IPsec-private-1" {
            tunnel true
            sap tunnel-1.private:1 {
                ipsec-tunnel "tunnel-1" {
                    admin-state enable
                    key-exchange {
                        dynamic {
                            ike-policy 1
                            ipsec-transform [1]
                            pre-shared-key "pass"
                        }
                    }
                    tunnel-endpoint {
                        local-gateway-address 10.10.10.1
                        remote-ip-address 10.10.20.1
                        delivery-service "IES-1"
                    }
                    security-policy {
                        id 1
                    }
                }
            }
        }
        interface "int-loopback-1" {
            loopback true
            ipv4 {
                primary {
                    address 192.168.1.1
                    prefix-length 32
                }
            }
        }
        static-routes {
            route 192.168.1.5/32 route-type unicast {
                ipsec-tunnel "tunnel-1" {
                    admin-state enable
                }
            }
        }
```

Configure S-2

On S-2, a local VPLS service 3 simulates a Layer 2 switch between CE-1, SeGW-3, and SeGW-4:

```
# on S-2:
configure {
    service {
        vpls "VPLS-3" {
            admin-state enable
            service-id 3
            customer "1"
            sap 1/1/c1/1:1 {
                description "to SAP in IES 1 on SeGW-3"
            }
            sap 1/1/c1/2:1000 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                description "to router interface in CE-1"
            }
            sap 1/1/c1/3:1 {
                description "to SAP in IES 1 on SeGW-4"
            }
        }
```

Configure P-5

P-5 simulates the core network router, connecting to SeGW-3 and SeGW-4. The configuration on P-5 includes the following:

- a loopback interface with address 192.168.1.5/32 in VPRN "VPRN-2", which is the destination address of the ping traffic from CE-1.
- an MP-BGP session for the VPN-IPv4 address family between P-5, SeGW-3, and SeGW-4.
- GRE spoke SDPs to connect to SeGW-3 and SeGW-4.

On P-5, the following router interfaces are configured in the base router. OSPF is used as IGP.

```
# on P-5:
configure {
    router "Base" {
        interface "int-P-5-SeGW-3" {
            port 1/1/c1/2:1000
            ipv4 {
                primary {
                    address 192.168.35.2
                    prefix-length 30
                }
            }
        }
        interface "int-P-5-SeGW-4" {
            port 1/1/c1/1:1000
            ipv4 {
                primary {
                    address 192.168.45.2
                    prefix-length 30
                }
            }
        }
        interface "system" {
            ipv4 {
                primary {
                    address 192.0.2.5
                    prefix-length 32
                }
            }
        }
        ospf 0 {
            admin-state enable
            area 0.0.0.0 {
                interface "int-P-5-SeGW-3" {
                }
                interface "int-P-5-SeGW-4" {
                }
                interface "system" {
                }
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

On P-5, the following GRE SDPs are configured toward SeGW-3 and SeGW-4:

```
configure {
    service {
        sdp 53 {
            admin-state enable
            description "GRE SDP toward SeGW-3"
            signaling off
            far-end {
                ip-address 192.0.2.3
            }
        }
        sdp 54 {
            admin-state enable
            description "GRE SDP toward SeGW-4"
            signaling off
            far-end {
                ip-address 192.0.2.4
            }
        }
```

VPRN "VPRN-2" is configured on P-5, as follows:

```
configure {
    service {
        vprn "VPRN-2" {
            admin-state enable
            service-id 2
            customer "1"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:2"
                    vrf-target {
                        community "target:64496:2"
                    }
                }
            }
            interface "int-loopback-1" {
                loopback true
                ipv4 {
                    primary {
                        address 192.168.1.5
                        prefix-length 32
                    }
                }
            }
            spoke-sdp 53:2 {
            }
            spoke-sdp 54:2 {
            }
        }
```

The BGP configuration on P-5 is as follows:

```
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            group "MPBGP" {
                type internal
                family {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    vpn-ipv4 true
                }
            }
            neighbor "192.0.2.3" {
                group "MPBGP"
            }
            neighbor "192.0.2.4" {
                group "MPBGP"
            }
        }
```

Configure IPSec tunnel on SeGW-3

The configuration on SeGW-3 is described in four consecutive sections. In this first section, the following is configured:

- the tunnel group, which must be in multi-active mode before MC-IPSec can be enabled.
- an interface "int-Redundant-1", which is a spoke-SDP terminated interface used for shunting.
- GRE SDP 34 toward SeGW-4 and GRE SDP 35 toward P-5.
- IPSec tunnel "tunnel-1" is the tunnel to CE-1; both SeGW-3 and SeGW-4 use the same local gateway address: 10.10.20.1.

The following configures tunnel group 1 on SeGW-3:

```
# on SeGW-3
configure {
    isa {
        tunnel-group 1 {
            admin-state enable
            isa-scale-mode tunnel-limit-2k
            ipsec-responder-only true
            multi-active {
                isa 1/2 { }
            }
        }
```

On SeGW-3, the following router interfaces are configured in the base router. A static route is configured toward CE-1. OSPF is the IGP used between SeGW-3, SeGW-4, and P-5.

```
configure {
    router "Base" {
        interface "int-SeGW-3-P-5" {
            port 1/1/1:1000
            ipv4 {
                primary {
                    address 192.168.35.1
                    prefix-length 30
                }
            }
        }
        interface "int-SeGW-3-SeGW-4" {
            port 1/1/3:1000
            ipv4 {
                primary {
                    address 192.168.34.1
                    prefix-length 30
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
        interface "system" {
            ipv4 {
                bfd {
                    admin-state enable
                }
                primary {
                    address 192.0.2.3
                    prefix-length 32
                }
            }
        }
        static-routes {
            route 10.10.10.0/24 route-type unicast {
                next-hop "172.16.1.100" {
                    admin-state enable
                }
            }
        }
        ospf 0 {
            admin-state enable
            area 0.0.0.0 {
                interface "int-SeGW-3-P-5" {
                }
                interface "int-SeGW-3-SeGW-4" {
                }
                interface "system" {
                }
            }
        }
```

The IPSec settings are as follows:

```
configure {
    ipsec {
        ike-policy 1 {
            ipsec-lifetime 7200
            ike-transform [1]
            ike-version-2 {
            }
        }
        ike-transform 1 {
            isakmp-lifetime 172800
        }
        ipsec-transform 1 {
        }
```

The GRE SDPs are configured as follows:

```
configure {
    service {
        sdp 34 {
            admin-state enable
            description "GRE SDP toward SeGW-4"
            signaling off
            far-end {
                ip-address 192.0.2.4
            }
        }
        sdp 35 {
            admin-state enable
            description "GRE SDP toward P-5"
            signaling off
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            far-end {
                ip-address 192.0.2.5
            }
        }
```

The public IES service is configured as follows. In a later step, a VRRP policy will be configured and applied.

```
configure {
    service {
        ies "IES-1" {
            admin-state enable
            service-id 1
            customer "1"
            interface "int-IPsec-Public-1" {
                static-tunnel-redundant-nexthop 192.168.34.2
                sap tunnel-1.public:1 {
                }
                ipv4 {
                    primary {
                        address 10.10.20.254
                        prefix-length 24
                    }
                }
            }
            interface "int-SeGW-3-S-2" {
                sap 1/1/2:1 {
                    description "SAP to switch S-2"
                }
                ipv4 {
                    primary {
                        address 172.16.1.252
                        prefix-length 24
                    }
                    vrrp 10 {
                        backup [172.16.1.254]
                        priority 200
                        ping-reply true
                    }
                }
            }
```

The private VPRN service "VPRN-2" is configured as follows:

```
configure {
    service {
        vprn "VPRN-2" {
            admin-state enable
            service-id 2
            customer "1"
            ipsec {
                security-policy 1 {
                    entry 10 {
                        local-ip {
                            address 192.168.1.5/32
                        }
                        remote-ip {
                            address 192.168.1.1/32
                        }
                    }
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:2"
                    vrf-target {
                        community "target:64496:2"
                    }
                }
            }
            interface "int-IPsec-Private-1" {
                tunnel true
                static-tunnel-redundant-nexthop 192.168.20.2
                sap tunnel-1.private:1 {
                    ipsec-tunnel "tunnel-1" {
                        admin-state enable
                        key-exchange {
                            dynamic {
                                ike-policy 1
                                ipsec-transform [1]
                                pre-shared-key "pass"
                            }
                        }
                        tunnel-endpoint {
                            local-gateway-address 10.10.20.1
                            remote-ip-address 10.10.10.1
                            delivery-service "IES-1"
                        }
                        security-policy {
                            id 1
                        }
                    }
                }
            }
            interface "int-Redundant-1" {
                ipv4 {
                    primary {
                        address 192.168.20.1
                        prefix-length 30
                    }
                }
                spoke-sdp 34:20 {
                    ingress {
                        vc-label 2049
                    }
                    egress {
                        vc-label 2048
                    }
                }
            }
            spoke-sdp 34:2 {
                description "SDP to SeGW-4"
            }
            spoke-sdp 35:2 {
                description "SDP to P-5"
            }
            static-routes {
                route 192.168.1.1/32 route-type unicast {
                    ipsec-tunnel "tunnel-1" {
                        admin-state enable
                    }
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Enable MC-IPSec for tunnel group 1 on SeGW-3

In this section, the following steps are described:

- Create a multi-chassis peer using the system address of SeGW-4.
- Enable MCS for IPSec and tunnel group 1.
- Enable MC-IPSec for the tunnel group with a configured priority 200.
- Bind a central BFD session to MC-IPSec from the system interface.

Multi-chassis peer 192.0.2.4 is configured and MCS and MC-IPSec are enabled for tunnel group 1:

```
# on SeGW-3:
configure {
    redundancy {
        multi-chassis {
            peer 192.0.2.4 {
                admin-state enable
                sync {
                    admin-state enable
                    ipsec true
                    tunnel-group 1 {
                        sync-tag "tag-1"
                    }
                }
                mc-ipsec {
                    bfd-liveness true
                    tunnel-group 1 {
                        admin-state enable
                        peer-group 1
                        priority 200
                    }
                }
            }
```

BFD is enabled for MC-IPSec in the preceding configuration. BFD is configured on the system interface 192.0.2.3:

```
configure {
    router "Base" {
        interface "system" {
            ipv4 {
                bfd {
                    admin-state enable
                }
                primary {
                    address 192.0.2.3
                    prefix-length 32
                }
            }
        }
```

Configure MC-IPSec-aware routing on SeGW-3

In this step, a route policy is defined and applied to VPRN "VPRN-2".

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Route policy "IPsec-to-MPBGP" exports static route 192.168.1.1/32 in VPRN "VPRN-2" to P-5. This policy sets the local preference of the prefix 192.168.1.1/32 according to the MC-IPSec state:

- for the **ipsec-master-with-peer** state: local preference 200
- for the **ipsec-non-master** state: local preference 100
- for the **ipsec-master-without-peer** state: local preference 200

The state **ipsec-master-without-peer** can be used to attract traffic to the designated primary chassis in case of "dual master" (meaning two chassis lose the MIMP connection in the base routing instance). In this example, SeGW-3 has local preference 200 and SeGW-4 has local preference 100 for **ipsec-master-without-peer**.

The route policy is configured as follows:

```
# on SeGW-3:
configure {
    policy-options {
        community "vprn2" {
            member "target:64496:2" { }
        }
        prefix-list "CE-1-Internal" {
            prefix 192.168.1.1/32 type exact {
            }
        }
        policy-statement "IPsec-to-MPBGP" {
            entry 10 {
                from {
                    prefix-list ["CE-1-Internal"]
                    state ipsec-master-with-peer
                }
                action {
                    action-type accept
                    local-preference 200
                    community {
                        add ["vprn2"]
                    }
                }
            }
            entry 20 {
                from {
                    prefix-list ["CE-1-Internal"]
                    state ipsec-non-master
                }
                action {
                    action-type accept
                    local-preference 100
                    community {
                        add ["vprn2"]
                    }
                }
            }
            entry 30 {
                from {
                    prefix-list ["CE-1-Internal"]
                    state ipsec-master-without-peer
                }
                action {
                    action-type accept
                    local-preference 200
                    community {
                        add ["vprn2"]
                    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
            }
            default-action {
                action-type accept
                community {
                    add ["vprn2"]
                }
            }
        }
```

The BGP configuration on SeGW-3 is as follows:

```
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            group "MPBGP" {
                type internal
                family {
                    vpn-ipv4 true
                }
            }
            neighbor "192.0.2.4" {
                group "MPBGP"
            }
            neighbor "192.0.2.5" {
                group "MPBGP"
            }
```

The route policy is applied as **vrf-export** in VPRN "VPRN-2":

```
configure {
    service {
        vprn "VPRN-2" {
            bgp-ipvpn {
                mpls {
                    vrf-export {
                        policy ["IPsec-to-MPBGP"]
                    }
```

Configure MC-IPSec-aware VRRP on SeGW-3

In this section, a VRRP policy is defined that uses the **mc-ipsec-non-forwarding** priority event to lower the in-use VRRP priority upon MC-IPSec switchover, which ensures VRRP and MC-IPSec have the same primary chassis. The VRRP instance needs to be in preempt mode.

This VRRP policy is only configured on the designated VRRP primary chassis SeGW-3, not on the standby chassis. The VRRP policy is applied to the interface "int-SeGW3-S-2" of IES "IES-1".

VRRP policy 1 is configured as follows:

```
# on SeGW-3:
configure {
    vrrp {
        policy 1 {
            priority-event {
                mc-ipsec-non-forwarding 1 {
                    priority {
                        priority-level 50
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        event-type explicit
                    }
                }
            }
        }
```

VRRP policy 1 is applied in VRRP instance 10 in the IES service:

```
configure {
    service {
        ies "IES-1" {
            interface "int-SeGW-3-S-2" {
                sap 1/1/2:1 {
                    description "SAP to switch S-2"
                }
                ipv4 {
                    primary {
                        address 172.16.1.252
                        prefix-length 24
                    }
                    vrrp 10 {
                        backup [172.16.1.254]
                        priority 200
                        ping-reply true
                        policy 1
                    }
                }
            }
---snip---
```

Configure SeGW-4

The configuration on the standby chassis SeGW-4 is similar, but with different priorities and without the VRRP policy.

The tunnel group is configured in multi-active mode:

```
# on SeGW-4:
configure {
    isa {
        tunnel-group 1 {
            admin-state enable
            isa-scale-mode tunnel-limit-2k
            ipsec-responder-only true
            multi-active {
                isa 1/2 { }
            }
        }
```

The MCS and MC-IPSec configuration is as follows:

```
configure {
    redundancy {
        multi-chassis {
            peer 192.0.2.3 {
                admin-state enable
                sync {
                    admin-state enable
                    ipsec true
                    tunnel-group 1 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        sync-tag "tag-1"
                    }
                }
                mc-ipsec {
                    bfd-liveness true
                    tunnel-group 1 {
                        admin-state enable
                        peer-group 1
                        priority 150
                    }
                }
            }
        }
```

The base router configuration on SeGW-4 includes the following router interfaces and a static route to CE-1. OSPF is used as IGP between SeGW-3, SeGW-4, and P-5.

```
configure {
    router "Base" {
        interface "int-SeGW-4-P-5" {
            port 1/1/2:1000
            ipv4 {
                primary {
                    address 192.168.45.1
                    prefix-length 30
                }
            }
        }
        interface "int-SeGW-4-SeGW-3" {
            port 1/1/3:1000
            ipv4 {
                primary {
                    address 192.168.34.2
                    prefix-length 30
                }
            }
        }
        interface "system" {
            ipv4 {
                bfd {
                    admin-state enable
                }
                primary {
                    address 192.0.2.4
                    prefix-length 32
                }
            }
        }
        static-routes {
            route 10.10.10.0/24 route-type unicast {
                next-hop "172.16.1.100" {
                    admin-state enable
                }
            }
        }
        ospf 0 {
            admin-state enable
            area 0.0.0.0 {
                interface "int-SeGW-4-P-5" {
                }
                interface "int-SeGW-4-SeGW-3" {
                }
                interface "system" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
            }
        }
```

The IPSec configuration is as follows:

```
configure {
    ipsec {
        ike-policy 1 {
            ipsec-lifetime 7200
            ike-transform [1]
            ike-version-2 {
            }
        }
        ike-transform 1 {
            isakmp-lifetime 172800
        }
        ipsec-transform 1 {
        }
```

The following route policy is configured on SeGW-4, The local preference is lower for the **ipsec-master-without-peer** state.

```
configure {
    policy-options {
        community "vprn2" {
            member "target:64496:2" { }
        }
        prefix-list "CE-1-Internal" {
            prefix 192.168.1.1/32 type exact {
            }
        }
        policy-statement "IPsec-to-MPBGP" {
            entry 10 {
                from {
                    prefix-list ["CE-1-Internal"]
                    state ipsec-master-with-peer
                }
                action {
                    action-type accept
                    local-preference 200
                    community {
                        add ["vprn2"]
                    }
                }
            }
            entry 20 {
                from {
                    prefix-list ["CE-1-Internal"]
                    state ipsec-non-master
                }
                action {
                    action-type accept
                    local-preference 100
                    community {
                        add ["vprn2"]
                    }
                }
            }
            entry 30 {
                from {
                    prefix-list ["CE-1-Internal"]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    state ipsec-master-without-peer
                }
                action {
                    action-type accept
                    local-preference 100
                    community {
                        add ["vprn2"]
                    }
                }
            }
            default-action {
                action-type accept
                community {
                    add ["vprn2"]
                }
            }
        }
```

The BGP configuration on SeGW-4 is as follows:

```
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            group "MPBGP" {
                type internal
                family {
                    vpn-ipv4 true
                }
            }
            neighbor "192.0.2.3" {
                group "MPBGP"
            }
            neighbor "192.0.2.5" {
                group "MPBGP"
            }
        }
```

The following GRE SDPs are configured:

```
configure {
    service {
        sdp 43 {
            admin-state enable
            description "GRE SDP toward SeGW-3"
            signaling off
            far-end {
                ip-address 192.0.2.3
            }
        }
        sdp 45 {
            admin-state enable
            description "GRE SDP toward P-5"
            signaling off
            far-end {
                ip-address 192.0.2.5
            }
        }
```

The public IES service is configured as follows:

```
configure {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    service {
        ies "IES-1" {
            admin-state enable
            service-id 1
            customer "1"
            interface "int-IPsec-Public-1" {
                static-tunnel-redundant-nexthop 192.168.34.1
                sap tunnel-1.public:1 {
                }
                ipv4 {
                    primary {
                        address 10.10.20.254
                        prefix-length 24
                    }
                }
            }
            interface "int-SeGW-4-S-2" {
                sap 1/1/1:1 {
                }
                ipv4 {
                    primary {
                        address 172.16.1.253
                        prefix-length 24
                    }
                    vrrp 10 {
                        backup [172.16.1.254]
                        ping-reply true
                    }
                }
            }
        }
```

The private VPRN service is configured as follows:

```
configure {
    service {
        vprn "VPRN-2" {
            admin-state enable
            service-id 2
            customer "1"
            ipsec {
                security-policy 1 {
                    entry 10 {
                        local-ip {
                            address 192.168.1.5/32
                        }
                        remote-ip {
                            address 192.168.1.1/32
                        }
                    }
                }
            }
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:2"
                    vrf-target {
                        community "target:64496:2"
                    }
                    vrf-export {
                        policy ["IPsec-to-MPBGP"]
                    }
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
            interface "int-IPsec-Private-1" {
                tunnel true
                static-tunnel-redundant-nexthop 192.168.20.1
                sap tunnel-1.private:1 {
                    ipsec-tunnel "tunnel-1" {
                        admin-state enable
                        key-exchange {
                            dynamic {
                                ike-policy 1
                                ipsec-transform [1]
                                pre-shared-key "pass"
                            }
                        }
                        tunnel-endpoint {
                            local-gateway-address 10.10.20.1
                            remote-ip-address 10.10.10.1
                            delivery-service "IES-1"
                        }
                        security-policy {
                            id 1
                        }
                    }
                }
            }
            interface "int-Redundant-1" {
                ipv4 {
                    primary {
                        address 192.168.20.2
                        prefix-length 30
                    }
                }
                spoke-sdp 43:20 {
                    ingress {
                        vc-label 2048
                    }
                    egress {
                        vc-label 2049
                    }
                }
            }
            spoke-sdp 43:2 {
                description "SDP to SeGW-3"
            }
            spoke-sdp 45:2 {
                description "SDP to P-5"
            }
            static-routes {
                route 192.168.1.1/32 route-type unicast {
                    ipsec-tunnel "tunnel-1" {
                        admin-state enable
                    }
                }
            }
        }
```

Verification

The following will be verified in this section:

- the MC-IPSec status and VRRP status on SeGW-3 and SeGW-4

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- the status of the IPSec tunnel on CE-1
- the status of the IPSec tunnel on the SeGWs

Verify the MC-IPSec status on SeGW-3 and SeGW-4

The following is verified:

- SeGW-3 is the primary chassis (**master**) and SeGW-4 is the standby for tunnel group 1 because SeGW-3 has the higher priority 200.
- SeGW-3 is the primary node for VRRP instance 10 and SeGW-4 is the backup.

SeGW-3 is the primary chassis in tunnel group 1 with priority 200:

```
[/]
A:admin@SeGW-3# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.4

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.4
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:09:10

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              200       Up             master
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
```

SeGW-4 is the standby chassis in tunnel group 1 with priority 150:

```
[/]
A:admin@SeGW-4# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.3

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.3
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:10:22

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              150       Up             standby
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
```

SeGW-3 is the primary node for VRRP instance 10:

```
[/]
A:admin@SeGW-3# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                  VR Id Own Adm  State       Base Pri   Msg Int
                                IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-3-S-2                  10    No  Up   Master       200       1
                                IPv4      Up   1           200        No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

SeGW-4 is backup for VRRP instance 10:

```
[/]
A:admin@SeGW-4# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                  VR Id Own Adm  State       Base Pri   Msg Int
                                IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-4-S-2                  10    No  Up   Backup       100       1
                                IPv4      Up   n/a         100        No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

Verify the IPSec tunnel on CE-1

The following is verified in this section:

- the connectivity between CE-1 and P-5
- the IPSec tunnel information

A ping command is launched from the loopback interface in VPRN "VPRN-2" on CE-1 to the loopback interface in VPRN "VPRN-2" on P-5:

```
[/]
A:admin@CE-1# ping 192.168.1.5 router-instance "VPRN-2"
PING 192.168.1.5 56 data bytes
64 bytes from 192.168.1.5: icmp_seq=1 ttl=63 time=2.44ms.
64 bytes from 192.168.1.5: icmp_seq=2 ttl=63 time=2.38ms.
64 bytes from 192.168.1.5: icmp_seq=3 ttl=63 time=2.38ms.
64 bytes from 192.168.1.5: icmp_seq=4 ttl=63 time=2.51ms.
64 bytes from 192.168.1.5: icmp_seq=5 ttl=63 time=2.50ms.
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
---- 192.168.1.5 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.38ms, avg = 2.44ms, max = 2.51ms, stddev = 0.053ms
```

The following command shows the IPSec tunnel information.

```
[/]
A:admin@CE-1# show ipsec tunnel

================================================================================
IPsec Tunnels
================================================================================
TunnelName                       LocalAddress      SvcId        Admn   Keying
  SapId                            RemoteAddress     DlvrySvcId   Oper   Sec
                                                                         Plcy
--------------------------------------------------------------------------------
tunnel-1                         10.10.10.1        2            Up     Dynamic
  tunnel-1.private:1               10.10.20.1        IES-1        Up     1
--------------------------------------------------------------------------------
IPsec Tunnels: 1
================================================================================
```

Verify the IPSec tunnel on the SeGWs

In this section, the following is verified:

- the MCS database is in-sync, so the tunnel status is up on both chassis.
- P-5 receives two VPN-IPv4 routes for prefix 192.168.1.1/32: the route from SeGW-3 has local preference 200; the route from SeGW-4 has local preference 100.

On both SeGWs, the IPSec tunnel with local address 10.10.20.1 and remote address 10.10.10.1 is up:

```
[/]
A:admin@SeGW-3# show ipsec tunnel

================================================================================
IPsec Tunnels
================================================================================
TunnelName                       LocalAddress      SvcId        Admn   Keying
  SapId                            RemoteAddress     DlvrySvcId   Oper   Sec
                                                                         Plcy
--------------------------------------------------------------------------------
tunnel-1                         10.10.20.1        2            Up     Dynamic
  tunnel-1.private:1               10.10.10.1        IES-1        Up     1
--------------------------------------------------------------------------------
IPsec Tunnels: 1
================================================================================
```

```
[/]
A:admin@SeGW-4# show ipsec tunnel

================================================================================
IPsec Tunnels
================================================================================
TunnelName                       LocalAddress      SvcId        Admn   Keying
  SapId                            RemoteAddress     DlvrySvcId   Oper   Sec
                                                                         Plcy
--------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
tunnel-1                          10.10.20.1        2            Up     Dynamic
  tunnel-1.private:1                10.10.10.1        IES-1        Up     1
-------------------------------------------------------------------------------
IPsec Tunnels: 1
===============================================================================
```

MCS is in sync on both SeGWs:

```
[/]
A:admin@SeGW-3# show redundancy multi-chassis sync

===============================================================================
Multi-chassis Peer Table
===============================================================================
Peer
-------------------------------------------------------------------------------
Peer IP Address         : 192.0.2.4
Description             : (Not Specified)
Authentication          : Disabled
Source IP Address       : 192.0.2.3
Admin State             : Enabled
Warm standby            : No
Remote warm standby     : No
Sub-mgmt options        :
  DHCP lease threshold  : Inactive
    Local / Remote      : -- / --
-------------------------------------------------------------------------------
Sync-status
-------------------------------------------------------------------------------
Client Applications     : IPsec
Sync Admin State        : Up
Sync Oper State         : Up
Sync Oper Flags         :
DB Sync State           : inSync
Num Entries             : 2
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
Rem Num Entries         : 2
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
===============================================================================
===============================================================================
```

```
[/]
A:admin@SeGW-4# show redundancy multi-chassis sync

===============================================================================
Multi-chassis Peer Table
===============================================================================
Peer
-------------------------------------------------------------------------------
Peer IP Address         : 192.0.2.3
Description             : (Not Specified)
Authentication          : Disabled
Source IP Address       : 192.0.2.4
Admin State             : Enabled
Warm standby            : No
Remote warm standby     : No
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Sub-mgmt options        :
  DHCP lease threshold  : Inactive
    Local / Remote      : -- / --
-------------------------------------------------------------------------------
Sync-status
-------------------------------------------------------------------------------
Client Applications     : IPsec
Sync Admin State        : Up
Sync Oper State         : Up
Sync Oper Flags         :
DB Sync State           : inSync
Num Entries             : 2
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
Rem Num Entries         : 2
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
===============================================================================
===============================================================================
```

The following command shows that P-5 received two VPN-IPv4 routes for prefix 192.168.1.1/32: one from SeGW-3 with local preference 200 and one from SeGW-4 with local preference 100:

```
[/]
A:admin@P-5# show router bgp routes vpn-ipv4
===============================================================================
 BGP Router ID:192.0.2.5        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv4 Routes
===============================================================================
Flag  Network                                           LocalPref   MED
      Nexthop (Router)                                  Path-Id     IGP Cost
      As-Path                                                       Label
-------------------------------------------------------------------------------
u*>i  64496:2:192.168.1.1/32                            200         None
      192.0.2.3                                         None        10
      No As-Path                                                    524286
*i    64496:2:192.168.1.1/32                            100         None
      192.0.2.4                                         None        10
      No As-Path                                                    524286
u*>i  64496:2:192.168.20.0/30                           100         None
      192.0.2.3                                         None        10
      No As-Path                                                    524286
*>i   64496:2:192.168.20.0/30                           100         None
      192.0.2.4                                         None        10
      No As-Path                                                    524286
u*>i  64496:2:192.168.20.1/32                           100         0
      192.0.2.3                                         None        10
      No As-Path                                                    524286
u*>i  64496:2:192.168.20.2/32                           100         0
      192.0.2.4                                         None        10
      No As-Path                                                    524286
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Routes : 6
===============================================================================
```

MC-IPSec failover scenarios

Two MC-IPSec failover scenarios are described in this section:

- MC-IPSec failover when MS-ISA is disabled
- MC-IPSec failover when the primary chassis SeGW-3 reboots

Failover when MS-ISA is disabled

Initially, MS-ISA is enabled, so SeGW-3 is the primary chassis and SeGW-4 is the standby:

```
[/]
A:admin@SeGW-3# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.4

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.4
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:09:10

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              200       Up             master
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
```

```
[/]
A:admin@SeGW-3# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                   VR Id Own Adm  State       Base Pri   Msg Int
                                 IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-3-S-2                   10    No  Up   Master       200       1
                                 IPv4      Up   1           200        No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

```
[/]
A:admin@SeGW-4# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.3
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.3
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:10:22

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              150       Up             standby
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
```

```
[/]
A:admin@SeGW-4# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                   VR Id Own Adm  State       Base Pri   Msg Int
                                 IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-4-S-2                   10    No  Up   Backup       100       1
                                 IPv4      Up   n/a         100        No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

The following command disables the MS-ISA on the primary chassis SeGW-3, which will trigger an MC-IPSec failover.

```
configure {
    card 1 {
        mda 2 {
            admin-state disable
```

With MS-ISA disabled, the MC-IPSec state of tunnel group 1 on SeGW-3 becomes **notEligible**, which means that the tunnel group is down, see the *7450 ESS, 7750 SR, and VSR Multiservice ISA and ESA Guide* for details description of MIMP states.:

```
[/]
A:admin@SeGW-3# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.4

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.4
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:09:10
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              200       Up             notEligible
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
```

SeGW-3 is backup for VRRP instance 10 with in-use priority 50, as per the VRRP policy 1:

```
[/]
A:admin@SeGW-3# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                   VR Id Own Adm  State       Base Pri   Msg Int
                                 IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-3-S-2                   10    No  Up   Backup       200       1
                                 IPv4      Up   1           50         No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

SeGW-4 is now the primary chassis in tunnel group 1. This is triggered by MC-IPSec failover, as per the **mc-ipsec-non-forwarding** event in VRRP policy 1.

```
[/]
A:admin@SeGW-4# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.3

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.3
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:10:22

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              150       Up             master
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
```

SeGW-4 is primary for VRRP instance 10;

```
[/]
A:admin@SeGW-4# show router vrrp instance
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
VRRP Instances
===============================================================================
Interface Name                  VR Id Own Adm  State       Base Pri   Msg Int
                                IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-4-S-2                  10    No  Up   Master       100       1
                                IPv4      Up   n/a         100        No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

The situation is restored by enabling MS-ISA on SeGW-3:

```
configure {
    card 1 {
        mda 2 {
            admin-state enable
```

MC-IPSec failover when primary chassis reboots

The following **tools** command on SeGW-3 triggers an MC-IPSec switchover:

```
tools perform redundancy multi-chassis mc-ipsec force-switchover tunnel-group 1

[/]
A:admin@SeGW-3# tools perform redundancy multi-chassis mc-ipsec force-switchover tunnel-group 1
WARNING! Forcing a mastership switchover may significantly impact traffic. Are you sure (y/n)?
 y
```

Before the failure condition takes place, SeGW-3 is the primary chassis for tunnel group 1:

```
[/]
A:admin@SeGW-3# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.4

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.4
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:09:10

===================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------
1              1              200       Up             master
-------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===================================================================
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

SeGW-3 is primary for VRRP instance 10:

```
[/]
A:admin@SeGW-3# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                  VR Id Own Adm  State       Base Pri   Msg Int
                                IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-3-S-2                  10    No  Up   Master       200       1
                                IPv4      Up   1           200        No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

SeGW-4 is the standby chassis for tunnel group 1:

```
[/]
A:admin@SeGW-4# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.3

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.3
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:10:22

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              150       Up             standby
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
```

The VRRP state on SeGW-4 is backup:

```
[/]
A:admin@SeGW-4# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                  VR Id Own Adm  State       Base Pri   Msg Int
                                IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-4-S-2                  10    No  Up   Backup       100       1
                                IPv4      Up   n/a         100        No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The following command reboots the primary chassis SeGW-3:

```
[/]
A:admin@SeGW-3# admin reboot card active now
```

While SeGW-3 reboots, the IPSec state of SeGW-4 becomes **eligible**:

```
[/]
A:admin@SeGW-4# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.3

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.3
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:10:22

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              150       Up             eligible
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
```

The VRRP state on SeGW-4 is primary (**master**):

```
[/]
A:admin@SeGW-4# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                  VR Id Own Adm  State       Base Pri   Msg Int
                                IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-4-S-2                  10    No  Up   Master       100       1
                                IPv4      Up   n/a         100        No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

When SeGW-3 comes up, the IPSec state of tunnel group 1 is **discovery**, which means that the system has not established the MIMP session with its peer yet.

```
[/]
A:admin@SeGW-3# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.4

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.4
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:24:41

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              200       Up             discovery
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
```

After a while, the preceding **show** command is repeated and the IPSec state for tunnel 1 on SeGW-3 is standby:

```
[/]
A:admin@SeGW-3# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.4

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.4
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:24:41

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              200       Up             standby
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
```

The VRRP state on SeGW-3 is backup:

```
[/]
A:admin@SeGW-3# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                   VR Id Own Adm  State       Base Pri   Msg Int
                                 IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-3-S-2                   10    No  Up   Backup       200       1
                                 IPv4      Up   1           50         No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

SeGW-4 is the primary chassis in MC-IPSec tunnel group 1:

```
[/]
A:admin@SeGW-4# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.3

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.3
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:10:22

===================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------
1              1              150       Up             master
-------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===================================================================
===============================================================================
```

SeGW-4 is the primary node for VRRP instance 10:

```
[/]
A:admin@SeGW-4# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                   VR Id Own Adm  State       Base Pri   Msg Int
                                 IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-4-S-2                   10    No  Up   Master       100       1
                                 IPv4      Up   n/a         100        No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

Configuration guidelines

The following is a list of guidelines for configuring MC-IPSec:

- To avoid high CPU load and issues in some complex cases, the following are suggestions for configuring the IKEv2 lifetime:
  - Both IKE_SA and CHILD_SA lifetime on MC-IPSec chassis (SeGW-3 and SeGW-4) should be around three times larger than on the IPSec peer CE-1.
  - With the first rule, the lifetime of the side with smaller lifetime (IPSec peer CE-1) should not be too small (these being the default values):
    - IKE_SA: >= 86400 seconds
    - CHILD_SA: >= 3600 seconds

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

  - With the first rule, on the side with smaller lifetime (IPSec peer CE-1), the IKE_SA lifetime must be at least 3 times larger than CHILD_SA lifetime.
- The IKE protocol is the control plane of IPSec, so IKE packets must be treated as high QoS priority in the end-to-end path of the public service. On the public interface, a SAP ingress QoS policy must be configured to ensure that IKE packets get high QoS priority.
- Configure **ipsec-responder-only true** under **tunnel-group** for static LAN-to-LAN tunnels.
- Enable dead peer detection (DPD) on the IPSec peer side (CE-1); disable DPD (default) on the MC-IPSec chassis side.
- The direct and redundant physical link between MC-IPSec chassis must be configured with sufficient bandwidth for MCS and shunting traffic, and proper QoS configuration to make sure the MIMP and MCS packets are treated as high priority traffic.
- The system time must be same on both MC-IPSec chassis.
- Make sure the protection status is **nominal** on both chassis before provoking a controlled switchover. The protection status can be displayed with the **show redundancy multi-chassis mc-ipsec peer ip-address <addr>** command.
- Wait at least five minutes between two consecutive switchovers if possible, to prevent a second switchover happening before the standby is ready to become the primary chassis.

Conclusion

MC-IPSec provides a stateful multi-chassis IPSec redundancy solution. This is very important in a carrier grade network, especially in applications such as mobile backhaul where high value mobile services run over IPSec tunnels.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

N:M MC-IPsec Redundancy

This chapter describes N:M MC-IPsec redundancy.

Topics in this chapter include:

- Applicability
- Overview
- Configuration
- Conclusion

Applicability

The information and MD-CLI configuration in this chapter are based on SR OS Release 22.10.R1.

The IPsec tunnel termination configuration described in this chapter requires an MS-ISA2 or an ESA server configured with a virtual machine. Configuration and setup for ISA2 or ESA are beyond the scope of this chapter; see the Multi-Chassis IPSec Redundancy chapter.

Overview

The N:M MC-IPsec redundancy model is a feature of the multi-chassis (MC) capabilities of SR OS when the router is deployed as Security Gateway (SeGW). N:M aims at enhancing the existing 1:1 redundancy model for IPsec tunnels. For the definition of N:M terminology and a description of its benefits, see the *7450 ESS, 7750 SR, and VSR Multiservice ISA and ESA Guide*.

The figure Figure 330: Three-node redundancy domain with a 2 DA + 1 DS model shows a three-node redundancy domain (RD) with the SeGWs SEG1, SEG2, and SEG3. SEG 1 and SEG 2 are designated active (DA) SeGWs and SEG 3 is designated standby (DS) SeGW.

Figure 330: Three-node redundancy domain with a 2 DA + 1 DS model

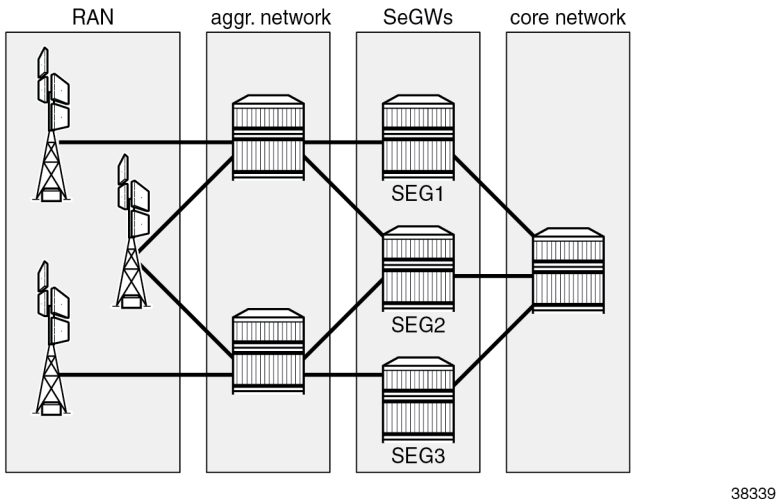

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Radio access network (RAN) elements are opening IPsec tunnels toward SeGW cluster tunnel endpoint IP addresses. The RAN, aggregation network, and core network are emulated with standard routing nodes. For this deployment, assume that connectivity between elements is established using routing protocols and, as for a classic SeGW router, the public side where traffic is encrypted is built on top of a public-side VPRN, while private side (clear-text traffic) is associated with another VPRN. ISA2 or ESA resources manage encryption and decryption operations across the VPRN boundary.

This chapter describes configuration of SeGW elements, as well as MD-CLI commands for tracking the functionality of N:M nodes in the same redundancy domain (RD).

Configuration

Assume that IP connectivity is established across the IP network elements in the architecture. It is beyond the scope of this chapter to describe how traffic is carried from the RAN to the SeGW or from SeGW to the mobile packet core. Among the protocols and techniques that are required to speed up convergence of routing, the bidirectional forwarding detection (BFD) protocol is especially useful to keep network convergence time in a range compatible with mobile traffic use case.

ISA2 or ESA setup for N:M

The nodes participating in the IPsec domain have a standard setup for ISA2 or ESA resources.

SEG1 and SEG 2 can each be configured like a classic SeGW, as follows:

```
[gl:/configure isa]
A:admin@SEG1# info
    tunnel-group 1 {
        admin-state enable
        isa-scale-mode tunnel-limit-2k
        ipsec-responder-only true
        multi-active {
            isa 1/2 { }
            active-isa-number 1
        }
        reassembly {
            max-wait-time 1200
        }
        stats-collection {
            isa-dp-cpu-usage true
        }
    }
```

The **active-isa-number** command specifies the number of active encryption and decryption elements. Nokia recommends implementing the same number of ISA2 and ESA resources among the nodes participating in the RD, which allows for the DS node to activate the same number of ISA2 or ESA resources when failover occurs. However, a failover can occur even if the DS node has a lower number of ISA2 or ESA resources available in its local pool. This allows operators to save costs, but if the ISA2 ore ESA resources on the initial DA nodes were fully loaded, the DS node cannot host all tunnels and the protection is only partial.

N:M redundancy allows DS nodes to cover multiple TGs, and therefore, multiple RDs. DS nodes may have more ISA2 or ESA resources than the DA nodes, because the DS nodes should be able to cover one or more DA node failures, with a maximum of 16.

The output from SEG2 is the same as for SEG1.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

SEG3 is configured as the DS node of the domain, where the configuration contains the **tunnel-member-pool** command:

```
[gl:/configure isa]
A:admin@SEG3# info
    tunnel-group 1 {
        admin-state enable
        isa-scale-mode tunnel-limit-2k
        ipsec-responder-only true
        multi-active {
            member-pool "MP1"
        }
        reassembly {
            max-wait-time 1200
        }
    }
    tunnel-member-pool "MP1" {
        isa 1/2 { }
    }
```

The **tunnel-member-pool** option defines the set of ISA2 or ESA resources used by the DS node during failures on active nodes. It is referenced in the tunnel group (TG) configuration, because multiple TGs can use the same tunnel member pool using the same set of ISA2 or ESA resources.

The output of the **show isa tunnel-member-pool** command lists ISA (ISA2 or ESA) members and their states. Under normal conditions, the ISA2 or ESA resource is not active on SEG3.

```
[gl:/configure isa]
A:admin@SEG3# /show isa tunnel-member-pool "MP1" detail

===============================================================================
ISA Tunnel Member Pool : MP1
Description            : (Not Specified)
Associated Tunnel Grps : 1
===============================================================================
Isa Members                  Active In Group     Last Configuration Change
-------------------------------------------------------------------------------
1/2                                              11/25/2022 12:10:14
-------------------------------------------------------------------------------
Number of Configured Entries: 1
Number of Active Entries: 0
===============================================================================
```

Redundancy domain configuration

The configuration of MC-IPsec as N:M starts by defining node roles and behavior. The configuration on SEG1 (with system IP address 192.0.2.1) is as follows:

```
[gl:/configure redundancy]
A:admin@SEG1# info
    multi-chassis {
        ipsec-domain 1 {
            admin-state enable
            designated-role active
            priority 250
            tunnel-group 1
        }
        peer 192.0.2.2  {
            admin-state enable
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            sync {
                admin-state enable
                ipsec true
            }
            mc-ipsec {
                bfd-liveness true
                domain 1 {
                    admin-state enable
                }
            }
        }
        peer 192.0.2.3 {
            admin-state enable
            sync {
                admin-state enable
                ipsec true
            }
            mc-ipsec {
                bfd-liveness true
                domain 1 {
                    admin-state enable
                }
            }
        }
    }
```

The preceding configuration example shows a multi-chassis IPsec domain, where the following domain characteristics have been specified:

- domain number – must be shared across all the nodes joining the redundancy domain (RD)
- designated role – DA or DS
- priority – required by the multi-chassis IPsec mastership protocol (MIMPv2) when an operationally active (OA) node must be elected. Setting a higher priority for an SeGW increases the likelihood of it being elected as the OA. In this case, SEG1 has the highest priority and DA role, so it is elected OA for RD 1.
- tunnel group – must be defined as per the ISA2 or ESA setup. The TG is always mapped to the RD in a 1:1 relationship
- peers – up to three peers can be added. While full-mesh peering between them is required, Nokia also recommends deploying highly redundant network paths between these peers.

  Each peer has its own CLI tree where the following characteristics must be defined:

  – the domain or domains the peer belongs to

  – the synchronization state for IPsec

  – whether BFD is applied to check peer liveliness.
- (optional) other parameters for keepalives, hold-time, and discovery-interval are configured with default values. Do not change these values unless a different setup is required under specific network conditions.

The configuration for the redundancy domain on SEG2 is the same as on SEG1, but with different IP addresses for peers and different priority:

```
A:admin@SEG2# info
    multi-chassis {
        ipsec-domain 1 {
            admin-state enable
            designated-role active
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            priority 240
            tunnel-group 1
        }
        peer 192.0.2.1 {
            admin-state enable
            sync {
                admin-state enable
                ipsec true
            }
            mc-ipsec {
                bfd-liveness true
                domain 1 {
                    admin-state enable
                }
            }
        }
        peer 192.0.2.3 {
            admin-state enable
            sync {
                admin-state enable
                ipsec true
            }
            mc-ipsec {
                bfd-liveness true
                domain 1 {
                    admin-state enable
                }
            }
        }
    }
```

The designated role of SEG2 is **active**, which means SEG2 behaves similarly to the 1:1 model where tunnel states are synchronized with SEG1 and immediately pushed to ISA2 or ESA resources. This behavior allows for a very quick failover when SEG1 experiences a failure.

The priority is 240, which is lower than for SEG1. As a result, SEG1 receives node role DA and is operationally active (OA) while SEG2 receives node role DA and is operationally standby (OS).

The RD configuration for DS SEG3 is as follows:

```
[gl:/configure redundancy multi-chassis]
A:admin@SEG3# info
    ipsec-domain 1 {
        admin-state enable
        designated-role standby
        priority 230
        tunnel-group 1
    }
    peer 192.0.2.1 {
        admin-state enable
        sync {
            admin-state enable
            ipsec true
        }
        mc-ipsec {
            bfd-liveness true
            domain 1 {
                admin-state enable
            }
        }
    }
    peer 192.0.2.2 {
        admin-state enable
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        sync {
            admin-state enable
            ipsec true
        }
        mc-ipsec {
            bfd-liveness true
            domain 1 {
                admin-state enable
            }
        }
    }
```

The peer configuration is similar to those of other nodes where BFD liveliness is enabled.

The designated role is standby (DS). The default value in the configuration is not shown from the **info** command.

The priority is 230 but the node role is DS. The DS node will not become OA because the DA role of SEG1 and SEG2 always prevails when electing the OA, regardless of priority value. Therefore, a DS node can become OA only if there are no DA nodes available in the domain.

After the setup of MC IPsec RD is completed across all the nodes, **show** commands can be used to track RD behavior and state:

```
A:admin@SEG1# show redundancy multi-chassis ipsec-domain 1

===============================================================================
Multi-Chassis IPsec Domain: 1
===============================================================================
Designated Role    : active            Priority          : 250
Tunnel Group       : 1                 Revertive         : false
Admin State        : Up                Protection Status : nominal
Router Id          : 192.0.2.1         Current Active    : 192.0.2.1
Activity State     : active
===============================================================================

===============================================================================
Domain 1 Adjacencies
===============================================================================
Peer                                    Oper    Remote       Remote
  Router-Id                             State   Activity     Designated
                                                State        Role
-------------------------------------------------------------------------------
192.0.2.2                                 Up      standby      active
  192.0.2.2
192.0.2.3                                 Up      standby      standby
  192.0.2.3
-------------------------------------------------------------------------------
Domain Adjacency Entries found: 2
===============================================================================

===============================================================================
Multi-Chassis Tunnel Statistics
===============================================================================
                             Static              Dynamic
-------------------------------------------------------------------------------
Installed                    0                   7
Installing                   0                   0
Standby Dormant              0                   0
Awaiting Config              0                   0
Failed                       0                   0
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The output shows important information about the redundancy domain:

- the designated role of the node – active or standby
- the activity state based on fault conditions – active or standby
- the protection status – "nominal" means that the nodes are synchronized.
- the domain adjacencies – list of peers and their activity state and designated role
- the tunnel statistics – in this case, seven dynamic tunnels are established

The same **show** command executed on SEG2 provides similar output, with differences for the priority and the designated role. The seven tunnels are shown in the "Installed" state because SEG2 is a DA node.

The same **show** command on DS SEG3 shows the following:

```
A:admin@SEG3# show redundancy multi-chassis ipsec-domain 1

===============================================================================
Multi-Chassis IPsec Domain: 1
===============================================================================
Designated Role    : standby            Priority           : 230
Tunnel Group       : 1                  Revertive          : false
Admin State        : Up                 Protection Status  : nominal
Router Id          : 192.0.2.3          Current Active     : 192.0.2.1
Activity State     : standby
===============================================================================

===============================================================================
Domain 1 Adjacencies
===============================================================================
Peer                                      Oper    Remote       Remote
  Router-Id                               State   Activity     Designated
                                                  State        Role
-------------------------------------------------------------------------------
192.0.2.1                                   Up      active       active
  192.0.2.1
192.0.2.2                                   Up      standby      active
  192.0.2.2
-------------------------------------------------------------------------------
Domain Adjacency Entries found: 2
===============================================================================

===============================================================================
Multi-Chassis Tunnel Statistics
===============================================================================
                             Static              Dynamic
-------------------------------------------------------------------------------
Installed                    0                   0
Installing                   0                   0
Standby Dormant              0                   7
Awaiting Config              0                   0
Failed                       0                   0
===============================================================================
```

Relevant information from the SEG3 CLI output, apart from the activity state, the designated role, and the peer’s state, is the tunnel state, which is now marked as "Standby Dormant".

Tunnels on SEG3 are not installed on the ISA2 or ESA; rather, they are stored in the router CPM and are kept ready to be offloaded on the ISA2 or ESA resources connected to the router. These tunnels are offloaded as soon as SEG3 becomes OA, following a node reboot, failure, or manual switchover.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Services configuration

The tunnels opened by RAN elements are terminated in a public-side VPRN IP address called TEIP (the public side can also be made on a IES service). Assume that the RAN nodes are using a single tunnel setup with a single IKE_SA, whereas the Child_SA's number is specific to the deployment. The configuration of this public side VPRN is the same for all three nodes and follows the standard SeGW setup:

```
[gl:/configure service vprn "100"]
A:admin@SEG1# info
    vprn "100" {
        admin-state enable
        description "public side"
        customer "1"
        ipsec {
            multi-chassis-shunt-interface "to_SEG2_Shunt" {
                next-hop {
                    address 10.1.12.2
                }
            }
            multi-chassis-shunt-interface "to_SEG3_Shunt" {
                next-hop {
                    address 10.1.13.2
                }
            }
            multi-chassis-shunting-profile "MCSPROF1" {
                peer 192.0.2.2 {
                    multi-chassis-shunt-interface "to_SEG2_Shunt"
                }
                peer 192.0.2.3 {
                    multi-chassis-shunt-interface "to_SEG3_Shunt"
                }
            }
        }
        interface "PUBLIC1" {
            multi-chassis-shunting-profile "MCSPROF1"
            sap tunnel-1.public:100 {
                ipsec-gateway "IPSECGW1" {
                    admin-state enable
                    default-tunnel-template 1
                    ike-policy 1
                    pre-shared-key "uCLxzS3PxoW0foPjmAKJ/Wv41hy6O3H76tg=" hash2
                    default-secure-service {
                        service-name "200"
                        interface "PRIVATE1"
                    }
                    local {
                        gateway-address 10.51.100.1
                    }
                }
            }
            ipv4 {
                primary {
                    address 198.51.100.2
                    prefix-length 24
                }
            }
        }
        interface "to_SEG2_Shunt" {
            spoke-sdp 2000:1 {
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            ipv4 {
                primary {
                    address 10.1.12.1
                    prefix-length 30
                }
            }
        }
        interface "to_SEG3_Shunt" {
            spoke-sdp 3000:1 {
            }
            ipv4 {
                primary {
                    address 10.1.13.1
                    prefix-length 30
                }
            }
        }
        ospf 0 {
            export-policy ["EXPORT_OSPF"]
        }
```

The parts of the configuration that are exclusive of N:M are those related to shunt-link setup.

The **multi-chassis-shunting-profile** command can be found under the **ipsec** configuration for the IES or VPRN service, where the multi-chassis shunting (MCS) profile is required to map each peer to a dedicated shunt interface. The MCS profile is referenced under the interface where the IPsec gateway is configured. In this scenario, peer 192.0.2.2 is reached through the to_SEG2_Shunt interface, which is defined under the same VPRN as an interface built on top of sdp:2000:1.

A full mesh of shunt interfaces is made across the RD, for both public and private side services.

```
A:admin@SEG1# show ipsec multi-chassis-shunt-interface service "100"

===============================================================================
IPsec Multi-Chassis Shunt Interfaces
===============================================================================
Service Id  MC Shunt Interface Name          Next Hop                Resolved
-------------------------------------------------------------------------------
100         to_SEG2_Shunt                    10.1.12.2               Yes
100         to_SEG3_Shunt                    10.1.13.2               Yes
-------------------------------------------------------------------------------
No. of IPsec MC Shunt Interfaces: 2
===============================================================================
```

The **show ipsec multi-chassis-shunt-interface service** command shows the liveliness of shunt interfaces and information on the next-hop resolution, whereas the **show ipsec multi-chassis-shunting-profile service** command provides a summary of the MCS profile and associated peers:

```
A:admin@SEG1# show ipsec multi-chassis-shunting-profile service "100"

===============================================================================
Multi-Chassis Shunting Profile Params Entries
===============================================================================
Service Id  MC Shunting Profile Name                MC Shunt Interface Name
            Peer
-------------------------------------------------------------------------------
100         MCSPROF1                                to_SEG2_Shunt
            192.0.2.2
100         MCSPROF1                                to_SEG3_Shunt
            192.0.2.3
-------------------------------------------------------------------------------
No. of IPsec MC Shunting Profile Params Entries: 2
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
```

The SDP full mesh must be configured on both sides, as shown in the figure Figure 331: SDP full mesh.

Figure 331: SDP full mesh

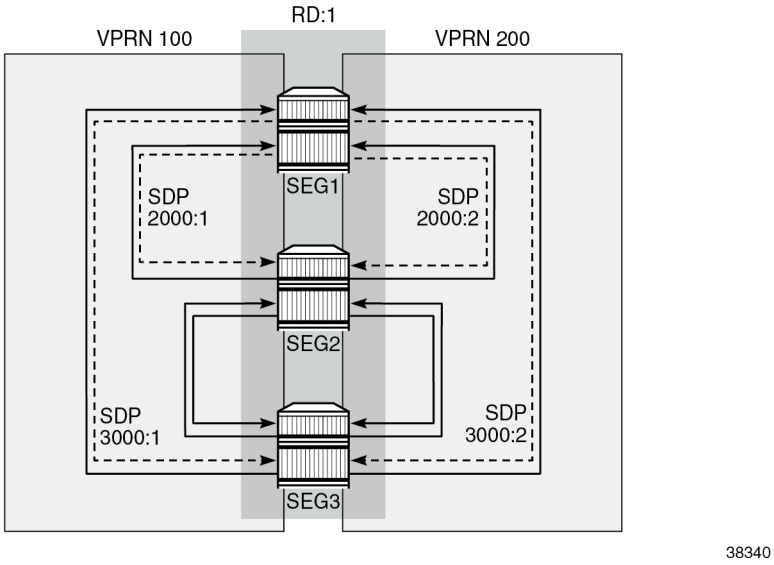

Note: Only the SDPs from SEG1 are shown with IDs.

The shunt link can be built from a standard spoke SDP or from a port-based interface. In this example, the following spoke SDPs are used in the public-side VPRN 100:

```
A:admin@SEG1# show service id "100" sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
2000:1           Spok     192.0.2.2       Up      Up        524285    524285
3000:1           Spok     192.0.2.3       Up      Up        524283    524285
-------------------------------------------------------------------------------
Number of SDPs : 2
-------------------------------------------------------------------------------
===============================================================================
```

The **show** output for the private-side VPRN 200 looks similar to that for the public-side VPRN, except for the SDP IDs and label values:

```
A:admin@SEG1# show service id "200" sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
2000:2           Spok     192.0.2.2       Up      Up        524284    524284
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
3000:2           Spok     192.0.2.3       Up      Up        524282    524284
-------------------------------------------------------------------------------
Number of SDPs : 2
-------------------------------------------------------------------------------
===============================================================================
```

There are no routing policy changes from the 1:1 MC-IPsec cluster, although this example could have a more complex routing setup, considering that the number of routers in a domain is higher than in the 1:1 model. The following configuration shows the SEG1-2-3 export policy used on the public side where the OSPF protocol is used under VPRN 100:

```
[gl:/configure policy-options policy-statement "EXPORT_OSPF"]
A:admin@SEG1# info
    description "EXPORT TEIP OSPF - PUBLIC SIDE"
    entry 10 {
        from {
            state ipsec-master-with-peer
            protocol {
                name [ipsec]
            }
        }
        action {
            action-type accept
            tag 100
            metric {
                set 30
            }
        }
    }
    entry 20 {
        from {
            state ipsec-non-master
            protocol {
                name [ipsec]
            }
        }
        action {
            action-type accept
            tag 100
            metric {
                set 190
            }
        }
    }
    entry 30 {
        from {
            state ipsec-master-without-peer
            protocol {
                name [ipsec]
            }
        }
        action {
            action-type accept
            tag 100
            metric {
                set 40
            }
        }
    }
    default-action {
        action-type reject
    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

On SEG2, only the metrics are different and are aligned with DA priorities:

```
[gl:/configure policy-options policy-statement "EXPORT_OSPF"]
A:admin@SEG2# info
    policy-statement "EXPORT_OSPF" {
        description "EXPORT TEIP OSPF - PUBLIC SIDE"
        entry 10 {
            from {
                state ipsec-master-with-peer
                protocol {
                    name [ipsec]
                }
            }
            action {
                action-type accept
                tag 200
                metric {
                    set 60
                }
            }
        }
        entry 20 {
            from {
                state ipsec-non-master
                protocol {
                    name [ipsec]
                }
            }
            action {
                action-type accept
                tag 200
                metric {
                    set 190
                }
            }
        }
        entry 30 {
            from {
                state ipsec-master-without-peer
                protocol {
                    name [ipsec]
                }
            }
            action {
                action-type accept
                tag 200
                metric {
                    set 50
                }
            }
        }
        default-action {
            action-type reject
        }
    }
```

On SEG3, the export policy is as follows:

```
[gl:/configure policy-options policy-statement "EXPORT_OSPF"]
A:admin@SEG3# info
    policy-statement "EXPORT_OSPF" {
        description "EXPORT TEIP OSPF - PUBLIC SIDE"
        entry 10 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            from {
                state ipsec-master-with-peer
                protocol {
                    name [ipsec]
                }
            }
            action {
                action-type accept
                tag 300
                metric {
                    set 90
                }
            }
        }
        entry 20 {
            from {
                state ipsec-non-master
                protocol {
                    name [ipsec]
                }
            }
            action {
                action-type accept
                tag 300
                metric {
                    set 195
                }
            }
        }
        entry 30 {
            from {
                state ipsec-master-without-peer
                protocol {
                    name [ipsec]
                }
            }
            action {
                action-type accept
                tag 300
                metric {
                    set 60
                }
            }
        }
        default-action {
            action-type reject
        }
    }
```

The export policy on the private-side VPRN is made with the same concept as the public side, but is not shown here.

Note: Parts of the configuration where the parameters remain the same as those in classic SeGW deployments (either stand-alone or 1:1) have not been added to this chapter. This information is described in the Multi-Chassis IPSec Redundancy chapter.

On the private side of SeGWs, a different VPRN is required, as per standard IPsec configuration. The private-side VPRN configuration on SEG1 is as follows:

```
[gl:/configure service vprn "200"]
A:admin@SEG1# info
    admin-state enable
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    description "private segw testing"
    customer "1"
    ipsec {
        multi-chassis-shunt-interface "to_SEG2_Shunt" {
            next-hop {
                address 10.2.12.2
            }
        }
        multi-chassis-shunt-interface "to_SEG3_Shunt" {
            next-hop {
                address 10.2.13.2
            }
        }
        multi-chassis-shunting-profile "MCSPROF1" {
            peer 192.0.2.2 {
                multi-chassis-shunt-interface "to_SEG2_Shunt"
            }
            peer 192.0.2.3 {
                multi-chassis-shunt-interface "to_SEG3_Shunt"
            }
        }
    }
    bgp-ipvpn {
        mpls {
            admin-state enable
            route-distinguisher "300:4"
        }
    }
    interface "PRIVATE1" {
        tunnel true
        multi-chassis-shunting-profile "MCSPROF1"
        sap tunnel-1.private:100 {
        }
    }
    interface "to_SEG2_Shunt" {
        ipv4 {
            primary {
                address 10.2.12.1
                prefix-length 30
            }
        }
        spoke-sdp 2000:2 {
        }
    }
    interface "to_SEG3_Shunt" {
        ipv4 {
            primary {
                address 10.2.13.1
                prefix-length 30
            }
        }
        spoke-sdp 3000:2 {
        }
    }
```

As the configuration shows, the same setup of shunt links is required on the private side to allow path resiliency in case of faults for the traffic going downstream from core toward the RAN.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Failure scenario – active node experiences a power failure

N:M can be triggered by different fault conditions, such as a complete node failure, an ISA2 or ESA failure, or a manual switchover executed with the **tools** command. In this scenario, complete node failures are simulated. When SEG1 experiences a node failure, SEG2 takes over. When SEG2 fails too, SEG3 takes over and remains the only node with active tunnels.

The initial scenario has SEG1 and SEG2 configured as DA nodes, while SEG3 is the DS node for the domain configured as **ipsec-domain 1**. The state can be verified with the **show redundancy multi-chassis ipsec-domain 1** command (as shown above in the Redundancy domain configurationsection).

As soon as SEG1 experiences a node failure, SEG2 takes over:

```
A:admin@SEG2# show redundancy multi-chassis ipsec-domain 1

===============================================================================
Multi-Chassis IPsec Domain: 1
===============================================================================
Designated Role    : active            Priority          : 240
Tunnel Group       : 1                 Revertive         : false
Admin State        : Up                Protection Status : notReady
Router Id          : 192.0.2.2         Current Active    : 192.0.2.2
Activity State     : active
===============================================================================

===============================================================================
Domain 1 Adjacencies
===============================================================================
Peer                                    Oper    Remote      Remote
  Router-Id                             State   Activity    Designated
                                                State       Role
-------------------------------------------------------------------------------
192.0.2.1                               Down    unknown     unknown
  0.0.0.0
192.0.2.3                               Up      standby     standby
  192.0.2.3
-------------------------------------------------------------------------------
Domain Adjacency Entries found: 2
===============================================================================

===============================================================================
Multi-Chassis Tunnel Statistics
===============================================================================
                            Static              Dynamic
-------------------------------------------------------------------------------
Installed                   0                   7
Installing                  0                   0
Standby Dormant             0                   0
Awaiting Config             0                   0
Failed                      0                   0
===============================================================================
```

Although the protection status, as seen from SEG2 and SEG3, is initially "notReady", it changes to "nominal" after few minutes. From the SEG2 and SEG3 point of view, SEG1 is unreachable, and its activity state remains unknown. Log 99 also records the failure event:

```
A:admin@SEG2# show log log-id 99

===============================================================================
Event Log 99 log-name 99
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Description : Default System Log
Memory Log contents  [size=500   next event=187  (not wrapped)]

186 2022/12/13 14:05:32.534 UTC WARNING: MC_REDUNDANCY #2047 Base MC-IPSEC-DOMAIN 1
"Protection status for the multi-chassis ipsec domain 1 changed to nominal"

185 2022/12/13 14:02:19.611 UTC MINOR: VRTR #2061 Base 192.0.2.1
"BFD: Local Discriminator 1 BFD session on node 192.0.2.1 is down due to noHeartBeat "

---snip---

179 2022/12/13 14:02:19.124 UTC WARNING: MC_REDUNDANCY #2004 Base
"The Sync status of peer 192.0.2.1 changed to outOfSync"

178 2022/12/13 14:02:18.746 UTC WARNING: MC_REDUNDANCY #2046 Base MC-IPSEC-DOMAIN 1
"Multi-chassis ipsec domain 1 local activity state changed from standby to active because an
 inter-chassis link went down. The active router in the domain is 192.0.2.2."
```

Next, SEG2 also experiences a full node failure, and SEG3 takes over:

```
A:admin@SEG3# show redundancy multi-chassis ipsec-domain 1

===============================================================================
Multi-Chassis IPsec Domain: 1
===============================================================================
Designated Role    : standby           Priority          : 230
Tunnel Group       : 1                 Revertive         : false
Admin State        : Up                Protection Status : notReady
Router Id          : 192.0.2.3         Current Active    : 192.0.2.3
Activity State     : eligible
===============================================================================

===============================================================================
Domain 1 Adjacencies
===============================================================================
Peer                                    Oper    Remote      Remote
  Router-Id                             State   Activity    Designated
                                                State       Role
-------------------------------------------------------------------------------
192.0.2.1                               Down    unknown     unknown
  0.0.0.0
192.0.2.2                               Down    unknown     unknown
  0.0.0.0
-------------------------------------------------------------------------------
Domain Adjacency Entries found: 2
===============================================================================

===============================================================================
Multi-Chassis Tunnel Statistics
===============================================================================
                             Static              Dynamic
-------------------------------------------------------------------------------
Installed                    0                   7
Installing                   0                   0
Standby Dormant              0                   0
Awaiting Config              0                   0
Failed                       0                   0
===============================================================================
```

Both SEG1 and SEG2 are seen as operationally down with an unknown activity state. On SEG3, the tunnel states have been copied from the CPM to the ISA2 or ESA entities and are now shown as "Installed", rather than "Standby Dormant". As soon as SEG1 or SEG2 are back up, the **revertive** flag configured

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

under the **ipsec-domain** command determines if the tunnels are kept on the current active DS node or if they are moved back to SEG1 ownership.

Failure scenario – using the tools command line

A planned failure condition is commonly seen when executing software upgrades or hardware maintenance on SeGW nodes, which leverages the **tools** command line utility to move tunnels toward other peering nodes.

The initial state is the same as for the previous example where SEG1 is initially the operationally active DA.

The following tools command triggers a switchover and therefore causes all the tunnels installed on the operationally active DA node to move on another node in the domain, selected by the **auto** flag in this case.

```
A:admin@SEG1# tools perform redundancy multi-chassis mc-ipsec force-switchover domain 1 auto
 now
```

To specify a peer IP address among those available in the domain, the **to <peer_ip>** option could be used instead of **auto**.

The following output shows the domain state as seen from SEG1 after the execution of the tools command:

```
A:admin@SEG1# show redundancy multi-chassis ipsec-domain 1

===============================================================================
Multi-Chassis IPsec Domain: 1
===============================================================================
Designated Role    : active             Priority           : 250
Tunnel Group       : 1                  Revertive          : false
Admin State        : Up                 Protection Status  : notReady
Router Id          : 192.0.2.1          Current Active     : 192.0.2.2
Activity State     : standby
===============================================================================

===============================================================================
Domain 1 Adjacencies
===============================================================================
Peer                                    Oper    Remote       Remote
  Router-Id                             State   Activity     Designated
                                                State        Role
-------------------------------------------------------------------------------
192.0.2.2                               Up      active       active
  192.0.2.2
192.0.2.3                               Up      standby      standby
  192.0.2.3
-------------------------------------------------------------------------------
Domain Adjacency Entries found: 2
===============================================================================

===============================================================================
Multi-Chassis Tunnel Statistics
===============================================================================
                            Static              Dynamic
-------------------------------------------------------------------------------
Installed                   0                   7
Installing                  0                   0
Standby Dormant             0                   0
Awaiting Config             0                   0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Failed                        0                   0
===============================================================================
```

As shown in the output, the current active node is SEG2 (192.0.2.2). The **auto** flag forced all the traffic to move across the second preferred active node in the domain, which is SEG2.

The protection status, as seen from SEG2, changes to "nominal" after a few minutes:

```
A:admin@SEG2# show redundancy multi-chassis ipsec-domain 1

===============================================================================
Multi-Chassis IPsec Domain: 1
===============================================================================
Designated Role    : active             Priority           : 240
Tunnel Group       : 1                  Revertive          : false
Admin State        : Up                 Protection Status  : nominal
Router Id          : 192.0.2.2          Current Active     : 192.0.2.2
Activity State     : active
===============================================================================

===============================================================================
Domain 1 Adjacencies
===============================================================================
Peer                                    Oper    Remote       Remote
  Router-Id                             State   Activity     Designated
                                                State        Role
-------------------------------------------------------------------------------
192.0.2.1                                 Up      standby      active
  192.0.2.1
192.0.2.3                                 Up      standby      standby
  192.0.2.3
-------------------------------------------------------------------------------
Domain Adjacency Entries found: 2
===============================================================================

===============================================================================
Multi-Chassis Tunnel Statistics
===============================================================================
                              Static              Dynamic
-------------------------------------------------------------------------------
Installed                     0                   7
Installing                    0                   0
Standby Dormant               0                   0
Awaiting Config               0                   0
Failed                        0                   0
===============================================================================
```

After maintenance operations on SEG1 have been completed and the node is operational (which can be verified using the **show** commands described in this chapter), the operator reverts services and traffic back to SEG1. For this purpose and in this specific example, the same **tools** command can be used. The **auto** flag selects SEG1, according to its highest priority in the domain. If more predictability is required in the selection choice, the **to <peer_ip>** flag can be used, as in this example:

```
A:admin@SEG2# tools perform redundancy multi-chassis mc-ipsec force-switchover domain 1 to
 192.0.2.1 now
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Conclusion

N:M adds a level of redundancy to an already efficient redundancy model; it ensures that RAN elements stay connected to the core network under a wide range of failure conditions. SR OS uses a full set of commands to implement this feature, available for both classic and MD-CLI. N:M also gives network engineers and architects the capability to deploy SeGW services with greater flexibility; for example, to deploy super-resilient SeGW clusters to serve high-density RAN areas, or to introduce cost-optimized solutions with an acceptable level of automated fault recovery.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Triple Play Service Delivery Architecture

This section provides TPSDA configuration information for the following topics:

- BNG Dual-Homing with EVPN VPWS in the Access Network and SRv6 Transport
- Diameter Base Protocol: Establishing a Diameter Peer Connection
- L2TP for Subscriber Access — LAC
- Vport-Based Load Balancing on a LAG

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

BNG Dual-Homing with EVPN VPWS in the Access Network and SRv6 Transport

This chapter describes BNG dual-homing with EVPN VPWS in the access network and SRv6 transport.

Topics in this chapter include:

- Applicability
- Overview
- Configuration
- Conclusion
- Appendix with configuration files

Applicability

The information and configuration in this chapter is based on SR OS Release 22.10.R1. The feature described in this chapter applies to all FP4 and above based SR systems; it is not supported on VSR.

Enhanced subscriber management (ESM) is prerequisite knowledge. This chapter can serve as an SRv6 primer for subscriber management users who are less familiar with this technology.

Overview

Segment routing over IPv6 dataplane (SRv6) is a technology that supports overlay network designs over a single dataplane protocol (IPv6). The size of an IPv6 address (128 bits) is large enough to carry more information than only the addressing of interfaces on a network device. SRv6 allows for additional information so that an IPv6 address can encode functions that extend beyond simple node reachability information. Mature routing protocols (IGP and BGP) with proven fast rerouting mechanisms continue to be used to disseminate SRv6 information across the networks. The mix of those fast-rerouting mechanisms, SRv6, and synchronized subscribers in dual-homed environment enhanced subscriber management with multi-chassis synchronization (ESM MCS) results in simplified and robust access networks that benefit service providers and their customers.

This chapter describes an example topology with:

- redundant BNGs
- synchronized subscribers
- an SRv6-based network providing access to the subscribers
- routing on the core side where the cost of advertised subscriber routes depends on the forwarding state —specifically, the subscriber routed redundancy protocol (SRRP) state—in the access network

Note: Although SRv6 is the transport technology described in this chapter, SRv6 is not a prerequisite for active/standby BNG with EVPN VPWS in the access network. Multi-chassis BNG redundancy can be realized with a variety of other transport technologies in the access network.

This chapter provides integrated configuration content for BNG and SRv6.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The figure Figure 332: Example topology shows the example topology used throughout this chapter. The terminology is as follows:

- The access network refers to all nodes downstream of the BNGs, including the interfaces and SAPs on the BNGs connected to that part of the network.
- The core part of the network refers to all nodes upstream of the BNGs, including the interfaces and SAPs on the BNGs connecting them to that part of the network.
- The term "EVPN VPWS" is used interchangeably with "Epipe" because the EVPN VPWS deployed in the access network is configured in SR OS nodes as an Epipe.

Note: Although terms such as "active or standby subscriber" or "active or standby node" are commonly used in redundant BNG topologies and may occasionally be used within this document, it is important to state that subscribers and nodes do not have redundancy states, such as "active" or "standby". Only SRRP has states that are here referred to as "active" and "standby".

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 332: Example topology

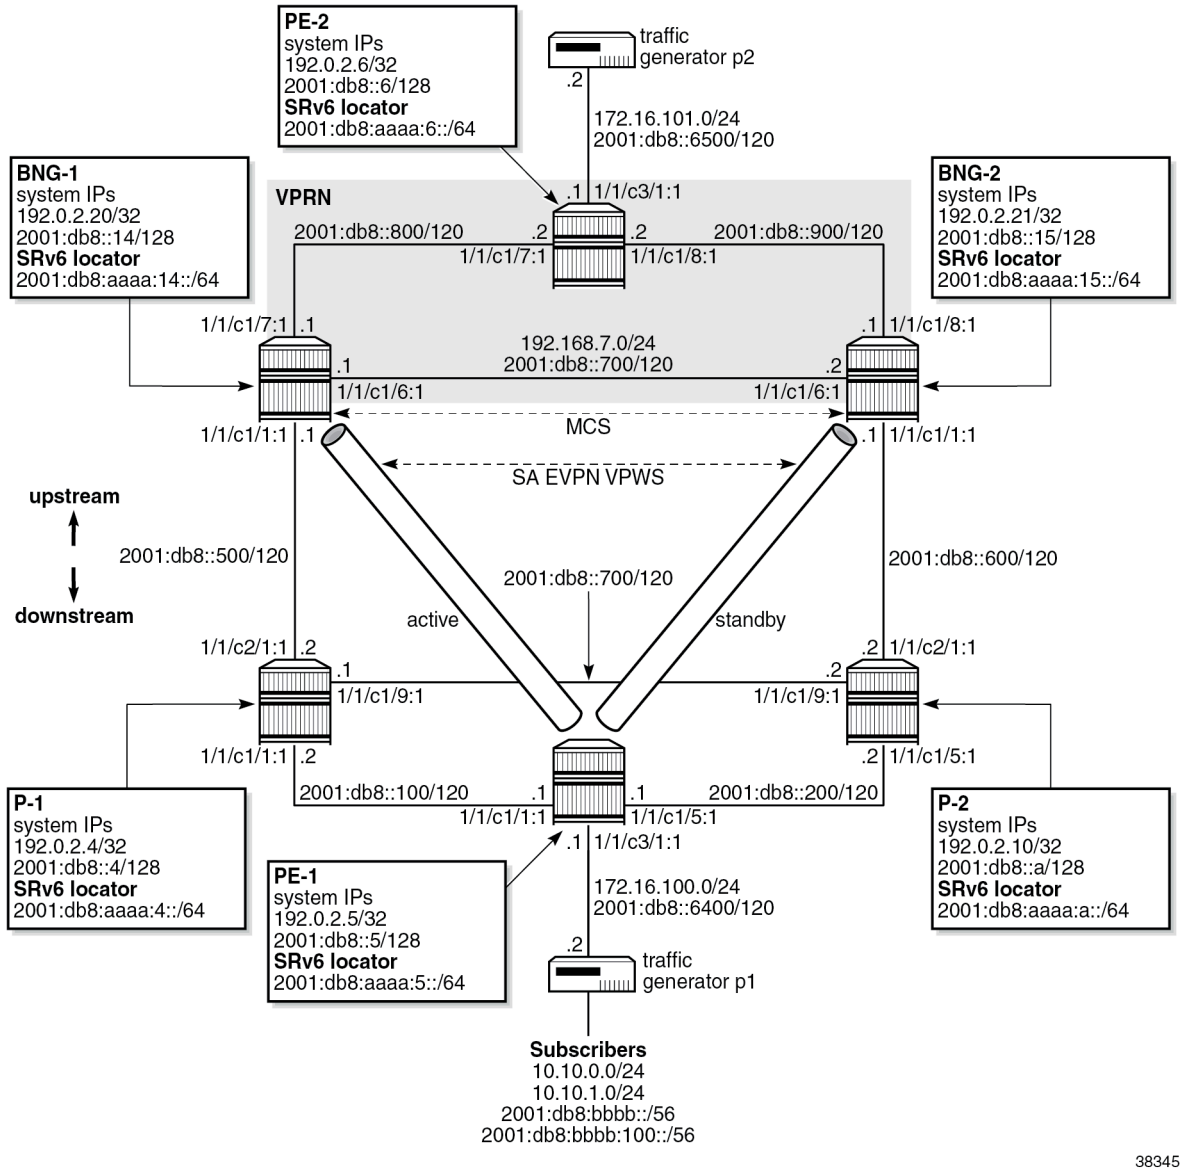

At a high level, the topology in the figure Figure 332: Example topology can be described in terms of access-related configuration, subscriber-related configuration, and core-related configuration.

Access-related configuration

- Subscribers are connected to the first-mile access node (PE-1), which is dual-homed to the two BNGs (BNG-1 and BNG-2).
- Pseudowire (PW) ports on the two BNGs are part of an Ethernet segment (ES) configured in single-active (SA) mode in the EVPN VPWS multihoming (MH) environment.
- The EVPN transport technology in the access network is SRv6.

Subscriber-related configuration

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- There are 10 IPoE and 10 PPPoE dual-stack sessions terminated on a PW-SAP in a VPRN on both BNGs.
- Each IPoE and PPPoE session has an IPv4 address, a DHCPv6 identity association for non-temporary addresses (IA-NA), and a DHCPv4 identity association for prefix delegation (IA-PD) addresses. IA-PD addresses are modeled as managed routes with IPv6 next hops.
- Each session is mapped to a separate subscriber.
- Subscriber sessions are synchronized between the two BNGs. A fully synchronized model offers redundancy protection with minimal loss during network outages.
- DHCPv4/v6 servers are instantiated on the BNGs and their pools are synchronized, except for DHCPv4 pools used for internal address assignment to PPPoEv4 sessions. IPv4 addresses for PPPoEv4 sessions are synchronized in DHCPv4 pools through PPPoE session synchronization and not directly through DHCP pools.
- The SRRP state is indirectly derived from the state of the dual-homed EVPN ES in the access network and does not rely on exchange of SRRP keepalive messages between the two BNGs.
- A redundant interface between the two BNGs provides a temporary path for subscriber traffic while the network is converging during switchovers.

Core-related configuration

- A VPRN with SRv6 transport is deployed in the core.
- SRRP-aware routing on the core side ensures that the core traffic follows the BNG with the active SRRP instance. In other words, traffic in both directions (upstream and downstream) is attracted to one BNG or the other based on the SRRP state.

The example topology provides protection against failures of network elements in the access and core networks, including the BNGs themselves.

Subscriber traffic is run between two ports of a traffic generator. For topology verification purposes, network failures are purposefully introduced and their effect on the subscriber traffic is examined.

Subscriber synchronization

The subscriber state is synchronized between the two BNGs, including DHCPv4/v6 lease state and PPPoEv4/v6 states.

DHCPv4/v6 pools on the local DHCP servers in both BNGs are synchronized in access-driven mode (except for the pools used for internal IP address allocation to PPPoEv4, which are synchronized indirectly through PPPoE session synchronization). In access driven mode, the DHCPv4/v6 servers in each BNG are configured identically and are attached to an interface with the same IP address. Both servers can allocate IP addresses at any time; however, to prevent duplication of IP addresses, access to only one server is allowed at any given time. This restriction is achieved through the redundancy model, which relies on the SA mode of operation on the ES coupled with SRRP.

The figure Figure 333: DHCP pool synchronization in access-driven mode shows a topology using the redundancy model, where subscribers have access only to BNG-1 via EVPN VPWS multihoming. The path of the EVPN VPWS on the left side is active while the path on the right side is standby.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 333: DHCP pool synchronization in access-driven mode

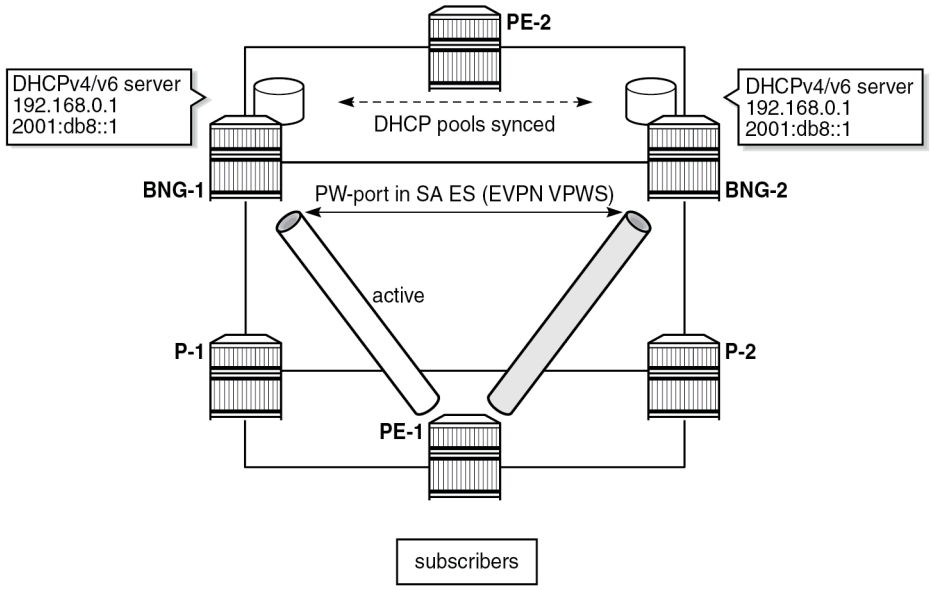

SRv6

This section provides a brief introduction to SRv6. For more information about SRv6, refer to the "Segment Routing over IPv6" chapter in the Segment Routing and PCE volume in the *7450 ESS, 7750 SR, and 7950 XRS MD-CLI Advanced Configuration Guide - Part I* and EVPN VPWS Services with SRv6 Transport chapter and to the *7750 SR and 7950 XRS Segment Routing and PCE User Guide*.

SRv6 allows for more flexible network management over a single IPv6 dataplane protocol, using the following concepts:

- traffic engineering, where traffic is securely steered through pre-selected network nodes on its way to the destination
- network programmability, which includes encoding higher-level functions (such as service identification) into the IPv6 address

At the dataplane level, SRv6-based traffic steering is implemented with the help of a routing extension header, called a segment routing header (SRH). Each application typically provides its own type of the routing header, as is the case here, with inherent support for security.

The SRH contains a list of segments, each of which has a unique segment identifier (SID), through which a packet is steered as it transits the network on its way to the destination. These segments may be nodes, adjacencies, bindings to downstream tunnels, and so on. Each node that owns a SID in the SRH needs to intercept the SRv6 packet for inspection and consequently update the IPv6 header. Transit nodes that are not in the SRH list do not intercept SRv6 packets because the packet is not addressed to them. Such transit nodes forward the packet to the destination based on a lookup of the destination address (DA) in the outer IPv6 header.

As well as identifying specific segments via IPv6 addresses in SRv6, SIDs carry additional information beyond node reachability. This additional information instructs the target node how to map arriving

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

IPv6 packets to a service. For example, a SID can be an IPv6 address that, in addition to carrying the address of the target node (SRv6 prefix or a locator), also carries specific EVPN VPWS information. This additional information is interpreted at the target node as “extract the payload of this packet and process it in the context of EVPN Epipe ID=10”. In this way, an arriving IPv6 packet is directly mapped into the corresponding service based on the IPv6 address, without the need for additional transport protocols such as MPLS with its VC labels.

The figure Figure 334: SRv6 SID shows the SID structure, where the length of the fields is configurable.

Figure 334: SRv6 SID

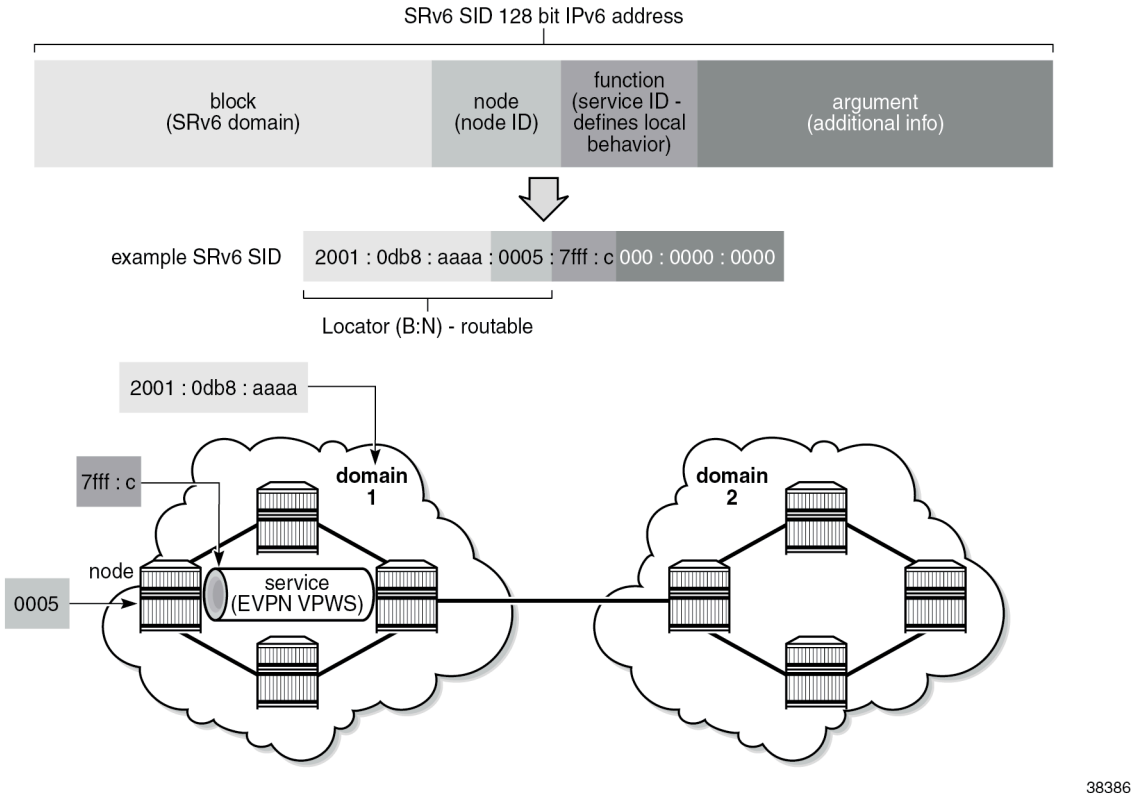

To differentiate between SIDs used in plain routing and the higher level SIDs that are used with services (such as EVPN and VPRN), the SIDs are sometimes referred to as transport SIDs and service SIDs. SIDs are advertised in routing protocols through segment routing extensions. Nokia SR OS supports such extensions in IS-IS for transport-related SIDs and in BGP for Layer 2 and Layer 3 service-related SIDs.

SIDs are an integral part of routing in SRv6, but SIDs are not interface addresses; SIDs are configured and allocated in respective **segment-routing-v6** contexts (outside of the interface configuration) to the Base router and service to which they pertain.

As the figure Figure 334: SRv6 SID shows, the block and the node form a locator field in the SID. The locator field is used for node reachability in basic destination-based routing. Typically, a node advertises IPv6 prefixes matching the locator function of its SIDs. In SR OS, when the locator prefix is advertised in IS-IS, it is encoded in an SRv6 Locator Sub-TLV, which is an SRv6 extension in IS-IS. A node can advertise multiple locators simultaneously (for example, one locator per Flexible-Algorithm).

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The block part of the locator defines an SRv6 domain. In an SRv6 domain, nodes share the same block prefix. SRv6 domains can be considered as administrative units that can be formed based on geography or some other logical entity, such as enterprise VPN. The node part of the locator is unique to each node.

The function part of the locator defines a local behavior on the node that owns the SID. Specific functions that can be associated with SIDs are described in RFC 8986 and RFC 8402. The End and End.X functions are used by IS-IS to create repair tunnels and backup paths (topology-independent loop-free-alternate (TI-LFA) and remote LFA). The combination of TI-LFA and remote LFA provides full coverage of any access network. The End function represents the node prefix which is reachable based on the shortest path. The End.X function represents links between the nodes (these links are router adjacencies which can be local or remote) and is used to specify the router adjacency (for example, for strict mode routing) out of which the frame is forwarded. Together, the End and End.X functions are used to install optimal and loop-free backup paths in the forwarding plane with achievable failover times of less than 50 ms.

The table Table 19: SRv6 endpoint behaviors shows the functions used in this chapter:

Table 19: SRv6 endpoint behaviors

Function	Description	Advertised in
End	SRv6 instantiation of a node SID	IS-IS SRv6 End SID sub-TLV
End.X	SRv6 instantiation of an adjacency SID	IS-IS SRv6 End.X SID sub-TLV
End.DX2	decapsulation and L2 cross-connect (L2VPN)	EVPN VPWS AD per-EVI route
End.DT4	decapsulation and specific IPv4 table lookup (IPv4-L3VPN)	BGP
End.DT6	decapsulation and specific IPv6 table lookup (IPv6-L3VPN)	BGP
End.B6.Encap.Red	function bound to an SRv6-policy with reduced encapsulation	N/A

The argument field is optional, and it can carry additional information related to the local function. It is set to 0 inSR OS.

Routing

IS-IS is used as IGP to disseminate node reachability information (the node's interface routes, BGP next hops, SRv6 locators, and node and adjacency SIDs), while BGP is used to disseminate SRv6 service SIDs and VPN and EVPN routes.

As shown in the figure Figure 335: Separation of routing domains, IS-IS is segmented into two instances:

- IS-IS instance 0 runs in the access network
- IS-IS instance 1 runs in the core part of the network, including the inter-chassis link (ICL) between the two BNGs

The system IPv6 addresses are included in both instances. There is no route exchange between the two IS-IS instances.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

IPv6 BGP peering is established directly (with no route reflector) between the service endpoints with respective BGP address families, as shown in the figure Figure 335: Separation of routing domains. IPv6 BGP peering is not a prerequisite and IPv4 BGP peering would work just as well.

Figure 335: Separation of routing domains

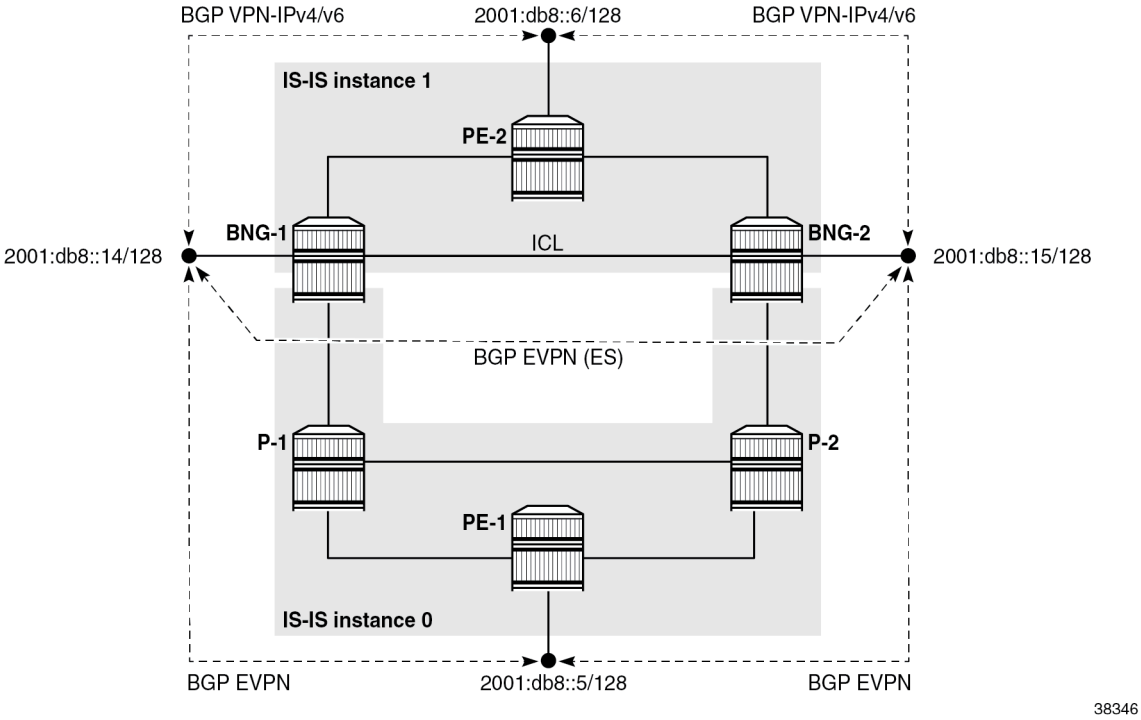

The ICL must be highly redundant but does not need to support high bandwidth. Its use is typically limited to subscriber synchronization and to shunt data traffic between BNGs during transient network conditions while the network reconverges during switchovers. For cost reductions, it may not be desirable to dimension this ICL to carry all subscriber traffic during prolonged periods of network failures. To minimize the use of ICL for data traffic, the topology shown in the figure Figure 335: Separation of routing domains relies on the separation of IS-IS routing domains between the access network and the core network, as described in the Examining failovers section. The hard separation of IS-IS routing domains is not a prerequisite or the only design option, but rather an arbitrary choice used in this chapter.

When the link between the nodes BNG-1 (with the active SRRP) and P-1 fails, the PE-1 BGP next hop becomes unreachable on BNG-1 because of the separation of IS-IS routing domains, which causes the active EVPN VPWS between BNG-1 and PE-1 to switch over to BNG-2. The switch to BNG-2, combined with SRRP-aware routing on the core side, ensures that the use of ICL is limited to transient conditions while routing is converging during the failure. The figure Figure 336: Optimal traffic flow during network failures shows the steady state traffic path after the network has converged: the path from PE-1 to PE-2 via P-2 and BNG-2.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 336: Optimal traffic flow during network failures

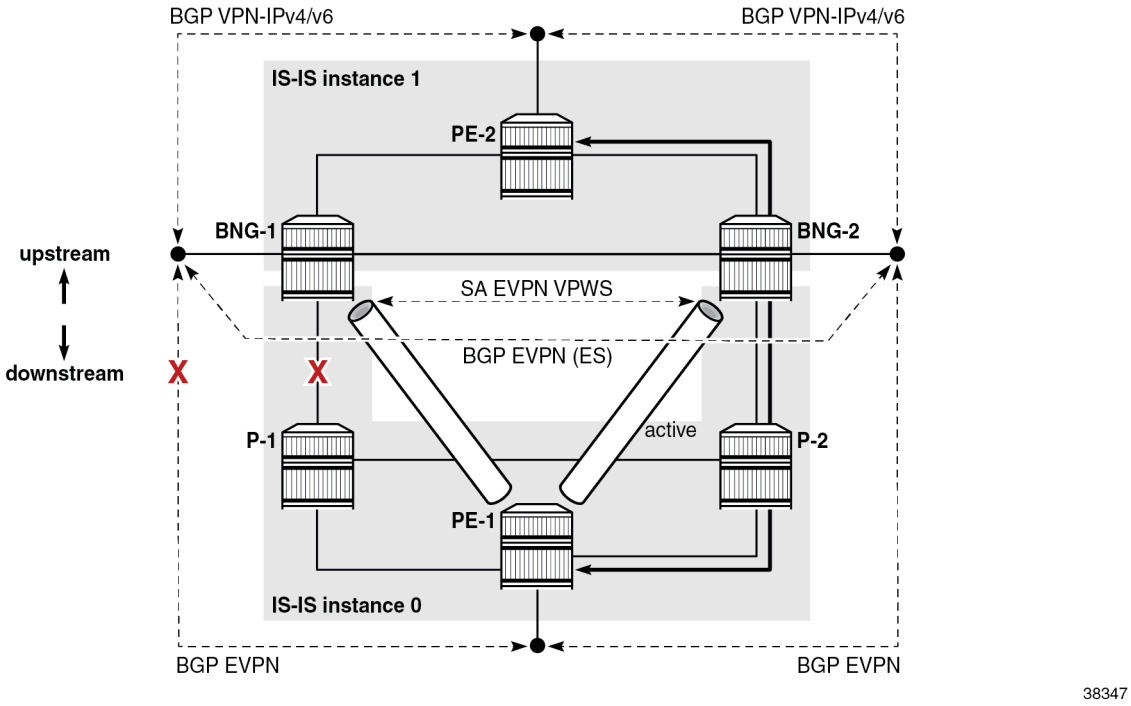

In an alternative design where all nodes are in the same IS-IS routing instance 0, the PE-1 node continues to be reachable from BNG-1 (via BNG-2 or PE-2) even after the access link on BNG-1 fails. As a result, BNG-1 remains the designated forwarder (DF) for the ES and SRRP on BNG-1 remains in the active state. This causes traffic in both directions to take a path over the ICL link or even to cross some of the core links twice, as shown in Figure 337: Suboptimal traffic flow during network failure . Traffic from PE-1 goes via P-2 to BNG-2, which sends the traffic to DF BNG-1 via PE-2 (the ICL is not used). BNG-1 forwards the traffic (back to) PE-2.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 337: Suboptimal traffic flow during network failure

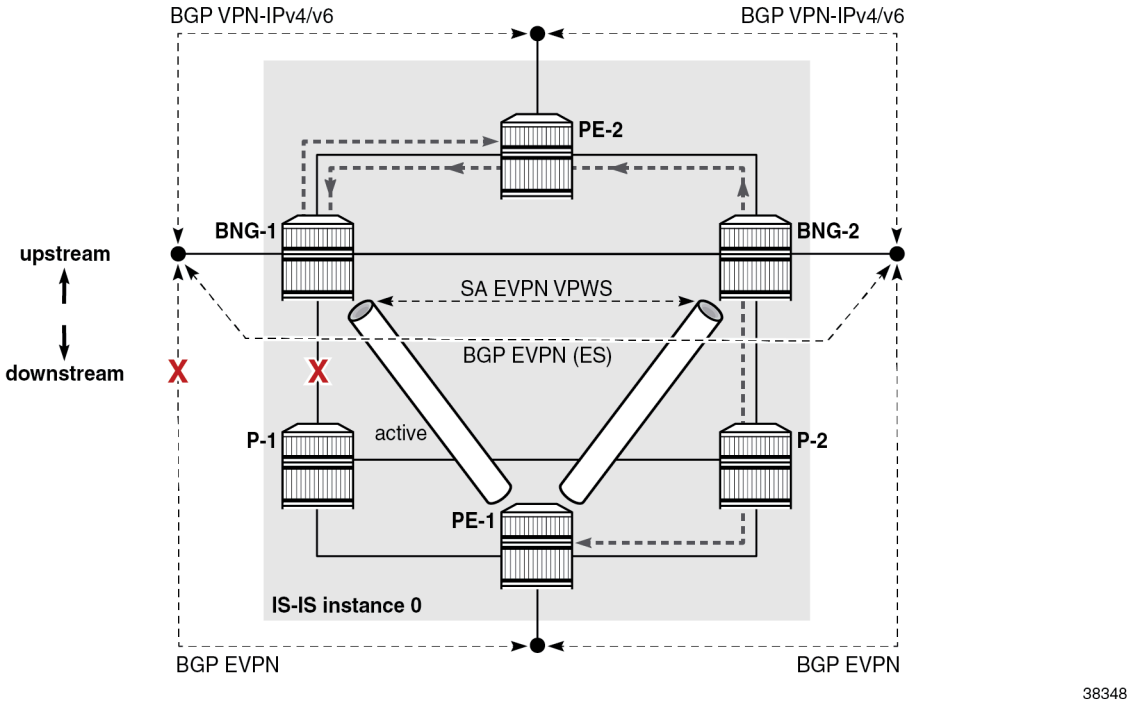

IPv4 addressing

Subscriber traffic in the overlay network utilizes both IPv4 and IPv6 address families, as part of dual-stack configuration. Comparatively, the underlay network is based on SRv6, with the advantage that IPv4 addressing is no longer needed.

Therefore, in the example topology, IPv4 addresses in the underlay network are removed from most interfaces. There are a few exceptions where local IPv4 addresses are kept, but without advertising them into the network, which means that routing protocols in the underlay network remain free of IPv4 addresses.

The following are exceptions where IPv4 addresses in the underlay network are still used:

- System IPv4 addresses are still used locally on each node for legacy reasons where some functionalities may still require it. For example, in EVPN VPWS multihoming, an IPv4 system address is required to derive the route distinguisher (RD) for the ES route (even if static RD for the EVPN VPWS is configured).
- Multi-chassis synchronization (MCS) peering between the two nodes is supported only over IPv4 addresses. For the MCS peering, the system IP addresses of the BNGs are used, with static routes over the IPv4 next hops configured on the direct link between the two BNG nodes. Because this link directly connects the two BNGs, no advertisement of the IPv4 prefix is required.
- The **redundant-Interface** command under the BNG **group-interface** context also requires IPv4 addressing. Those IPv4 addresses are also not advertised into the network.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

SRv6 policy

For illustrative purposes, the example uses a static SR policy of type SRv6 in the access network. This SR policy is configured with segment list entries, which are node SIDs that are to be visited on the way to the final destination.

The SR policy is treated as a routing entity whose configured endpoint is installed as an entry in the tunnel table. The SR policy becomes the protocol owner of such an entry in the routing or tunnel table. The tunnel entry can be used for BGP next-hop resolution.

Configuring hops in the SR policy triggers the generation of an SRH, which is further described in the Configuration section.

The SR policy in the example steers traffic onto a longer path, despite a shorter path with lower cost being available. This behavior is shown in the figure Figure 338: Static SRv6 policy paths where traffic from the BNGs to PE-1, and vice versa, always flows through both P-1 and P-2.

Figure 338: Static SRv6 policy paths

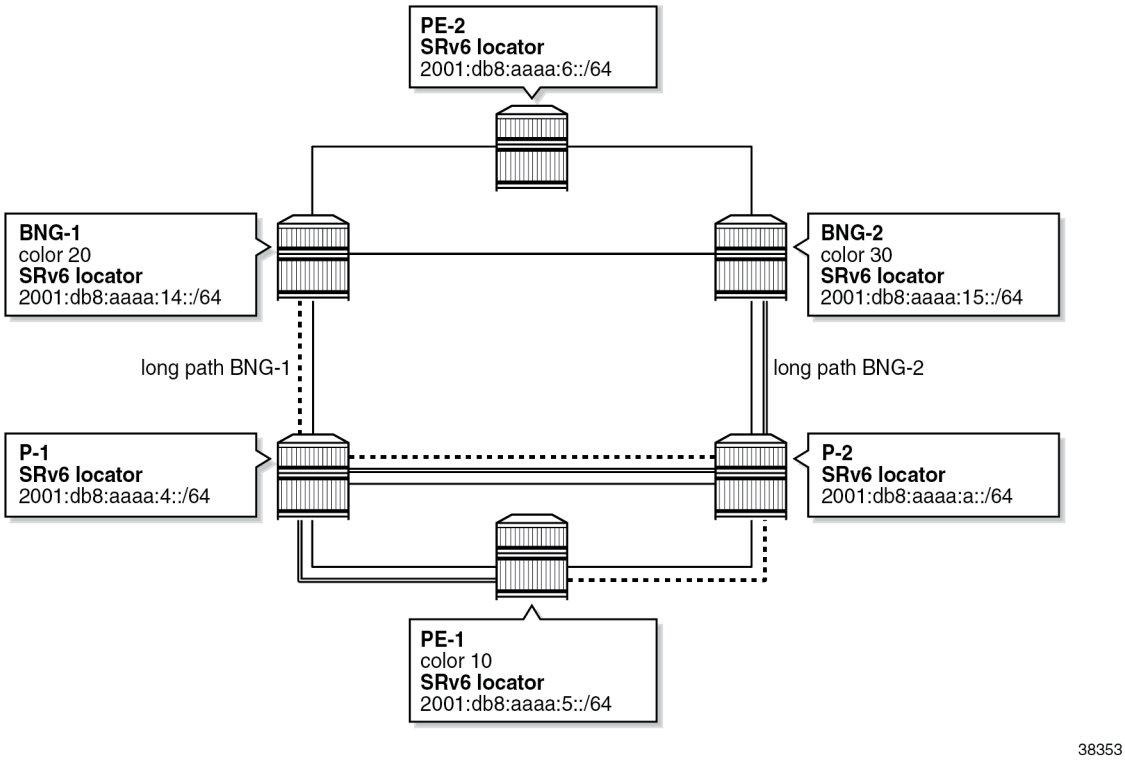

Examining failovers

The configuration used in this chapter has been tested by running traffic and incurring network failures at various points. Traffic was monitored after the failures to ensure the traffic was flowing through a predicted path.

Traffic is run bidirectionally between two ports on a traffic generator, representing the subscribers on the 10.10.*x*.*x* networks and the Internet on the 172.16.101.*x* network.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

BNG-1 advertises subscriber IPv4 and IPv6 subnets into the core network (VPRN) with lower cost than BNG-2, because SRRP is active on BNG-1 and standby on BNG-2.

In the access network, the ES in EVPN VPWS on BNG-1 is selected as DF due to a higher preference.

Traffic through the access network is steered by the SRH, which is inserted by the static SR policy configured on both BNGs and PE-1. In the upstream direction, traffic follows the path depicted in 8: from traffic generator port 1 via PE-1, P-2, P-1, BNG-1, and PE-2 to traffic generator port 2. Downstream traffic follows the same path in reverse direction.

Figure 339: Baseline traffic flow

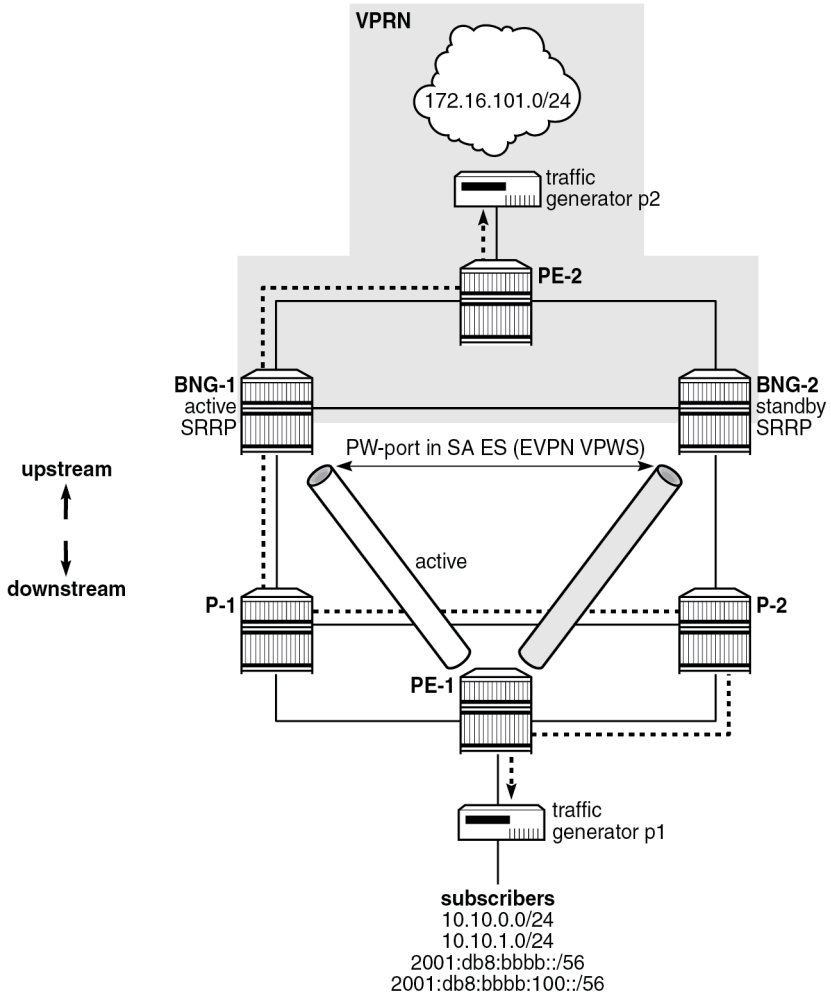

Three switchovers are examined:

- switchover due to failure of the link between BNG-1 and P-1
- switchover due to failure of the network link connecting BNG-1 to PE-2
- switchover due to failure of the entire BNG-1 node

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

In all three cases, the failover time depends on the failure detection speed and convergence of the routing protocols. Nokia encourages operators to run performance-related tests, which includes measuring failover times, in their own test environment on service routers with Nokia native hardware (non-server based service routers).

Failure detection speed can, in some cases, be improved by bi-directional forwarding detection (BFD). In the example topology, BFD is configured for detection of BGP peer failures.

Switchover caused by the failure of the access link on BNG-1

The figure Figure 340: Access port failure shows a switchover scenario where the access port on BNG-1 toward P-1 is disabled and all traffic is diverted to BNG-2 as indicated. The traffic detour over the link between P-1 and P-2 is taken because the static SR policies favor longer paths over shorter.

Figure 340: Access port failure

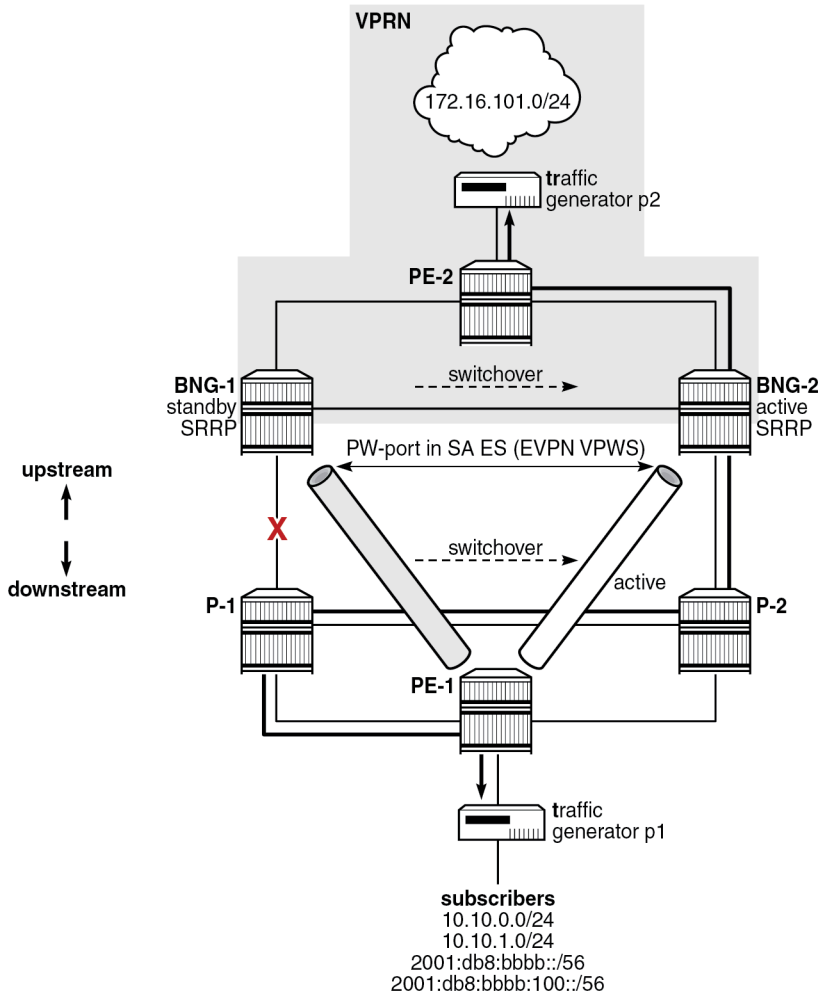

The following events trigger the traffic switchover from BNG-1 to BNG-2:

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- When the access port on BNG-1 is disabled, the BGP peering connection between BNG-1 and PE-1 is lost because of the separation of IS-IS routing domains. The BGP connection is required for the advertisement of the overlay EVPN routes.
- On BNG-1, the speed of detecting the BGP connection failure is locally driven by the speeds of the link loss detection and the consequent IS-IS algorithm that is run to determine that an alternate path to the BGP peer (PE-1) does not exist.
- On PE-1, the speed of detecting the BGP connection failure is driven by the speed of IS-IS link state advertisement (LSA) propagation.
- The severed BGP connection between BNG-1 and PE-1 triggers the DF election process on the ES between the two BNGs, which is driven by BGP. BNG-2 is selected as the DF and BNG-1 as the non-DF.
- The newly elected DF on BNG-2 performs the following actions:
  - It activates EVPN VPWS toward PE-1 so that traffic can be sent to it and be received from it. (More explicitly, it advertises an Auto-Discovery route with the Primary flag (P-bit) set to 1 to indicate that it is active, and the Backup flag (B-bit) set to zero.)
  - It triggers activation of SRRP on BNG-2 (but does not activate it). SRRP continues to run through its own state machine, which involves waiting for three times the SRRP transmit interval (SRRP keepalive messages) before it becomes active. Due to this delay in becoming active, it is still important to have a short SRRP transmit interval configured under SRRP, even if SRRP keepalive messages are not exchanged between the two BNGs, as in this example. In other words, SRRP activation is driven by the DF outcome of the EVPN ES election process, which then triggers the SRRP activation process with its own state machine.
- PE-1 does not wait for the BGP update from BNG-2 that contains the new DF decision. PE-1 activates the stored backup route from BNG-2 as soon as it realizes, via IS-IS, that the BGP connection to BNG-1 is lost.
- Lastly, the activation of SRRP on BNG-2 triggers the advertisement of subscriber routes into the core with a lower cost than that advertised by BNG-1. After the core network converges, traffic in the downstream direction is diverted from the link between PE-2 and BNG-1 to the link between PE-2 and BNG-2. While the routing is converging, traffic in the downstream direction might be flowing from BNG-1 to BNG-2 through the redundant interface (over the ICL) between the two BNGs.

Switchover caused by the network link failure between BNG-1 and PE-2

In the figure Figure 341: Network port failure, the port on BNG-1 toward PE-2 is disabled, which results in a switchover due to the network link failure between BNG-1 and PE-2. All traffic is diverted to BNG-2. The traffic detour over the link between P-1 and P-2 is taken because of the static SR policies on BNGs and PE-1.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 341: Network port failure

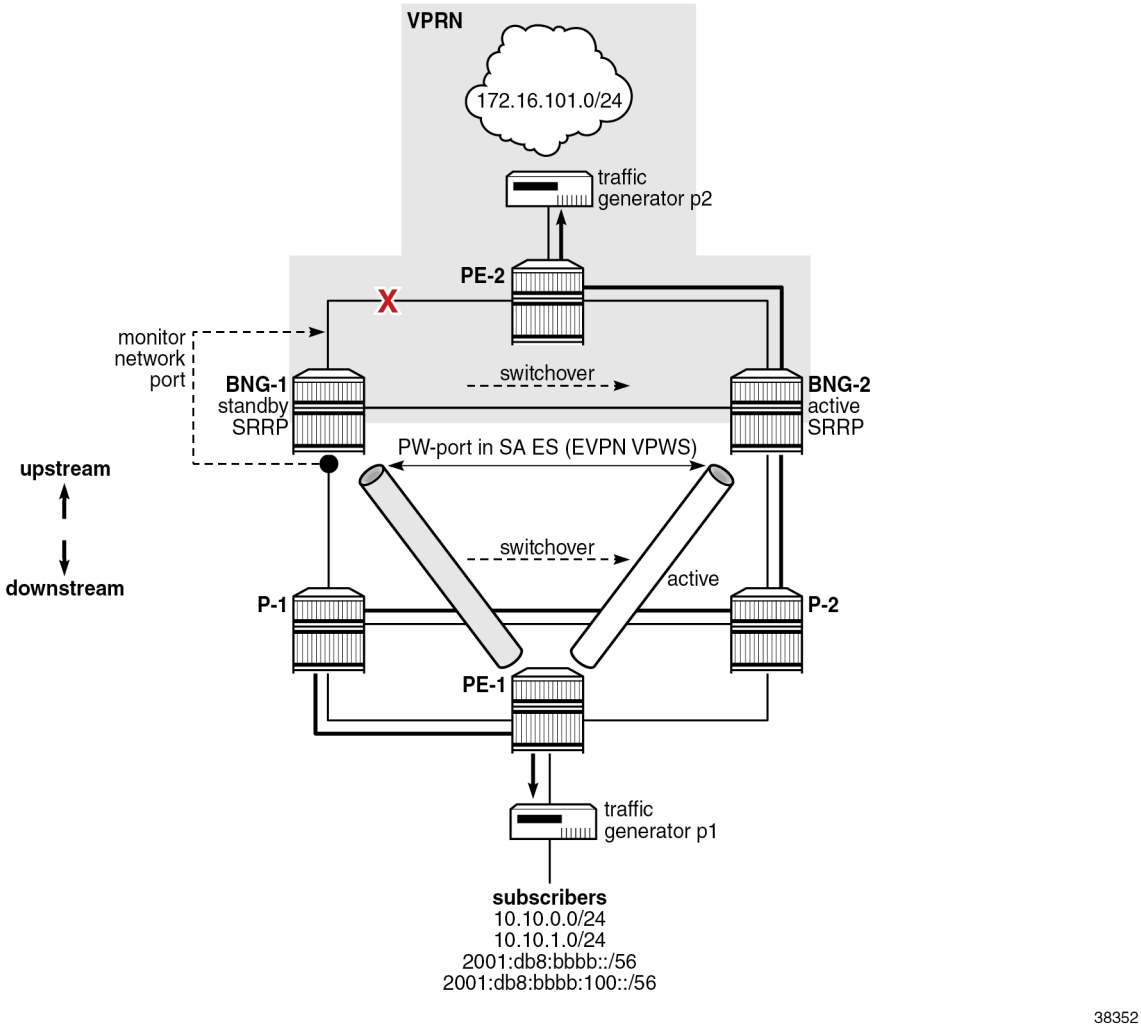

A failure of the core link on BNG-1 diverts downstream traffic to BNG-2 in the core network, which uses VPRN. But in the access network, which uses EVPN, traffic still flows through BNG-1, which remains the DF on EVPN ES and where SRRP is active. To prevent a scenario where the inter-chassis link is used to carry traffic between the two BNGs during the failure, a feedback loop from the core port to the access port is required. In this way, the state change of the core port is reflected in the access port and, as result, traffic is completely diverted from BNG-1 to BNG-2. This feedback loop is implemented through an operational group associated with a core port and monitored by the access port.

The remainder of the switchover mechanics are the same as those in the Switchover caused by the failure of the access link on BNG-1 section.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Switchover caused by the BNG-1 node failure

After the loss of the BGP peering connection to BNG-1, BNG-2 becomes the DF for the EVPN ES. The BGP peering connection loss can be detected by a loss of link on BNG-2, by BFD, or by IS-IS LSA propagation.

BNG-2 becoming DF has the same effects as already described in the previous two cases.

Recovery

To restore the network to its state from before the failure, the operators have the option to revert traffic automatically to the primary paths that were available before the failure. The revertive behavior involves another switchover, which is associated with possible traffic loss. Whether to enable revertive behavior is a choice between the cost of another switchover measured by affected user experience due to a small traffic loss and the cost of keeping the traffic on backup paths which may not be optimized for such traffic.

To ensure minimal traffic loss when revertive behavior is enabled, the network must have fully converged after the recovery, before the switchover to the previous state occurs. For example, upon reboot, the recovered BNG must be fully integrated into the network; that is, all subscribers must be synchronized between the two BNGs, the BGP peering sessions must be reestablished, and routes exchanged. The synchronization between the BNGs is independent of the exchange of BGP routes and these events may occur in any sequence. For this reason, bootup timers must be configured properly to make sure that all relevant processes are completed before the switchover occurs. Bootup timers are described in the Configuration section.

Configuration

The complete configuration files for all six nodes in this example topology are provided in the Appendix.

The examples in this section focus on the Forwarding Path Extension (FPE)-based PW port in multihomed EVPN environments with SRv6 transport.

The configuration shown in this section is mainly for the two nodes that represent the two endpoints of the EVPN VPWS: BNG-1 and PE-1.

The configuration blocks, together with the output of the **show** commands, cover the following topics:

- FPE
- SRv6
- IS-IS routing related to SRv6
- BGP routing related to SRv6
- EVPN VPWS (including ES and PW port) and service-related SRv6 configuration
- VPRN-related SRv6 configuration
- multi-chassis redundancy configuration
- DHCP server configuration
- subscriber management configuration

While many of these topics are described in other chapters, the purpose of this chapter is to show how they interact together.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The less important parts in the configuration blocks or the output from the **show** commands are removed for brevity.

FPE

In SR OS Release 22.10, the example setup requires three FPE-based port cross-connects (PXCs) on each BNG, one for each of:

- the SRv6 origin
- the SRv6 termination
- the PW port

All three PXCs are port-based and they are set up on the same port. Selecting port-based PXCs in this example versus MAC- or internal-based PXCs is an arbitrary choice. Either type of PXC can be used.

The PXC configuration on BNG-1 is as follows. The configuration on BNG-2 is identical. Other nodes do not have a PW port and require only two PXCs for SRv6.

```
[pr:/configure port-xc]
A:admin@bng-1#
    pxc 1 {
        admin-state enable
        description "fpe srv6 origin"
        port-id 1/1/c4/1
    }
    pxc 2 {
        admin-state enable
        description "fpe srv6 termination"
        port-id 1/1/c4/1
    }
    pxc 3 {
        admin-state enable
        description "fpe pw-port"
        port-id 1/1/c4/1
    }
```

The logical PXC ports are configured as follows:

```
[pr:/configure]
A:admin@bng-1#
    port pxc-1.a {
        admin-state enable
    }
    port pxc-1.b {
        admin-state enable
    }
    port pxc-2.a {
        admin-state enable
    }
    port pxc-2.b {
        admin-state enable
    }
    port pxc-3.a {
        admin-state enable
    }
    port pxc-3.b {
        admin-state enable
    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
}
```

The following FPE configuration is required for PW port and SRv6 configuration:

```
[pr:/configure]
A:admin@bng-1#
    fwd-path-ext {
        sdp-id-range {
            start 17500
            end 17600
        }
        fpe 1 {
            description "srv6 origination"
            path {
                pxc 1
            }
            application {
                srv6 {
                    type origination
                }
            }
        }
        fpe 2 {
            description "srv6 termination"
            path {
                pxc 2
            }
            application {
                srv6 {
                    type termination
                }
            }
        }
        fpe 3 {
            description "pw-port on a single port"
            path {
                pxc 3
            }
            application {
                pw-port-extension { }
            }
        }
```

SRv6

SRv6 is configured on several levels:

- the base routing context level (which is the focus of this section)
- the IS-IS routing protocol level
- the services level

The static SRv6 policy is configured in the base routing context and is examined in the context of the Epipe to which it is applied.

Relevant to the example setup, the global SRv6 configuration under the base router provides:

- an SRv6 association with an FPE
- an SRv6 locator

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- the source address
- the allocation and behavior of End and End.X SIDs

Note: Nokia encourages that the source address is configured at the Base routing context level. Although some entities, such as EVPN or VPRN services, provide the configuration option for the source address, other entities, such as the static SR policy, do not. In these cases, the entities inherit the source address from the global configuration provided here. Without the configured source address, SRv6 is not operational.

```
[pr:/configure router "Base" segment-routing segment-routing-v6]
A:admin@bng-1#
    origination-fpe [1]
    source-address 2001:db8::14
    locator "bng-1-loc" {
        admin-state enable
        block-length 48
        termination-fpe [2]
        prefix {
            ip-prefix 2001:db8:aaaa:14::/64
        }
    }
    base-routing-instance {
        locator "bng-1-loc" {
            function {
                end 1 {
                    srh-mode usp
                }
                end-x-auto-allocate usp protection protected { }
            }
        }
    }
```

The following **show** command displays a summary of SRv6 on a global level:

```
A:admin@bng-1# /show router segment-routing-v6 summary
===============================================================================
Segment Routing v6
===============================================================================
Origination FPE              : 1
Source IPv6 Address          : 2001:db8::14
===============================================================================
Locator                                                          Admin State
    Prefix
-------------------------------------------------------------------------------
bng-1-loc                                                        Up
    2001:db8:aaaa:14::/64
---snip---
```

The following **show** output lists the node and adjacency SIDs:

```
A:admin@bng-1# /show router segment-routing-v6 base-routing-instance
===============================================================================
Segment Routing v6 Base Routing Instance
===============================================================================
Locator
  Type       Function    SID                                     Status/InstId
    SRH-mode Protection  Interface
-------------------------------------------------------------------------------
bng-1-loc
  End                  1 2001:db8:aaaa:14:0:1000::               ok
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    USP
-------------------------------------------------------------------------------
  Auto-allocated End.X: USP Protected,
-------------------------------------------------------------------------------
  End.X          *524287 2001:db8:aaaa:14:7fff:f000::            1
    USP        Protected int-1-bng-1-pe-2
    ISIS Level: L2 Mac Address: b4:a0:01:01:00:07 Nbr Sys Id: 1920.0000.2006
  End.X          *524288 2001:db8:aaaa:14:8000::                 0
    USP        Protected int-1-bng-1-p-1
    ISIS Level: L2 Mac Address: b4:9e:01:01:00:0b Nbr Sys Id: 1920.0000.2004
  End.X          *524289 2001:db8:aaaa:14:8000:1000::            1
    USP        Protected int-1-bng-1-bng-2
    ISIS Level: L2 Mac Address: b4:98:01:01:00:06 Nbr Sys Id: 1920.0000.2021
-------------------------------------------------------------------------------
Legend: * - System allocated
```

The InstId field for End.X SIDs shows the IS-IS instance in which they are configured.

SRv6 locators are used for node reachability and are, through route advertisements, installed in the route tables of other nodes. For example, the following output shows the SRv6 locator of BNG-1 is installed in the IPv6 routing table of BNG-2:

```
A:admin@bng-2# /show  router route-table ipv6

===============================================================================
IPv6 Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
---snip---
2001:db8:aaaa:14::/64                         Remote  ISIS(1)   02h14m11s  18
       2001:db8:aaaa:14::/64 (tunneled:SRV6-ISIS)                  20
---snip---
```

IS-IS routing related to SRv6

As described in the Routing section, IS-IS is split into two instances. Instance 0, as indicated by its configured interfaces, is active on the access side. Instance 1 is active on the core side, including the link between the two BNGs.

The **segment-routing-v6** section in the IS-IS configuration block references the SRv6 locator that IS-IS, together with its basic functions (node and adjacency SIDs), advertises into the network.

Some of the IS-IS configuration outside of the **isis segment-routing-v6** context is directly pertinent to SRv6. For example, the SRv6 locator is only advertised when the **wide-metrics** command is configured.

BNG-2 is configured similarly to BNG-1. P-1, P-2, and PE-1 have only IS-IS instance 0 configured, while PE-2 has only IS-IS instance 1 configured. Although only a portion of the configuration on BNG-1 is provided, the IS-IS configuration on other nodes in respective IS-IS instances is similar with different naming for interfaces and for the SRv6 locator.

Between the two IS-IS instances in BNG-1, the only configuration difference is that they reference different interfaces. In the following example, only one repeated part of the IS-IS instance 1 configuration is shown:

```
[pr:/configure router "Base"]
    isis 0 {
        admin-state enable
        advertise-passive-only false
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        advertise-router-capability area
        ipv6-routing native
        level-capability 2
        traffic-engineering true
        area-address [49.0001]
        loopfree-alternate {
            remote-lfa {
            }
            ti-lfa {
            }
        }
        traffic-engineering-options {
            ipv6 true
            application-link-attributes {
            }
        }
        segment-routing-v6 {
            admin-state enable
            locator "bng-2-loc" {
                level-capability 2
                level 2 {
                    metric 10
                }
            }
        }
        interface "int-1-bng-2-p-2" {
        }
        interface "system" {
        }
        level 2 {
            wide-metrics-only true
        }
    }
    isis 1 {
        ---snip---
        interface "int-1-bng-2-bng-1" {
        }
        interface "int-1-bng-2-pe-2" {
        }
        interface "system" {
        }
        ---snip---
```

The following **show** command confirms that IS-IS adjacencies in instance 1 are up. A similar command can be used for IS-IS instance 0.

```
A:admin@bng-1# /show  router isis 1 adjacency

===============================================================================
Rtr Base ISIS Instance 1 Adjacency
===============================================================================
System ID               Usage State Hold Interface                    MT-ID
-------------------------------------------------------------------------------
bng-2                   L2    Up    24   int-1-bng-1-bng-2            0
pe-2                    L2    Up    9    int-1-bng-1-pe-2             0
-------------------------------------------------------------------------------
Adjacencies : 2
```

Further SRv6-related information in IS-IS, such as exchanged locators or node SIDs, can be explored by using commands under the following hierarchy:

```
A:admin@bng-1# /show  router isis 1 segment-routing-v6
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The following **show** command lists the locators that are present in IS-IS instance 1 on BNG-1. These locators can be local or learned through IS-IS. For BNG-1, locator 2001:db8:aaaa:14::/64 is local.

```
A:admin@bng-1# /show  router isis 1 segment-routing-v6 locator

===============================================================================
Rtr Base ISIS Instance 1 SRv6 Locator Table
===============================================================================
Prefix                              AdvRtr                        MT     Lvl/Typ
 AttributeFlags                      Tag                           Flags  Algo
-------------------------------------------------------------------------------
2001:db8:aaaa:6::/64                pe-2                          0      2/Int.
  -                                   0                             -      0
2001:db8:aaaa:14::/64               bng-1                         0      2/Int.
  -                                   0                             -      0
2001:db8:aaaa:15::/64               bng-2                         0      2/Int.
  -                                   0                             -      0
-------------------------------------------------------------------------------
```

BGP peering

IBGP is used to exchange EVPN and VPRN routes between the BNGs and the edge nodes (PE-1 and PE-2). IBGP peering sessions with corresponding address families are shown in the figure Figure 335: Separation of routing domains.

In the EVPN part of the network, BGP is used between the two BNGs to elect the DF on the ES, and between the BNGs and PE-1 to exchange the reachability EVPN routes. PE-1 uses the advertised EVPN route from the BNGs to set up and activate an EVPN VPWS connection toward the active (DF) BNG.

In the VPRN, the two BNGs exchange VPN-IPv4 and VPN-IPv6 routes (subscriber and Internet routes from the traffic generator) with PE-2.

BGP is not configured on nodes P-1 and P-2.

Consider the following when configuring BGP:

- All BGP peers are IPv6 peers.
- The **advertise-ipv6-next-hops** command for the VPN IPv4 and VPN IPv6 address families must be explicitly enabled. By default, only IPv4 next hops are advertised, which in this case do not exist. The counterpart command for EVPN address family is enabled under the **bgp-evpn** configuration for the corresponding Epipe: the **route-next-hop system-ipv6** command is configured in the **configure service epipe <..> bgp-evpn segment-routing-v6** context.
- The **extended-nh-encoding** command must be enabled to allow IPv6 next hops for the VPN IPv4 address family.
- The SR policy used for traffic steering is deployed in the access network.
- BFD is enabled for faster failure detection of all BGP neighbors.

For the SR policy to take effect, the target node referenced in the **endpoint** command must advertise its routes with a color extended community, in the format community:color, that matches the one configured in the policy. The configured color extended community advertised by BNG-1 is color-20. Rather than exporting this color extended community through a VSI-export policy at the service level, for simplicity reasons, it is exported at the global BGP level. This way, the color extended community can be easily added to the existing communities, such as the route target (RT) community in the Epipe.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Note: With the alternative approach relying on the VSI export policy, the RT community would need to be explicitly readvertised along with the color extended community.

The routing policy name used to export the color extended community is "pol-color-20". It is applied only toward the BGP peer PE-1 (neighbor 2001:db8::5). To activate the export policy, the **vpn-apply-export** command must be specifically enabled in MD-CLI.

The BGP configuration on BNG-1 is as follows:

```
[pr:/configure router "Base"]
A:admin@bng-1#
    bgp {
        admin-state enable
        vpn-apply-export true
        vpn-apply-import true
        rapid-withdrawal true
        rapid-update {
            vpn-ipv4 true
            vpn-ipv6 true
            evpn true
        }
        extended-nh-encoding {
            vpn-ipv4 true
        }
        advertise-ipv6-next-hops {
            vpn-ipv6 true
            vpn-ipv4 true
        }
        group "evpn" {
            peer-as 64500
            local-address 2001:db8::14
            bfd-liveness true
            family {
                evpn true
            }
        }
        group "ipvpn" {
            peer-as 64500
            local-address 2001:db8::14
            bfd-liveness true
            family {
                vpn-ipv4 true
                vpn-ipv6 true
            }
        }
        neighbor "2001:db8::5" {
            group "evpn"
            export {
                policy ["pol-color-20"]
            }
        }
        neighbor "2001:db8::6" {
            group "ipvpn"
        }
        neighbor "2001:db8::15" {
            group "evpn"
        }
    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The following shows the export policy "pol-color-20", which adds the community "color-20" to all BGP route advertisements with the family type EVPN originating from the service with the tag 11. For this purpose, the EVPN VPWS (Epipe) is explicitly tagged with tag 11, as shown in the Epipe configuration section.

```
[pr:/configure policy-options]
A:admin@bng-1#
    community "color-20" {
        member "color:00:20" { }
    }
    policy-statement "pol-color-20" {
            entry 10 {
            from {
                family [evpn]
                tag 11
            }
            action {
                action-type accept
                community {
                    add ["color-20"]
                }
            }
        }
    }
```

BFD is enabled on the interface advertised as the next hop in BGP updates:

```
[pr:/configure router "Base"]
A:admin@bng-1#
    autonomous-system 64500
    router-id 192.0.2.20
    interface "system" {
        ipv4 {
            primary {
                address 192.0.2.20
                prefix-length 32
            }
        }
        ipv6 {
            bfd {
                admin-state enable
                transmit-interval 100
                receive 100
                multiplier 2
            }
            address 2001:db8::14 {
                prefix-length 128
            }
        }
    }
```

The BGP configurations on nodes BNG-2, PE-1, and PE-2 are similar in that both PE-1 and PE-2 nodes only interact with the two BNGs within their respective address families.

The following command shows a summary of the BGP status. The output of this command has been shortened to show only that the peers are communicating and exchanging routes. The advertised routes are examined in the service configuration sections.

```
A:admin@bng-1# show router bgp summary
===============================================================================
 BGP Router ID:192.0.2.20       AS:64500       Local AS:64500
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
BGP Admin State         : Up          BGP Oper State              : Up
---snip---
===============================================================================
Neighbor
Description
                   AS PktRcvd InQ  Up/Down   State|Rcv/Act/Sent (Addr Family)
                      PktSent OutQ
-------------------------------------------------------------------------------
2001:db8::5
                64500   11617    0 04d00h46m 1/1/3 (Evpn)
                        11621    0
2001:db8::6
                64500   11618    0 04d00h46m 1/1/2 (VpnIPv4)
                        11620    0           1/1/2 (VpnIPv6)
2001:db8::15
                64500   11621    0 04d00h46m 3/3/3 (Evpn)
                        11620    0
-------------------------------------------------------------------------------
```

EVPN VPWS (including ES and PW port) and related SRv6 configuration

EVPN VPWS is configured as an Epipe. The three instances to be configured are:

- SRv6
- BGP
- BGP-EVPN

The SRv6 instance is configured at the top level of the Epipe and represents the dataplane instantiation of SRv6. In Release 22.10, SRv6 can be used only with the EVPN control plane (BGP-EVPN). Under this hierarchy:

- The SRv6 locator is referenced.
- The function part (End.DX2) required for creating the service SID is defined. The function value can be allocated statically or automatically; in this case, it is allocated automatically.

The BGP instance is configured at the top level of the Epipe (**bgp 1** in the example), where RT and RD can be optionally configured (otherwise they are auto-derived) and where VSI policies are applied.

The BGP-EVPN instance is configured at the top level of the Epipe, with information related to EVPN control plane signaling. The number 1 in the **segment-routing-v6 1** configuration under the **bgp-evpn** context refers to the **bgp 1** instance. Consider the following aspects of this configuration:

- The source IPv6 address of the SRv6 tunnel is, in this case, the system IP address (2001:db8::14). If the source address is not defined in the configuration, the one from the global SRv6 configuration is used. Without the source address defined in one of those two places, the SRv6 transport for the service is non-operational.
- The Epipe is tagged with a tag 11 (0xb), which is used to identify the Epipe in the BGP export policy for the color extended community, as described in the BGP peering section.
- The **resolution fallback-tunnel-to-route-table** command forces the system to first check if the next hops for the received EVPN routes are present (or resolved) in the tunnel table, before falling back to the routing table for the resolution. This is required because the endpoint in the SR policy, along with the color, is installed in the tunnel table. A valid next-hop resolution via SR policy is a trigger for the addition of the SRH to the packet header. Without the SR policy, the SRH is not added, and the packet takes the shortest path to the final destination based on plain destination-based routing.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- The **route-next-hop** command identifies the next hop that is used in EVPN route advertisements. In this case, it is the system IPv6 address, but it can be any reachable locally configured IPv6 address.

The benefit of using the multi-instance provisioning model via three provisioning instances (SRv6, BGP, and BGP-EVPN) within the Epipe is that it offers more configuration flexibility when it comes to provisioning multiple data and control planes in the same service.

The Epipes on the two BNGs are configured with the same Ethernet tags. When the AD per-EVI routes with those tags are advertised, the AD per-EVI route advertised from BNG-1 has the P-bit set to 1 and the B-bit set to 0, making it the active route on PE-1. BNG-2 advertises the same AD per-EVI route (but with a different next-hop) with the value of the P-bit set to 0 and the B-bit to 1, making it the standby route on the PE-1.

```
[pr:/configure service epipe "evpn-dual-homing"]
A:admin@bng-1#
    admin-state enable
    service-id 11
    customer "1"
    segment-routing-v6 1 {
        locator "bng-1-loc" {
            function {
                end-dx2 {
                }
            }
        }
    }
    bgp 1 {
    }
    bgp-evpn {
        evi 11
        local-attachment-circuit "bng" {
            eth-tag 2
        }
        remote-attachment-circuit "access" {
            eth-tag 1
        }
        segment-routing-v6 1 {
            admin-state enable
            default-route-tag 0xb
            source-address 2001:db8::14
            resolution fallback-tunnel-to-route-table
            srv6 {
                instance 1
                default-locator "bng-1-loc"
            }
            route-next-hop {
                system-ipv6
            }
        }
    }
```

The preceding configuration defines the remote EVPN destination (or connection to the remote node over SRv6) in the Epipe. The local termination point in the Epipe is a PW port. The PW port is FPE-based and associated with the Epipe, as follows:

```
[pr:/configure pw-port 1]
A:admin@bng-1#
    encap-type qinq
    epipe "evpn-dual-homing" {
        admin-state enable
        fpe-id 3
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        oper-up-on-mh-standby true
    }
```

The **oper-up-on-mh-standby** command is related to an optimization of the ES in EVPN VPWS multihoming. The effect of this command is that the PW port and its SAPs remain operationally up when the associated ES is non-DF. This effect results in faster recovery during switchovers when the PW port has a large number of PW SAPs in the subscriber-management environment. In other words, during a non-DF to DF transition (standby to active), the system does not have to wait additional time to bring up thousands of PW SAPs from a down to an up state before it starts forwarding traffic.

The geo-redundant BNG setup in this chapter relies on two superimposed redundancy mechanisms:

- EVPN multihoming deployed in the access network-based DF election on the ES
- SRRP-based redundancy in ESM

These two redundancy mechanisms are, by themselves, independent of each other. However, to provide predictable results, they must be made co-dependent where one mechanism drives the other. In this context, EVPN MH as a technology deployed in the access network is used to detect the failures and to drive the state of SRRP. SRRP follows the states of the EVPN MH through operational groups.

ESM geo-redundancy based on SRRP does not support active-active mode of operation for a set of paired SRRP instances. As a result, the EVPN MH must be configured in the single-active (SA) mode of operation. The load-balancing between the two BNGs is achieved by distributing the activity of multiple SRRP instances across the two BNGs, some of them being active on one side and some of them on the other. In other words, the granularity of load balancing in ESM is per SRRP instance.

The configuration for coupling EVPN MH and SRRP through operational groups is shown in the following three configuration snippets, involving the definition of the operational group, the SA ES, and the SRRP messaging SAP.

The following configures the operational group "ES-1" with hold times of 0 s to prevent any reaction time delays during switchovers. Only the up time is set explicitly to 0 (from a default value of 4 s). The default value for a hold down time is already set to 0.

```
[pr:/configure service oper-group "ES-1"]
A:admin@bng-1#
    hold-time {
        up 0
    }
```

The following configures SA ES "ES-1" as part of the operational group "ES-1" and associated with PW-port 1. The state of the operational group "ES-1" depends on the state of the ES. If the ES is elected as the DF, then the operational group is up. Otherwise, it is down.

```
[pr:/configure service system]
A:admin@bng-1#
    bgp {
        evpn {
            ethernet-segment "ES-1" {
                admin-state enable
                esi 0x01010101010101010101
                orig-ip 2001:db8::14
                route-next-hop 2001:db8::14
                multi-homing-mode single-active
                oper-group "ES-1"
                df-election {
                    es-activation-timer 0
                    service-carving-mode manual
                    manual {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        preference {
                            mode revertive
                            value 150
                        }
                    }
                }
                association {
                    pw-port 1 {
                        pw-port-headend true
                    }
                }
            }
        }
    }
```

In the preceding ES configuration, the **orig-ip** command and the **route-next-hop** command must be explicitly configured in pure IPv6 environments. By default, the ES route is advertised with an IPv4 address as the next hop, which is not reachable in this example.

Under the **df-election** context, the configuration of the preference value influences which side becomes active and which becomes standby and the configuration of the preference mode indicates whether to enable revertive behavior, as explained in the Recovery section. In this example, the **es-activation-timer** command is set to 0 s, which means that the DF election occurs immediately after the failure is detected. Configuring this timer to a value greater than 0 allows some time for the peer routes to be received and collected before DF election is run. This configuration is applicable to environments with multiple routes and nodes in the same ES, but can be ignored in this example topology with only two BNGs.

The following configures the SRRP messaging SAP that tracks the state of the ES (active/standby) by monitoring the operational group "ES-1". The state of the SRRP messaging SAP is up if the ES is a DF, and down if it is a non-DF. An SRRP instance with a messaging SAP in a down state assumes the INIT state, rendering its side as standby. The pairing SRRP instance on the other BNG becomes active.

```
[pr:/configure service vprn "dual-homing" subscriber-interface "sub-int-1" group-interface
 "group-int-1"]
A:admin@bng-1#
    srrp 1 {
        admin-state enable
        keep-alive-interval 2
        message-path pw-1:2.4094
    }
    sap pw-1:2.4094 {
        monitor-oper-group "ES-1"
    }
```

The boot timer, whose purpose is to wait for routing to converge after a node reboots, is also important when revertive mode is enabled in the ES. The boot timer accounts for completion of subscriber synchronization between the two BNGs upon reboot. In this example, there is a wait time of 2 minutes after the node reboot for the network to converge before the ES is considered stable and ready for DF election.

```
*[pr:/configure redundancy]
A:admin@bng-1#
    bgp-evpn {
        ethernet-segment {
            boot-timer 120
        }
    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The multihomed Epipe configuration on PE-1 (the subscriber-connecting access node) is as follows:

```
[pr:/configure service epipe "dual-homing"]
A:admin@pe-1# info
    admin-state enable
    service-id 11
    customer "1"
    segment-routing-v6 1 {
        locator "pe-1-loc" {
            function {
                end-dx2 {
                }
            }
        }
    }
    sap 1/1/c3/1:*.* {
    }
    bgp-evpn {
        evi 11
        local-attachment-circuit "access" {
            eth-tag 1
        }
        remote-attachment-circuit "bng" {
            eth-tag 2
        }
        segment-routing-v6 1 {
            admin-state enable
            default-route-tag 0xb
            source-address 2001:db8::5
            resolution fallback-tunnel-to-route-table
            srv6 {
                instance 1
                default-locator "pe-1-loc"
            }
            route-next-hop {
                system-ipv6
            }
        }
    }
```

SR policy

An SR policy is identified by the tuple <`headend, color, endpoint`>, and its origin can be derived from the path computation element protocol (PCEP), BGP, or may be static via configuration.

In this example, a static SR policy is used in the access network as a traffic-engineering tool to guide traffic through a predetermined network path. This path is configured in the SR policy **segment-list** command as a list of next hop SIDs that are programmed into the SRH.

The SR policy is not explicitly applied to an object in a classical sense, such as routing policies, which are applied in routing and service contexts. Instead, the `end-node` and the `color` parameters configured in the SR policy are programmed (or activated) in the tunnel table. As such, the SR policy is used for the BNG advertised next-hop resolution in the tunnel table. Forcing the next-hop resolution through the tunnel-table is configured via the **resolution route-table | tunnel-table | fallback-tunnel-to-route-table** command under the Epipe configuration.

The **segment-list** command references the node-SIDs, while the **endpoint** configured in the SR policy is not a SID but instead a regular IPv6 address of the destination node. When multiple routing policies with

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

the same endpoint and color are defined, the one with the highest `preference` prevails and is installed in the tunnel-table.

The **head-end local** command signifies that the SR policy is locally defined and activated, as opposed to advertised to BGP peers.

The binding SID is a mandatory local SID associated with the SR policy. Its function value is `end-b6-encaps-red`.

The reachability of the first SID in the SR policy must be validated before the SR policy is activated.

Multiple segment lists can be defined in the SR policy. Traffic between these segment lists can be distributed according to the configured weight (not shown in the following configuration).

The SR policy does not have the configuration option for the **source-address**. Instead, it uses the one configured for the segment-routing configuration in the global routing context. Without the **source-address** configured in the global SRv6 context, the SR policy is not activated.

```
[pr:/configure router "Base" segment-routing sr-policies]
A:admin@bng-1#
    admin-state enable
    static-policy "to-pe-1-long-path" {
        admin-state enable
        color 20
        endpoint 2001:db8::5
        preference 150
        head-end local
        type srv6
        segment-routing-v6 {
            binding-sid 1 {
                locator {
                    locator-name "bng-1-loc"
                    function end-b6-encaps-red
                }
            }
        }
        segment-list 1 {
            admin-state enable
            segment 1 {
                srv6-sid 2001:db8:aaaa:4:0:1000::
            }
            segment 2 {
                srv6-sid 2001:db8:aaaa:a:0:1000::
            }
        }
    }
```

To verify that the SR policy has been activated:

```
A:admin@pe-1# show router segment-routing sr-policies static


===============================================================================
SR-Policies Path
===============================================================================
-------------------------------------------------------------------------------
Type            : srv6
Active          : Yes                   Owner           : static
Color           : 20
Head            : 0.0.0.0               Endpoint Addr   : 2001:db8::14
RD              : 0                     Preference      : 150
SRv6 BSID 1     : 2001:db8:aaaa:5:0:7000::
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
TunnelId        : 917506                Age             : 189018
Origin ASN      : 0                     Origin          : 0.0.0.0
NumReEval       : 1                     ReEvalReason    : route-add
NumActPathChange: 0                     Last Change     : 11/01/2022 08:39:33
Maintenance Policy: N/A

Path Segment Lists:
Segment-List    : 1                     Weight          : 1
S-BFD State     : Down                  S-BFD Transitio*: 0
Num Segments    : 2                     Last Change     : 11/01/2022 08:39:30
  Seg 1 SID   : 2001:db8:aaaa:a:0:1000::                 State : resolved-up
  Seg 2 SID   : 2001:db8:aaaa:4:0:1000::                 State : N/A

-------------------------------------------------------------------------------
Type            : srv6
Active          : Yes                   Owner           : static
Color           : 30
Head            : 0.0.0.0               Endpoint Addr   : 2001:db8::15
RD              : 0                     Preference      : 150
SRv6 BSID 1     : 2001:db8:aaaa:5:0:6000::
TunnelId        : 917507                Age             : 189036
Origin ASN      : 0                     Origin          : 0.0.0.0
NumReEval       : 1                     ReEvalReason    : route-add
NumActPathChange: 0                     Last Change     : 11/01/2022 08:39:33
Maintenance Policy: N/A

Path Segment Lists:
Segment-List    : 1                     Weight          : 1
S-BFD State     : Down                  S-BFD Transitio*: 0
Num Segments    : 2                     Last Change     : 11/01/2022 08:39:30
  Seg 1 SID   : 2001:db8:aaaa:4:0:1000::                 State : resolved-up
  Seg 2 SID   : 2001:db8:aaaa:a:0:1000::                 State : N/A
```

Verify routes

In this section, **show** commands are used to verify forwarding and show routing-related information. The outputs from the following **show** commands focus primarily (but not exclusively) on nodes BNG-1 and PE-1, which are the endpoints of the active leg of EVPN VPWS multihoming.

The following command shows the SRv6 information on BNG-1. The locator SID is advertised to other nodes via IS-IS. The locator SID is installed in the routing table of all other nodes and is used for reachability information.

```
A:admin@bng-1# show router segment-routing-v6 locator
===============================================================================
Locator bng-1-loc
===============================================================================
Admin State                 : Up
Prefix                      : 2001:db8:aaaa:14::/64
Block Length                : 48
Label Block                 :
Function Length             : 20
Flex Algorithm              : 0
Termination FPE             : 2
Static Function
  Max Entries               : 1
  Label-Block               :
---snip---
```

The following command lists different SIDs allocated to BNG-1:

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- one node SID (End), advertised in IS-IS
- three adjacency SIDs (End.X), representing the BNG-1 connections to adjacent nodes BNG-2, PE-2, and P-1 (these SIDs are used to build fast failover alternate paths and are also advertised in IS-IS)
- one EVPN SID (End.DX2), advertised via BGP-EVPN for EVPN routes
- two VPN SIDs (End.DT4 and End.DT6), advertised via BGP-IPVPN for VPN-IPv4 and VPN-IPv6 routes
- one binding SID (End.b6.encaps.red), for the SRv6 static segment-policy (this SID is not advertised in this case)

```
A:admin@bng-1# show router segment-routing-v6 local-sid

===============================================================================
Segment Routing v6 Local SIDs
===============================================================================
SID                                                 Type           Function
  Locator
  Context
-------------------------------------------------------------------------------
2001:db8:aaaa:14:0:1000::                           End            1
  bng-1-loc
  Base
2001:db8:aaaa:14:0:2000::                           End.b6.encaps* 2
  bng-1-loc
  None
2001:db8:aaaa:14:7fff:c000::                        End.DT6        524284
  bng-1-loc
  SvcId: 10 Name: dual-homing
2001:db8:aaaa:14:7fff:d000::                        End.DT4        524285
  bng-1-loc
  SvcId: 10 Name: dual-homing
2001:db8:aaaa:14:7fff:e000::                        End.DX2        524286
  bng-1-loc
  SvcId: 11 Name: evpn-dual-homing
2001:db8:aaaa:14:7fff:f000::                        End.X          524287
  bng-1-loc
  None
2001:db8:aaaa:14:8000::                            End.X          524288
  bng-1-loc
  None
2001:db8:aaaa:14:8000:1000::                        End.X          524289
  bng-1-loc
  None
```

The following route table on BNG-1 shows:

- system IPv6 interfaces from all other nodes learned through IS-IS
- IPv6 prefixes on links between all nodes learned through IS-IS
- SRv6 locators from all nodes learned through IS-IS
- all local SIDs except for service SIDs, which are part of services and not global routing

BNGs are connected to both the core network (VPRN) and the access network (EVPN), so network addresses from all nodes are visible on a BNG. By contrast, the PE-1 node in the EVPN part of the network sees only routes in the EVPN part of the network, but not the routes in the VPRN part of the network (Figure 332: Example topology).

```
A:admin@bng-1# show router route-table ipv6

===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
IPv6 Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
2001:db8::4/128                                 Remote  ISIS      02d03h06m  18
       fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                        10
2001:db8::5/128                                 Remote  ISIS      02d03h06m  18
       fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                        20
2001:db8::6/128 [L]                             Remote  ISIS(1)   02d03h06m  18
       fe80::b6a0:ffff:fe00:0-"int-1-bng-1-pe-2"                       10
2001:db8::a/128                                 Remote  ISIS      02d03h06m  18
       fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                        20
2001:db8::14/128                                Local   Local     02d03h06m  0
       system                                                           0
2001:db8::15/128 [L]                            Remote  ISIS(1)   02d03h06m  18
       fe80::b697:ffff:fe00:0-"int-1-bng-1-bng-2"                      10
2001:db8::100/120                               Remote  ISIS      02d03h06m  18
       fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                        20
2001:db8::200/120                               Remote  ISIS      02d03h06m  18
       fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                        30
2001:db8::300/120                               Remote  ISIS      02d03h06m  18
       fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                        20
2001:db8::400/120                               Remote  ISIS      02d03h06m  18
       fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                        30
2001:db8::500/120                               Local   Local     02d03h06m  0
       int-1-bng-1-p-1                                                  0
2001:db8::600/120                               Remote  ISIS      02d03h06m  18
       fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                        30
2001:db8::700/120                               Local   Local     02d03h06m  0
       int-1-bng-1-bng-2                                                0
2001:db8::800/120                               Local   Local     02d03h06m  0
       int-1-bng-1-pe-2                                                 0
2001:db8::900/120 [L]                           Remote  ISIS(1)   02d03h06m  18
       fe80::b6a0:ffff:fe00:0-"int-1-bng-1-pe-2"                       20
2001:db8::a00/120                               Remote  ISIS      02d03h06m  18
       fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                        20
2001:db8::6400/120                              Remote  ISIS      02d03h06m  18
       fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                        30
2001:db8::6500/120 [L]                          Remote  ISIS(1)   02d03h06m  18
       fe80::b6a0:ffff:fe00:0-"int-1-bng-1-pe-2"                       20
2001:db8:aaaa:4::/64                            Remote  ISIS      02d03h06m  18
       2001:db8:aaaa:4::/64 (tunneled:SRV6-ISIS)                       20
2001:db8:aaaa:5::/64                            Remote  ISIS      02d03h06m  18
       2001:db8:aaaa:5::/64 (tunneled:SRV6-ISIS)                       30
2001:db8:aaaa:6::/64                            Remote  ISIS(1)   02d03h06m  18
       2001:db8:aaaa:6::/64 (tunneled:SRV6-ISIS)                       20
2001:db8:aaaa:a::/64                            Remote  ISIS      02d03h06m  18
       2001:db8:aaaa:a::/64 (tunneled:SRV6-ISIS)                       30
2001:db8:aaaa:14::/64                           Local   SRV6      02d03h06m  3
       fe80::201-"_tmnx_fpe_2.a"                                        0
2001:db8:aaaa:14:0:1000::/128                   Local   SRV6      02d03h06m  3
       Black Hole                                                       0
2001:db8:aaaa:14:0:2000::/128                   Local   SRV6-Pol* 02d03h06m  14
       2001:db8::5 (tunneled:SRV6-Policy:917506)                       1
2001:db8:aaaa:14:7fff:f000::/128                Local   ISIS(1)   02d03h06m  18
       2001:db8:aaaa:14:7fff:f000:: (tunneled:SRV6-ISIS)               10
2001:db8:aaaa:14:8000::/128                     Local   ISIS      02d03h06m  18
       2001:db8:aaaa:14:8000:: (tunneled:SRV6-ISIS)                    10
2001:db8:aaaa:14:8000:1000::/128                Local   ISIS(1)   02d03h06m  18
       2001:db8:aaaa:14:8000:1000:: (tunneled:SRV6-ISIS)               10
2001:db8:aaaa:15::/64                           Remote  ISIS(1)   02d03h06m  18
       2001:db8:aaaa:15::/64 (tunneled:SRV6-ISIS)                      20
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The state of the ES "ES-1" on the two BNGs is DF for BNG-1 and NDF for BNG-2, as follows:

```
A:admin@bng-1# /show  service  id "evpn-dual-homing" ethernet-segment "ES-1"
---snip---
Pw-Port               Eth-Seg                          Status
-------------------------------------------------------------------------------
1                     ES-1                             DF



A:admin@bng-2# /show  service  id "evpn-dual-homing" ethernet-segment "ES-1"
---snip---
Pw-Port               Eth-Seg                          Status
-------------------------------------------------------------------------------
1                     ES-1                             NDF
```

PE-1 receives the following two AD per-EVI EVPN routes: the first AD per-EVI route is received from BNG-1 as primary (P=1, B=0), while the second AD per-EVI route is received from BNG-2 as backup (P=0, B=1):

```
A:admin@pe-1# show router  bgp routes evpn auto-disc tag 2 hunt
===============================================================================
 BGP Router ID:192.0.2.5        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Auto-Disc Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Nexthop        : 2001:db8::14
From           : 2001:db8::14
Res. Nexthop   : fe80::b69e:ffff:fe00:0
Local Pref.    : 100                    Interface Name : int-1-pe-1-p-1
AIGP Metric    : None                   IGP Cost       : 30
Community      : color:00:20 target:64500:11
                 l2-attribute:MTU: 1514 C: 0 P: 1 B: 0
Originator Id  : None                   Peer Router Id : 192.0.2.20
Flags          : Used Valid Best IGP
Route Source   : Internal
EVPN type      : AUTO-DISC
ESI            : 01:01:01:01:01:01:01:01:01:01
Tag            : 2
Route Dist.    : 192.0.2.20:11
MPLS Label     : 524286
Route Tag      : 0
Last Modified  : 02d03h38m
SRv6 TLV Type  : SRv6 L2 Service TLV (6)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:aaaa:14::
Full Sid       : 2001:db8:aaaa:14:7fff:e000::
Behavior       : End.DX2 (21)
SRv6 SubSubTLV : SRv6 SID Structure (1)
Loc-Block-Len  : 48                     Loc-Node-Len   : 16
Func-Len       : 20                     Arg-Len        : 0
Tpose-Len      : 20                     Tpose-offset   : 64
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Nexthop        : 2001:db8::15
From           : 2001:db8::15
Res. Nexthop   : fe80::b6a1:ffff:fe00:0
Local Pref.    : 100                    Interface Name : int-1-pe-1-p-2
AIGP Metric    : None                   IGP Cost       : 30
Community      : color:00:30 target:64500:11
                 l2-attribute:MTU: 1514 C: 0 P: 0 B: 1
Originator Id  : None                   Peer Router Id : 192.0.2.21
Flags          : Used Valid Best IGP
Route Source   : Internal
EVPN type      : AUTO-DISC
ESI            : 01:01:01:01:01:01:01:01:01:01
Tag            : 2
Route Dist.    : 192.0.2.21:11
MPLS Label     : 524286
Route Tag      : 0
Last Modified  : 02d03h38m
SRv6 TLV Type  : SRv6 L2 Service TLV (6)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:aaaa:15::
Full Sid       : 2001:db8:aaaa:15:7fff:e000::
Behavior       : End.DX2 (21)
SRv6 SubSubTLV : SRv6 SID Structure (1)
Loc-Block-Len  : 48                     Loc-Node-Len   : 16
Func-Len       : 20                     Arg-Len        : 0
Tpose-Len      : 20                     Tpose-offset   : 64
```

When an IPv4 packet from the traffic generator enters the SAP of the Epipe on PE-1, the packet goes through the Epipe toward the Epipe's other endpoint, which is the EVPN destination on BNG-1. BNG-1 is the primary EVPN destination because PE-1 receives an active (P=1, B=0) AD-per EVI route with the tag 2 from BNG-1. This route has:

- color 20
- next-hop 2001:db8::14 (BNG-1)
- service SID 2001:db8:aaaa:14:7fff:e000:: (End.DX2 on BNG-1)

The service SID is the final destination of the packet and it is inserted in the SRH as the first segment. The received next-hop in the route must be resolved before the route is validated and installed in the routing table. The first attempt to resolve the next-hop for this route is performed through the tunnel table, as indicated by the **resolution fallback-tunnel-to-route-table** command. The following command shows the IPv6 tunnel table with an entry for the next-hop address 2001:db8::14 and color 20, as received in the AD per-EVI route. The owner of this entry is SRv6-policy and the encapsulation is SRv6.

```
A:admin@pe-1# show router tunnel-table ipv6

===============================================================================
IPv6 Tunnel Table (Router: Base)
===============================================================================
Destination                                 Owner     Encap TunnelId  Pref
Nexthop                                     Color           Metric
-------------------------------------------------------------------------------
2001:db8::14/128                            srv6-pol  SRV6  917506    14
  fpe_1.a                                     20              0
2001:db8::15/128                            srv6-pol  SRV6  917507    14
  fpe_1.a                                     30              0
2001:db8:aaaa:4::/64 [L]                    srv6-isis SRV6  524290    0
  fe80::b69e:ffff:fe00:0-"int-1-pe-1-p-1"                     20
2001:db8:aaaa:5:8000:1000::/128 [L]         srv6-isis SRV6  524289    0
  fe80::b69e:ffff:fe00:0-"int-1-pe-1-p-1"                     10
2001:db8:aaaa:5:8000:2000::/128 [L]         srv6-isis SRV6  524292    0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
  fe80::b6a1:ffff:fe00:0-"int-1-pe-1-p-2"                          10
2001:db8:aaaa:a::/64 [L]                       srv6-isis SRV6  524293    0
  fe80::b6a1:ffff:fe00:0-"int-1-pe-1-p-2"                          20
2001:db8:aaaa:14::/64 [L]                      srv6-isis SRV6  524297    0
  fe80::b69e:ffff:fe00:0-"int-1-pe-1-p-1"                          30
2001:db8:aaaa:15::/64 [L]                      srv6-isis SRV6  524294    0
  fe80::b6a1:ffff:fe00:0-"int-1-pe-1-p-2"                          30
---snip---
```

A more detailed output for the next-hop entry in the tunnel table reveals two additional SIDs, which are the SIDs from the configured SR policy:

- 2001:db8:aaaa:a:0:1000:: is the node SID for P-2
- 2001:db8:aaaa:4:0:1000:: is the node SID for P-1

```
A:admin@pe-1# show  router tunnel-table ipv6 detail

Tunnel Table (Router: Base)
---snip---
Destination      : 2001:db8::14/128
NextHop          : fpe_1.a
NextHop Weight   : 1
Tunnel Flags     : has-color
Age              : 03d02h15m                    Color            : 20
CBF Classes      : (Not Specified)
Owner            : srv6-pol                     Encap            : SRV6
Tunnel ID        : 917506                       Preference       : 14
Tunnel SRV6 SID  : 2001:db8:aaaa:a:0:1000::     Tunnel Metric    : 0
                 : 2001:db8:aaaa:4:0:1000::
Tunnel MTU       :  -                   Max Label Stack  : 2
---snip---
```

The presence of the two node SIDs for P-2 and P-1 implies that the nodes P-2 and P-1 must be visited, in that order, on the path from PE-1 to BNG-1. This path is accomplished by inserting the SRH header into the packet. The trigger for the SRH creation and insertion into the packet is the next-hop resolution via SR policy.

Because the **End.B6.Encap.Red** function for the binding SID in the SR policy reduces the encapsulation, as defined in RFC 8986, the node SID for P-2 is not part of the SRH, even though it is listed as the first node to be visited in the segment list of the SR policy. As the first node to be visited, its SID is directly included in the DA in the SRv6 header of the packet without repeating itself in the SRH. The SRH contains two SIDs: the node SID for P-1 (2001:db8:aaaa:4:0:1000::) and the service SID (2001:db8:aaaa:14:7fff:e000:) which represents the Epipe on BNG-1, retrieved from the received EVPN AD-per-EVI route. These two SIDs are, in turn, copied to the IPv6 DA field of the packet by the visited nodes, as described in the "Segment Routing over IPv6" chapter in the Segment Routing and PCE volume in the *7450 ESS, 7750 SR, and 7950 XRS MD-CLI Advanced Configuration Guide - Part I*.

The longest match lookup for the DA 2001:db8:aaaa:a:0:1000:: (P-2 node-SID) reveals the next hop leading to its SRv6 locator (P-2 locator):

```
A:admin@pe-1#  show router route-table ipv6 2001:db8:aaaa:a:0:1000:: extensive
===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix             : 2001:db8:aaaa:a::/64
  Protocol              : ISIS
  Age                   : 03d02h40m
  Preference            : 18
  Next-Hop              : 2001:db8:aaaa:a::/64 (SRV6-ISIS tunnel)
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    ---snip---
```

The tunnel table for next hop 2001:db8:aaaa:a::/64 (P-2 SRv6 locator) points to the link-local IPv6 address of the adjacent node that advertised this next hop (P-2 SRv6 locator) in IS-IS:

```
A:admin@pe-1# show router tunnel-table ipv6 2001:db8:aaaa:a::/64 detail
===============================================================================
Tunnel Table (Router: Base)
Destination      : 2001:db8:aaaa:a::/64 [L]
NextHop          : fe80::b6a1:ffff:fe00:0-"int-1-pe-1-p-2"
Tunnel Flags     : has-lfa
Age              : 03d02h42m
CBF Classes      : (Not Specified)
Owner            : srv6-isis (0)        Encap            : SRV6
Tunnel ID        : 524293               Preference       : 0
Tunnel SRV6 SID  :  -                   Tunnel Metric    : 20
Tunnel MTU       : 8894                 Max Label Stack  : 0
-------------------------------------------------------------------------------
```

The neighbor cache on PE-1 shows the mapping between the link local IPv6 of the adjacent next hop and its MAC address, which is the destination MAC address in the packet:

```
A:admin@pe-1# show router neighbor
===============================================================================
Neighbor Table (Router: Base)
===============================================================================
IPv6 Address                                Interface
   MAC Address               State         Expiry          Type         RTR
-------------------------------------------------------------------------------
---snip---
fe80::b6a1:ffff:fe00:0                      int-1-pe-1-p-2
   b4:a1:01:01:00:05         REACHABLE     00h00m25s       Dynamic      Yes
---snip---
```

In summary:

- PE-1 receives an AD per-EVI from BNG-1 with the address of the final destination (service SID), the route color, and the next-hop of the advertised route.
- The next-hop resolution in the tunnel table points to the SR policy with SRv6 encapsulation.
- The SR policy is configured with a list of node SIDs that need to be visited, with the first SID in the list copied to the destination IPv6 address of the packet.
- The remaining SID in the SR policy along with the service SID advertised in the route are populated in the SRH.
- The destination MAC address is resolved through a series of lookups that lead to the directly connected IS-IS node that advertised the SRv6 locator for the first node to be visited.

Without the SR policy, the SRH would not be inserted into the packet. As a result, the final destination (the service SID on BNG-1 for EVPN, 2001:db8:aaaa:14:7fff:e000::) would immediately be populated in the DA field of the packet. The destination-based lookup for this DA gives us a different next hop for the packet (P-1 instead of P-2):

```
A:admin@pe-1# show router route-table ipv6 2001:db8:aaaa:14:7fff:e000:: extensive
Route Table (Router: Base)
Dest Prefix             : 2001:db8:aaaa:14::/64
  Protocol              : ISIS
  Age                   : 04d07h47m
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
  Preference           : 18
  Next-Hop             : 2001:db8:aaaa:14::/64 (SRV6-ISIS tunnel)
---snip---

A:admin@pe-1# show router tunnel-table 2001:db8:aaaa:14::/64 detail
Tunnel Table (Router: Base)
Destination      : 2001:db8:aaaa:14::/64 [L]
NextHop          : fe80::b69e:ffff:fe00:0-"int-1-pe-1-p-1"
---snip---
Owner            : srv6-isis (0)       Encap            : SRV6
Tunnel ID        : 524297              Preference       : 0
Tunnel SRV6 SID  :  -                  Tunnel Metric    : 30
Tunnel MTU       : 8894                Max Label Stack  : 0
```

VPRN related SRv6 configuration

The SRv6 locator and functions in a VPRN are configured in a similar way as in the EVPN.

The export VRF policy advertises the subscriber subnets on the active SRRP node with a better metric than the same routes on the node with standby SRRP. The result is that traffic in the downstream direction is attracted to the node with the active SRRP instance.

```
[pr:/configure service vprn "dual-homing"]
A:admin@bng-1#
    admin-state enable
    service-id 10
    customer "1"
    segment-routing-v6 1 {
        locator "bng-1-loc" {
            function {
                end-dt4 {
                }
                end-dt6 {
                }
            }
        }
    }
    bgp-ipvpn {
        segment-routing-v6 1 {
            admin-state enable
            route-distinguisher "192.0.2.20:10"
            source-address 2001:db8::14
            vrf-target {
                community "target:64500:10"
            }
            vrf-export {
                policy ["srrp-aware-routing"]
            }
            srv6 {
                instance 1
                default-locator "bng-1-loc"
            }
        }
    }
    bgp {
    }
```

```
[pr:/configure policy-options policy-statement "srrp-aware-routing"]
A:admin@bng-1# info
    description "vrf-export; advertising sub-if routes based on srrp state"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    entry 10 {
        from {
            state srrp-master
            protocol {
                name [direct]
            }
        }
        action {
            action-type accept
            local-preference 150
            community {
                add ["dual-homing"]
            }
        }
    }
    entry 20 {
        from {
            state srrp-non-master
            protocol {
                name [direct]
            }
        }
        action {
            action-type accept
            local-preference 100
            community {
                add ["dual-homing"]
            }
        }
    }
```

Multi-chassis redundancy configuration

The full configuration under the **configure redundancy multi-chassis peer <IP address>** context is provided in the configuration files in the Appendix.

This section emphasizes on the time synchronization of each node to the same clock, which is an aspect of the subscriber synchronization that is easily overlooked. Without times that match on both BNGs, the subscribers are not synchronized. In the example, simple network timing protocol (SNTP) is used.

```
[pr:/configure system time]
A:admin@bng-1#
    prefer-local-time true
    zone {
        standard {
            name cst
        }
    }
    dst-zone "CDT" {
        end {
            day sunday
            month november
            hours-minutes "02:00"
        }
        start {
            day sunday
            month march
            hours-minutes "02:00"
        }
    }
    sntp {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        admin-state enable
        server 135.227.160.253 {
        }
    }
```

DHCP server configuration

DHCP configuration shows that all pools are synchronized via the **peer** command, except for the pool for PPPoEv4 sessions. The IP addresses in the pool for PPPoE sessions are synchronized indirectly as part of the PPPoE session synchronization.

The second important aspect of the DHCP redundancy in access driven mode is that the DHCP servers are associated with interfaces configured with the same IP address on both BNGs (see the Appendix).

```
[pr:/configure service vprn "dual-homing"]
A:admin@bng-1#
    dhcp-server {
        dhcpv4 "dhcpv4" {
            admin-state enable
            pool-selection {
                use-gi-address {
                    scope pool
                }
                use-pool-from-client {
                }
            }
            pool "dhcpv4-1" {
                max-lease-time 1200
                failover {
                    admin-state enable
                    peer 192.0.2.21 {
                        sync-tag "dhcp4"
                    }
                }
                subnet 10.10.0.0/24 {
                    address-range 10.10.0.10 end 10.10.0.100 {
                        failover-control-type access-driven
                    }
                }
            }
            pool "pppoev4-1" {
                max-lease-time 1200
                subnet 10.10.1.0/24 {
                    address-range 10.10.1.10 end 10.10.1.100 {
                        failover-control-type access-driven
                    }
                }
            }
        }
        dhcpv6 "dhcpv6" {
            admin-state enable
            pool-selection {
                use-pool-from-client {
                }
            }
            pool "dhcpv6-1" {
                delegated-prefix {
                    minimum 56
                }
                failover {
                    admin-state enable
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    peer 192.0.2.21 {
                        sync-tag "dhcp6"
                    }
                }
                prefix 2001:db8:bbbb::/56 {
                    failover-control-type access-driven
                    preferred-lifetime 900
                    valid-lifetime 1200
                    renew-time 600
                    rebind-time 1000
                    prefix-type {
                        wan-host true
                    }
                }
                prefix 2001:db8:bbbb:100::/56 {
                    failover-control-type access-driven
                    preferred-lifetime 900
                    valid-lifetime 1200
                    renew-time 600
                    rebind-time 1000
                    prefix-type {
                        pd true
                    }
                }
            }
        }
    }


[pr:/configure service vprn "dual-homing" interface "loopback-1"]
A:admin@bng-1#
    admin-state enable
    loopback true
    ipv4 {
        local-dhcp-server "dhcpv4"
        primary {
            address 192.168.0.1
            prefix-length 32
        }
    }
    ipv6 {
        local-dhcp-server "dhcpv6"
        address 2001:db8::1 {
            prefix-length 128
        }
    }
```

Subscriber management configuration

The subscriber management configuration is provided in the configuration files shown in the Appendix.

The following output shows that there are 20 dual-stack subscribers instantiated:

```
A:admin@bng-1# show service active-subscribers summary

===============================================================================
Active Subscriber table summary
===============================================================================
Total Count       : 20
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The following output, showing a hierarchical view of each subscriber type, is shortened for brevity.

```
A:admin@bng-1# show service active-subscribers hierarchy

===============================================================================
Active Subscribers Hierarchy
===============================================================================
-- ipoe-ds-1
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.11] - sla:sla-profile-1
       |
       +-- IPOE-session - mac:00:13:01:00:00:01 - svc:10
           |
           |-- 10.10.0.10 - DHCP
           |
           +-- 2001:db8:bbbb:14::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:15::/64 - DHCP6-PD-MR

---snip---

-- pppoe-1
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.1] - sla:sla-profile-1
       |
       +-- PPP-session - mac:00:11:01:00:00:01 - sid:1 - svc:10
           |
           |-- 10.10.1.10 - IPCP
           |
           +-- 2001:db8:bbbb::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:1::/64 - DHCP6-PD-MR

---snip---
```

Summary of show commands

Show commands are powerful tools when building and troubleshooting networks. An extensive set of **show** command outputs relevant to the example topology described in this chapter are provided in separate files for BNG-1, P-1, and PE-1. The following is the list of these commands. Not all commands are applicable to all nodes; for example, commands related to subscriber management do not apply to P and PE nodes.

The following commands show ports and router interfaces:

```
show port
show router <id> interface
```

The following command shows FPE information:

```
show fwd-path-ext fpe <id>
```

The following commands show EVPN VPWS information:

```
show service id <name> base
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
show service id <name> bgp-evpn
show service system bgp-evpn ethernet-segment name <name>
show service system bgp-evpn ethernet-segment name <name> evi evi-1 <id>
show router bgp routes evpn eth-seg detail
show router bgp routes evpn auto-disc tag <id> detail
show service id <name> ethernet-segment <name>
show service id <name> segment-routing-v6 detail
show service id <name> segment-routing-v6 instance 1 destinations
show service id <name> segment-routing-v6 instance 1 end-dx2
show service id <name> segment-routing-v6 instance 1 locator <name>
```

The following commands show VPRN information:

```
show service id <name> base
show service id <name> bgp-ipvpn segment-routing-v6
show service id <name> segment-routing-v6 instance 1 locator <name>
show router <id> route-table
show router <id> route-table ipv6
show router bgp routes vpn-ipv4 hunt
show router bgp routes vpn-ipv6 hunt
```

The following commands show SRv6-related information:

```
show router segment-routing-v6 base-routing-instance all
show router segment-routing-v6 local-sid context  "Base"
show router segment-routing-v6 local-sid context <id>
```

The following commands show routing and forwarding information:

```
show router route-table ipv6
show router tunnel-table ipv6
show router fp-tunnel-table 1 ipv6
```

The following commands show BGP and IS-IS information:

```
show router bgp  summary
show router isis  adjacency
show router isis database
show router isis 1 adjacency
show router isis 1 database
show router isis <instance-id> adjacency
show router isis <instance-id> database
```

The following commands show redundancy-related information:

```
show redundancy multi-chassis sync peer <ip-address> detail
tools dump redundancy multi-chassis sync-database
```

The following commands show subscriber-related information:

```
show service active-subscribers summary
show service active-subscribers hierarchy
show service id <id> subscriber-hosts detail
show service id <id> ipoe session
show service id <id> pppoe session
show service id <id> dhcp lease-state
show service id <id> dhcp6 lease-state
show srrp <id> detail
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The following commands show DHCP information:

```
show router <id> dhcp local-dhcp-server <name> summary
show router <id> dhcp6 local-dhcp-server <name> summary
show router <id> dhcp local-dhcp-server <name> leases
show router <id> dhcp6 local-dhcp-server <name> leases
```

Conclusion

Resilient overlay networks that support traffic engineering can be built with the simplicity of a single dataplane protocol (IPv6) using SRv6. When SRv6 is deployed in the access network in combination with BNG dual-homing, subscribers are protected against network failures with minimal recovery times. This chapter shows one example of such network.

Appendix

The following configurations are included.

- bng-1-config
- bng-2-config
- p-1-config
- p-2-config
- pe-1-config
- pe-2-config
- users.radius

bng-1-config

```
# TiMOS-C-22.10.R1 cpm/x86_64 Nokia 7750 SR Copyright (c) 2000-2022 Nokia.
# All rights reserved. All use subject to applicable license agreements.
# Built on Sun Oct 30 14:49:55 PDT 2022 by builder in /builds/c/2210B/R1/panos/main/sros
# Configuration format version 22.10 revision 0

# Generated 2022-11-01T10:41:30.6-05:00 by admin from 135.231.208.32
# Commit ID 5
#   Committed 2022-11-01T10:37:59.8-05:00 by admin (MD-CLI) from 135.231.208.32
# Commit ID 4
#   Committed 2022-11-01T10:36:36.8-05:00 by admin (MD-CLI) from 135.231.208.32
# Commit ID 3
#   Committed 2022-11-01T10:33:20.1-05:00 by admin (MD-CLI) from 135.231.208.32
# Commit ID 2
#   Committed 2022-11-01T10:32:34.0-05:00 by admin (MD-CLI) from 135.231.208.32
# Commit ID 1
#   Committed 2022-11-01T10:04:53.3-05:00 by system (MD-CLI) from Console
#   Log       "System booted version C-22.10.R1."

configure {
    aaa {
        radius {
            server-policy "radius-server-1" {
                servers {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                router-instance "Base"
                server 1 {
                    server-name "free-radius-1"
                }
            }
        }
    }
    card 1 {
        card-type iom5-e
        mda 1 {
            mda-type me6-100gb-qsfp28
        }
        mda 2 {
            mda-type me6-100gb-qsfp28
        }
    }
    fwd-path-ext {
        sdp-id-range {
            start 17500
            end 17600
        }
        fpe 1 {
            description "srv6 origination"
            path {
                pxc 1
            }
            application {
                srv6 {
                    type origination
                }
            }
        }
        fpe 2 {
            description "srv6 termination"
            path {
                pxc 2
            }
            application {
                srv6 {
                    type termination
                }
            }
        }
        fpe 3 {
            description "pw-port on a single port"
            path {
                pxc 3
            }
            application {
                pw-port-extension {
                }
            }
        }
    }
    log {
        filter "1001" {
            named-entry "10" {
                description "Collect only events of major severity or higher"
                action forward
                match {
                    severity {
                        gte major
                    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
            }
        }
        log-id "100" {
            description "Default Serious Errors Log"
            filter "1001"
            source {
                main true
            }
            destination {
                memory {
                    max-entries 500
                }
            }
        }
        log-id "99" {
            description "Default System Log"
            source {
                main true
            }
            destination {
                memory {
                    max-entries 500
                }
            }
        }
    }
    policy-options {
        community "color-20" {
            member "color:00:20" { }
        }
        community "dual-homing" {
            member "target:64500:10" { }
        }
        policy-statement "pol-color-20" {
            entry 10 {
                from {
                    family [evpn]
                    tag 12
                }
                action {
                    action-type accept
                    community {
                        add ["color-20"]
                    }
                }
            }
            entry 20 {
                from {
                    family [evpn]
                    tag 11
                }
                action {
                    action-type accept
                    community {
                        add ["color-20"]
                    }
                }
            }
        }
        policy-statement "srrp-aware-routing" {
            description "vrf-export; advertizing sub-if routes based on srrp state"
            entry 10 {
                from {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    state srrp-master
                    protocol {
                        name [direct]
                    }
                }
                action {
                    action-type accept
                    local-preference 150
                    community {
                        add ["dual-homing"]
                    }
                }
            }
            entry 20 {
                from {
                    state srrp-non-master
                    protocol {
                        name [direct]
                    }
                }
                action {
                    action-type accept
                    local-preference 100
                    community {
                        add ["dual-homing"]
                    }
                }
            }
        }
        policy-statement "sub-mgmt-routes" {
            description "this is not applied, but if it was - it would work fine through vrf-
export"
            entry 10 {
                from {
                    protocol {
                        name [sub-mgmt]
                    }
                }
                action {
                    action-type accept
                    community {
                        add ["dual-homing"]
                    }
                }
            }
        }
    }
    port 1/1/c1 {
        admin-state enable
        connector {
            breakout c10-10g
        }
    }
    port 1/1/c1/1 {
        admin-state enable
        description "to p-1 access; mutually exclusive with lag-2"
        monitor-oper-group "network-port"
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/2 {
        admin-state disable
        description "lag-2 to p-1 access; mutually exclusive with port 1/1/c1/1"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/3 {
        admin-state disable
        description "lag-2 to p-1 access; mutually exclusive with port 1/1/c1/1"
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/4 {
        admin-state enable
        description "spare port to p-1 access"
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/5 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/6 {
        admin-state enable
        description "to bng-2"
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/7 {
        admin-state enable
        description "to pe-2 network"
        oper-group "network-port"
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/8 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/9 {
        admin-state enable
        ethernet {
            mode hybrid
            dot1x {
                tunneling true
            }
        }
    }
    port 1/1/c1/10 {
        admin-state enable
        description "RADIUS and mirroring"
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2 {
        admin-state enable
        connector {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            breakout c1-100g
        }
    }
    port 1/1/c2/1 {
        admin-state enable
        ethernet {
            mode hybrid
            encap-type dot1q
            dot1x {
                tunneling true
            }
        }
    }
    port 1/1/c4 {
        admin-state enable
        connector {
            breakout c4-10g
        }
    }
    port 1/1/c4/1 {
        admin-state enable
        description "PXC 1,2,3; fpe 1,2,3; srv6 orig, term, pw-port single port"
        ethernet {
            mode hybrid
            dot1x {
                tunneling true
            }
        }
    }
    port 1/1/c4/2 {
        admin-state enable
    }
    port 1/1/c4/3 {
        admin-state enable
    }
    port 1/1/c4/4 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port pxc-1.a {
        admin-state enable
    }
    port pxc-1.b {
        admin-state enable
    }
    port pxc-2.a {
        admin-state enable
    }
    port pxc-2.b {
        admin-state enable
    }
    port pxc-3.a {
        admin-state enable
    }
    port pxc-3.b {
        admin-state enable
    }
    port-xc {
        pxc 1 {
            admin-state enable
            description "fpe srv6 origin"
            port-id 1/1/c4/1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
        pxc 2 {
            admin-state enable
            description "fpe srv6 termination"
            port-id 1/1/c4/1
        }
        pxc 3 {
            admin-state enable
            description "fpe pw-port single port"
            port-id 1/1/c4/1
        }
    }
    pw-port 1 {
        encap-type qinq
        epipe "evpn-dual-homing" {
            admin-state enable
            fpe-id 3
            oper-up-on-mh-standby true
        }
    }
    redundancy {
        multi-chassis {
            peer 192.0.2.21 {
                admin-state enable
                sync {
                    admin-state enable
                    local-dhcp-server true
                    srrp true
                    sub-mgmt {
                        ipoe true
                        pppoe true
                    }
                    tags {
                        pw-port 1 {
                            sync-tag "pw-port-1"
                        }
                    }
                    track-srrp 1 {
                    }
                }
            }
        }
        bgp-evpn {
            ethernet-segment {
                boot-timer 120
            }
        }
    }
    router "Base" {
        autonomous-system 64500
        router-id 192.0.2.20
        interface "int-1-bng-1-bng-2" {
            port 1/1/c1/6:1
            gre-termination true
            ipv4 {
                primary {
                    address 192.168.7.1
                    prefix-length 24
                }
            }
            ipv6 {
                address 2001:db8::701 {
                    prefix-length 120
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
        }
        interface "int-1-bng-1-p-1" {
            port 1/1/c1/1:1
            ipv6 {
                address 2001:db8::501 {
                    prefix-length 120
                }
            }
        }
        interface "int-1-bng-1-pe-2" {
            port 1/1/c1/7:1
            ipv6 {
                address 2001:db8::801 {
                    prefix-length 120
                }
            }
        }
        interface "system" {
            ipv4 {
                primary {
                    address 192.0.2.20
                    prefix-length 32
                }
            }
            ipv6 {
                bfd {
                    admin-state enable
                    transmit-interval 100
                    receive 100
                    multiplier 2
                }
                address 2001:db8::14 {
                    prefix-length 128
                }
            }
        }
        interface "to-radius" {
            port 1/1/c1/10:114
            ipv4 {
                primary {
                    address 192.168.114.20
                    prefix-length 24
                }
            }
        }
        bgp {
            admin-state enable
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            rapid-update {
                vpn-ipv4 true
                vpn-ipv6 true
                evpn true
            }
            extended-nh-encoding {
                vpn-ipv4 true
            }
            advertise-ipv6-next-hops {
                vpn-ipv6 true
                vpn-ipv4 true
            }
            group "evpn" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                peer-as 64500
                local-address 2001:db8::14
                bfd-liveness true
                family {
                    evpn true
                }
            }
            group "ipvpn" {
                peer-as 64500
                local-address 2001:db8::14
                bfd-liveness true
                family {
                    vpn-ipv4 true
                    vpn-ipv6 true
                }
            }
            neighbor "2001:db8::5" {
                group "evpn"
                export {
                    policy ["pol-color-20"]
                }
            }
            neighbor "2001:db8::6" {
                group "ipvpn"
            }
            neighbor "2001:db8::15" {
                group "evpn"
            }
        }
        isis 0 {
            admin-state enable
            advertise-passive-only false
            advertise-router-capability area
            ipv6-routing native
            level-capability 2
            traffic-engineering true
            area-address [49.0001]
            loopfree-alternate {
                remote-lfa {
                }
                ti-lfa {
                }
            }
            traffic-engineering-options {
                ipv6 true
                application-link-attributes {
                }
            }
            segment-routing-v6 {
                admin-state enable
                locator "bng-1-loc" {
                    level-capability 2
                    level 2 {
                        metric 10
                    }
                }
            }
            interface "int-1-bng-1-p-1" {
            }
            interface "system" {
            }
            level 2 {
                wide-metrics-only true
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
        isis 1 {
            admin-state enable
            advertise-passive-only false
            advertise-router-capability area
            ipv6-routing native
            level-capability 2
            traffic-engineering true
            area-address [49.0001]
            loopfree-alternate {
                remote-lfa {
                }
                ti-lfa {
                }
            }
            traffic-engineering-options {
                ipv6 true
                application-link-attributes {
                }
            }
            segment-routing-v6 {
                admin-state enable
                locator "bng-1-loc" {
                    level-capability 2
                    level 2 {
                        metric 10
                    }
                }
            }
            interface "int-1-bng-1-bng-2" {
            }
            interface "int-1-bng-1-pe-2" {
            }
            interface "system" {
            }
            level 2 {
                wide-metrics-only true
            }
        }
        radius {
            server "free-radius-1" {
                address 192.168.114.2
                secret "HDqTwZYSvEu934VNhUQy/pubZxTKpSDzvHg=" hash2
                accept-coa true
            }
        }
        segment-routing {
            sr-policies {
                admin-state enable
                static-policy "to-pe-1-long-path" {
                    admin-state enable
                    color 10
                    endpoint 2001:db8::5
                    preference 150
                    head-end local
                    type srv6
                    segment-routing-v6 {
                        binding-sid 1 {
                            locator {
                                locator-name "bng-1-loc"
                                function end-b6-encaps-red
                            }
                        }
                    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    segment-list 1 {
                        admin-state enable
                        segment 1 {
                            srv6-sid 2001:db8:aaaa:4:0:1000::
                        }
                        segment 2 {
                            srv6-sid 2001:db8:aaaa:a:0:1000::
                        }
                    }
                }
            }
            segment-routing-v6 {
                origination-fpe [1]
                source-address 2001:db8::14
                locator "bng-1-loc" {
                    admin-state enable
                    block-length 48
                    termination-fpe [2]
                    prefix {
                        ip-prefix 2001:db8:aaaa:14::/64
                    }
                }
                base-routing-instance {
                    locator "bng-1-loc" {
                        function {
                            end 1 {
                                srh-mode usp
                            }
                            end-x-auto-allocate usp protection protected { }
                        }
                    }
                }
            }
        }
        static-routes {
            route 192.0.2.21/32 route-type unicast {
                next-hop "192.168.7.2" {
                    admin-state enable
                }
            }
        }
    }
    service {
        oper-group "ES-1" {
            hold-time {
                up 0
            }
        }
        oper-group "network-port" {
            hold-time {
                up 0
            }
        }
        system {
            bgp {
                evpn {
                    ethernet-segment "ES-1" {
                        admin-state enable
                        esi 0x01010101010101010101
                        orig-ip 2001:db8::14
                        route-next-hop 2001:db8::14
                        multi-homing-mode single-active
                        oper-group "ES-1"
                        df-election {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        es-activation-timer 0
                        service-carving-mode manual
                        manual {
                            preference {
                                mode revertive
                                value 150
                            }
                        }
                    }
                    association {
                        pw-port 1 {
                            pw-port-headend true
                        }
                    }
                }
            }
        }
        epipe "evpn-dual-homing" {
            admin-state enable
            service-id 11
            customer "1"
            segment-routing-v6 1 {
                locator "bng-1-loc" {
                    function {
                        end-dx2 {
                        }
                    }
                }
            }
            bgp 1 {
            }
            bgp-evpn {
                evi 11
                local-attachment-circuit "bng" {
                    eth-tag 2
                }
                remote-attachment-circuit "access" {
                    eth-tag 1
                }
                segment-routing-v6 1 {
                    admin-state enable
                    default-route-tag 0xb
                    source-address 2001:db8::14
                    resolution fallback-tunnel-to-route-table
                    srv6 {
                        instance 1
                        default-locator "bng-1-loc"
                    }
                    route-next-hop {
                        system-ipv6
                    }
                }
            }
        }
        sdp 1 {
            admin-state enable
            signaling off
            far-end {
                ip-address 192.168.7.2
            }
        }
        vpls "capture-sap" {
            admin-state enable
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            service-id 5
            customer "1"
            capture-sap pw-1:*.* {
                radius-auth-policy "radius-1"
                track-srrp 1
                trigger-packet {
                    dhcp true
                    dhcp6 true
                    pppoe true
                }
                ipoe-session {
                    admin-state enable
                    ipoe-session-policy "ipoe-session-policy-1"
                }
                pppoe {
                    policy "ppp-policy-1"
                }
            }
        }
        vprn "dual-homing" {
            admin-state enable
            service-id 10
            customer "1"
            segment-routing-v6 1 {
                locator "bng-1-loc" {
                    function {
                        end-dt4 {
                        }
                        end-dt6 {
                        }
                    }
                }
            }
            bgp-ipvpn {
                segment-routing-v6 1 {
                    admin-state enable
                    route-distinguisher "192.0.2.20:10"
                    source-address 2001:db8::14
                    vrf-target {
                        community "target:64500:10"
                    }
                    vrf-export {
                        policy ["srrp-aware-routing"]
                    }
                    srv6 {
                        instance 1
                        default-locator "bng-1-loc"
                    }
                }
            }
            bgp {
            }
            interface "loopback-1" {
                admin-state enable
                loopback true
                ipv4 {
                    local-dhcp-server "dhcpv4"
                    primary {
                        address 192.168.0.1
                        prefix-length 32
                    }
                    neighbor-discovery {
                        local-proxy-arp false
                        remote-proxy-arp false
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
                dhcp {
                    admin-state enable
                }
            }
            ipv6 {
                local-dhcp-server "dhcpv6"
                address 2001:db8::1 {
                    prefix-length 128
                }
            }
        }
        redundant-interface "red-int-bng-1-bng-2" {
            admin-state enable
            description "static to system ip"
            spoke-sdp 1:1 {
                admin-state enable
                ingress {
                    vc-label 1000
                }
                egress {
                    vc-label 1000
                }
            }
            ipv4 {
                primary {
                    address 192.168.11.1
                    prefix-length 24
                    remote-ip 192.168.11.2
                }
            }
        }
        dhcp-server {
            dhcpv4 "dhcpv4" {
                admin-state enable
                pool-selection {
                    use-gi-address {
                        scope pool
                    }
                    use-pool-from-client {
                    }
                }
                pool "dhcpv4-1" {
                    max-lease-time 1200
                    failover {
                        admin-state enable
                        peer 192.0.2.21 {
                            sync-tag "dhcp4"
                        }
                    }
                    subnet 10.10.0.0/24 {
                        address-range 10.10.0.10 end 10.10.0.100 {
                            failover-control-type access-driven
                        }
                    }
                }
                pool "pppoev4-1" {
                    max-lease-time 1200
                    subnet 10.10.1.0/24 {
                        address-range 10.10.1.10 end 10.10.1.100 {
                            failover-control-type access-driven
                        }
                    }
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
            dhcpv6 "dhcpv6" {
                admin-state enable
                pool-selection {
                    use-pool-from-client {
                    }
                }
                pool "dhcpv6-1" {
                    delegated-prefix {
                        minimum 56
                    }
                    failover {
                        admin-state enable
                        peer 192.0.2.21 {
                            sync-tag "dhcp6"
                        }
                    }
                    prefix 2001:db8:bbbb::/56 {
                        failover-control-type access-driven
                        preferred-lifetime 900
                        valid-lifetime 1200
                        renew-time 600
                        rebind-time 1000
                        prefix-type {
                            wan-host true
                        }
                    }
                    prefix 2001:db8:bbbb:100::/56 {
                        failover-control-type access-driven
                        preferred-lifetime 900
                        valid-lifetime 1200
                        renew-time 600
                        rebind-time 1000
                        prefix-type {
                            pd true
                        }
                    }
                }
            }
        }
        subscriber-interface "sub-int-1" {
            admin-state enable
            ipv4 {
                address 10.10.0.1 {
                    prefix-length 24
                    gateway 10.10.0.254
                    track-srrp 1
                }
                address 10.10.1.1 {
                    prefix-length 24
                    gateway 10.10.1.254
                    track-srrp 1
                }
            }
            ipv6 {
                delegated-prefix-length variable
                prefix 2001:db8:bbbb::/56 {
                    track-srrp 1
                    host-type wan
                }
                prefix 2001:db8:bbbb:100::/56 {
                    track-srrp 1
                    host-type pd
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    link-local-address {
                        address fe80::b696:ffff:fe00:0
                    }
                }
                group-interface "group-int-1" {
                    admin-state enable
                    radius-auth-policy "radius-1"
                    oper-up-while-empty true
                    redundant-interface "red-int-bng-1-bng-2"
                    dynamic-routes-track-srrp {
                    }
                    ipv4 {
                        neighbor-discovery {
                            remote-proxy-arp true
                            populate false
                        }
                        dhcp {
                            admin-state enable
                            server [192.168.0.1]
                            trusted true
                            gi-address 10.10.0.1
                            match-circuit-id true
                            option-82 {
                                action keep
                                vendor-specific-option {
                                    pool-name true
                                }
                            }
                            lease-populate {
                                max-leases 100
                            }
                            client-applications {
                                dhcp true
                            }
                        }
                    }
                    ipv6 {
                        auto-reply {
                            neighbor-solicitation true
                            router-solicitation true
                        }
                        dhcp6 {
                            pd-managed-route {
                            }
                            relay {
                                admin-state enable
                                server ["2001:db8::1"]
                                client-applications {
                                    dhcp true
                                    ppp true
                                }
                            }
                        }
                        router-advertisements {
                            admin-state enable
                            force-mcast ip-mac
                            options {
                                managed-configuration true
                                reachable-time 10000
                                retransmit-timer 3
                            }
                        }
                        router-solicit {
                            admin-state disable
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    }
                }
                ipoe-session {
                    admin-state enable
                    ipoe-session-policy "ipoe-session-policy-1"
                    sap-session-limit 100
                    force-auth {
                        cid-change false
                        rid-change false
                    }
                }
                pppoe {
                    admin-state enable
                    policy "ppp-policy-1"
                    session-limit 100
                    sap-session-limit 100
                }
                local-address-assignment {
                    admin-state enable
                    ipv4 {
                        server "dhcpv4"
                        client-applications {
                            ppp true
                        }
                    }
                }
                srrp 1 {
                    admin-state enable
                    keep-alive-interval 2
                    message-path pw-1:2.4094
                }
                sap pw-1:2.4094 {
                    monitor-oper-group "ES-1"
                }
            }
        }
    }
}
sfm 1 {
    sfm-type m-sfm6-7/12
}
subscriber-mgmt {
    ipoe-session-policy "ipoe-session-policy-1" {
    }
    sub-profile "sub-profile-1" {
    }
    sla-profile "sla-profile-1" {
    }
    sub-ident-policy "sub-ident-policy-1" {
        sla-profile-map {
            use-direct-map-as-default true
        }
        sub-profile-map {
            use-direct-map-as-default true
        }
    }
    ppp-policy "ppp-policy-1" {
        max-sessions-per-mac 100
        allow-same-circuit-id-for-dhcp true
        ncp-renegotiation ignore
    }
    radius-authentication-policy "radius-1" {
        password "ncd8qyrNUMhYfa2SfrUqHMDZ9IXn3sVSmYBzbw==" hash2
        pppoe-access-method pap-chap
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            radius-server-policy "radius-server-1"
            user-name {
                format circuit-id
            }
            include-radius-attribute {
                circuit-id true
                nas-identifier true
            }
        }
        msap-policy "msap-policy-1" {
            sub-sla-mgmt {
                subscriber-limit 100
                sub-ident-policy "sub-ident-policy-1"
                defaults {
                    subscriber-id {
                        sap-id
                    }
                }
            }
            ies-vprn-only-sap-parameters {
                anti-spoof next-hop-ip-and-mac-addr
                ingress {
                    qos {
                        queuing-type service
                    }
                }
            }
        }
    }
    system {
        name "bng-1"
        management-interface {
            configuration-mode model-driven
            cli {
                cli-engine [md-cli classic-cli]
            }
            yang-modules {
                nokia-submodules true
                nokia-combined-modules false
            }
            snmp {
                admin-state disable
            }
        }
        login-control {
            idle-timeout none
        }
        security {
            aaa {
                local-profiles {
                    profile "administrative" {
                        default-action permit-all
                        entry 10 {
                            match "configure system security"
                            action permit
                        }
                        entry 20 {
                            match "show system security"
                            action permit
                        }
                        entry 30 {
                            match "tools perform security"
                            action permit
                        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    entry 40 {
                        match "tools dump security"
                        action permit
                    }
                    entry 50 {
                        match "admin system security"
                        action permit
                    }
                    entry 100 {
                        match "configure li"
                        action deny
                    }
                    entry 110 {
                        match "show li"
                        action deny
                    }
                    entry 111 {
                        match "clear li"
                        action deny
                    }
                    entry 112 {
                        match "tools dump li"
                        action deny
                    }
                    netconf {
                        base-op-authorization {
                            action true
                            cancel-commit true
                            close-session true
                            commit true
                            copy-config true
                            create-subscription true
                            delete-config true
                            discard-changes true
                            edit-config true
                            get true
                            get-config true
                            get-data true
                            get-schema true
                            kill-session true
                            lock true
                            validate true
                        }
                    }
                }
                profile "default" {
                    entry 10 {
                        match "exec"
                        action permit
                    }
                    entry 20 {
                        match "exit"
                        action permit
                    }
                    entry 30 {
                        match "help"
                        action permit
                    }
                    entry 40 {
                        match "logout"
                        action permit
                    }
                    entry 50 {
                        match "password"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        action permit
                    }
                    entry 60 {
                        match "show config"
                        action deny
                    }
                    entry 65 {
                        match "show li"
                        action deny
                    }
                    entry 66 {
                        match "clear li"
                        action deny
                    }
                    entry 67 {
                        match "tools dump li"
                        action deny
                    }
                    entry 68 {
                        match "state li"
                        action deny
                    }
                    entry 70 {
                        match "show"
                        action permit
                    }
                    entry 75 {
                        match "state"
                        action permit
                    }
                    entry 80 {
                        match "enable-admin"
                        action permit
                    }
                    entry 90 {
                        match "enable"
                        action permit
                    }
                    entry 100 {
                        match "configure li"
                        action deny
                    }
                }
            }
        }
        ssh {
            server-cipher-list-v2 {
                cipher 190 {
                    name aes256-ctr
                }
                cipher 192 {
                    name aes192-ctr
                }
                cipher 194 {
                    name aes128-ctr
                }
                cipher 200 {
                    name aes128-cbc
                }
                cipher 205 {
                    name 3des-cbc
                }
                cipher 225 {
                    name aes192-cbc
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    }
                    cipher 230 {
                        name aes256-cbc
                    }
                }
                client-cipher-list-v2 {
                    cipher 190 {
                        name aes256-ctr
                    }
                    cipher 192 {
                        name aes192-ctr
                    }
                    cipher 194 {
                        name aes128-ctr
                    }
                    cipher 200 {
                        name aes128-cbc
                    }
                    cipher 205 {
                        name 3des-cbc
                    }
                    cipher 225 {
                        name aes192-cbc
                    }
                    cipher 230 {
                        name aes256-cbc
                    }
                }
                server-mac-list-v2 {
                    mac 200 {
                        name hmac-sha2-512
                    }
                    mac 210 {
                        name hmac-sha2-256
                    }
                    mac 215 {
                        name hmac-sha1
                    }
                    mac 220 {
                        name hmac-sha1-96
                    }
                    mac 225 {
                        name hmac-md5
                    }
                    mac 240 {
                        name hmac-md5-96
                    }
                }
                client-mac-list-v2 {
                    mac 200 {
                        name hmac-sha2-512
                    }
                    mac 210 {
                        name hmac-sha2-256
                    }
                    mac 215 {
                        name hmac-sha1
                    }
                    mac 220 {
                        name hmac-sha1-96
                    }
                    mac 225 {
                        name hmac-md5
                    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    mac 240 {
                        name hmac-md5-96
                    }
                }
            }
            user-params {
                local-user {
                    user "admin" {
                        password "$2y$10$TQrZlpBDra86.qoexZUzQeBXDY1FcdDhGWdD9lLxMuFyPVSm0OGy6"
                        access {
                            console true
                        }
                        console {
                            member ["administrative"]
                        }
                    }
                }
            }
        }
        time {
            prefer-local-time true
            zone {
                standard {
                    name cst
                }
            }
            dst-zone "CDT" {
                end {
                    day sunday
                    month november
                    hours-minutes "02:00"
                }
                start {
                    day sunday
                    month march
                    hours-minutes "02:00"
                }
            }
            sntp {
                admin-state enable
                server 135.227.160.253 {
                }
            }
        }
    }
}


# Finished 2022-11-01T10:41:30.6-05:00
```

bng-2-config

```
# TiMOS-C-22.10.R1 cpm/x86_64 Nokia 7750 SR Copyright (c) 2000-2022 Nokia.
# All rights reserved. All use subject to applicable license agreements.
# Built on Sun Oct 30 14:49:55 PDT 2022 by builder in /builds/c/2210B/R1/panos/main/sros
# Configuration format version 22.10 revision 0

# Generated 2022-11-01T10:41:32.5-05:00 by admin from 135.231.208.32
# Commit ID 1
#   Committed 2022-11-01T10:04:58.4-05:00 by system (MD-CLI) from Console
#   Log       "System booted version C-22.10.R1."

configure {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    aaa {
        radius {
            server-policy "radius-server-1" {
                servers {
                    router-instance "Base"
                    server 1 {
                        server-name "free-radius-1"
                    }
                }
            }
        }
    }
    card 1 {
        card-type iom5-e
        mda 1 {
            mda-type me6-100gb-qsfp28
        }
        mda 2 {
            mda-type me6-100gb-qsfp28
        }
    }
    fwd-path-ext {
        sdp-id-range {
            start 17500
            end 17600
        }
        fpe 1 {
            description "srv6 origination"
            path {
                pxc 1
            }
            application {
                srv6 {
                    type origination
                }
            }
        }
        fpe 2 {
            description "srv6 termination"
            path {
                pxc 2
            }
            application {
                srv6 {
                    type termination
                }
            }
        }
        fpe 3 {
            description "pw-port on a single port"
            path {
                pxc 3
            }
            application {
                pw-port-extension {
                }
            }
        }
    }
    log {
        filter "1001" {
            named-entry "10" {
                description "Collect only events of major severity or higher"
                action forward
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                match {
                    severity {
                        gte major
                    }
                }
            }
        }
        log-id "100" {
            description "Default Serious Errors Log"
            filter "1001"
            source {
                main true
            }
            destination {
                memory {
                    max-entries 500
                }
            }
        }
        log-id "99" {
            description "Default System Log"
            source {
                main true
            }
            destination {
                memory {
                    max-entries 500
                }
            }
        }
    }
    policy-options {
        community "color-30" {
            member "color:00:30" { }
        }
        community "dual-homing" {
            member "target:64500:10" { }
        }
        policy-statement "pol-color-30" {
            entry 10 {
                from {
                    family [evpn]
                    tag 11
                }
                action {
                    action-type accept
                    community {
                        add ["color-30"]
                    }
                }
            }
        }
        policy-statement "srrp-aware-routing" {
            description "vrf-export; advertizing sub-if routes based on srrp state"
            entry 10 {
                from {
                    state srrp-master
                    protocol {
                        name [direct]
                    }
                }
                action {
                    action-type accept
                    local-preference 150
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    community {
                        add ["dual-homing"]
                    }
                }
            }
            entry 20 {
                from {
                    state srrp-non-master
                    protocol {
                        name [direct]
                    }
                }
                action {
                    action-type accept
                    local-preference 100
                    community {
                        add ["dual-homing"]
                    }
                }
            }
        }
        policy-statement "sub-mgmt-routes" {
            description "this is not applied, but if it was - it would work fine through vrf-
export"
            entry 10 {
                from {
                    protocol {
                        name [sub-mgmt]
                    }
                }
                action {
                    action-type accept
                    community {
                        add ["dual-homing"]
                    }
                }
            }
        }
    }
    port 1/1/c1 {
        admin-state enable
        connector {
            breakout c10-10g
        }
    }
    port 1/1/c1/1 {
        admin-state enable
        description "to p-2 access; mutually exclusive with lag-2"
        monitor-oper-group "network-port"
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/2 {
        admin-state disable
        description "lag-2 to p-2 access; mutually exclusive with port 1/1/c1/1"
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/3 {
        admin-state disable
        description "lag-2 to p-2 access; mutually exclusive with port 1/1/c1/1"
        ethernet {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            mode hybrid
        }
    }
    port 1/1/c1/4 {
        admin-state enable
        description "spare port to p-2 access"
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/5 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/6 {
        admin-state enable
        description "to bng-1"
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/7 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/8 {
        admin-state enable
        description "to pe-2 network"
        oper-group "network-port"
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/9 {
        admin-state enable
        ethernet {
            mode hybrid
            dot1x {
                tunneling true
            }
        }
    }
    port 1/1/c1/10 {
        admin-state enable
        description "RADIUS and mirroring"
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c4 {
        admin-state enable
        connector {
            breakout c4-10g
        }
    }
    port 1/1/c4/1 {
        admin-state enable
        description "PXC 1,2,3; fpe 1,2,3; srv6 orig, term, pw-port single port"
        ethernet {
            mode hybrid
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            dot1x {
                tunneling true
            }
        }
    }
    port 1/1/c4/4 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port pxc-1.a {
        admin-state enable
    }
    port pxc-1.b {
        admin-state enable
    }
    port pxc-2.a {
        admin-state enable
    }
    port pxc-2.b {
        admin-state enable
    }
    port pxc-3.a {
        admin-state enable
    }
    port pxc-3.b {
        admin-state enable
    }
    port-xc {
        pxc 1 {
            admin-state enable
            description "fpe srv6 origin"
            port-id 1/1/c4/1
        }
        pxc 2 {
            admin-state enable
            description "fpe srv6 termination"
            port-id 1/1/c4/1
        }
        pxc 3 {
            admin-state enable
            description "fpe pw-port single port"
            port-id 1/1/c4/1
        }
    }
    pw-port 1 {
        encap-type qinq
        epipe "evpn-dual-homing" {
            admin-state enable
            fpe-id 3
            oper-up-on-mh-standby true
        }
    }
    redundancy {
        multi-chassis {
            peer 192.0.2.20 {
                admin-state enable
                sync {
                    admin-state enable
                    local-dhcp-server true
                    srrp true
                    sub-mgmt {
                        ipoe true
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    pppoe true
                }
                tags {
                    pw-port 1 {
                        sync-tag "pw-port-1"
                    }
                }
                track-srrp 1 {
                }
            }
        }
    }
    bgp-evpn {
        ethernet-segment {
            boot-timer 120
        }
    }
}
router "Base" {
    autonomous-system 64500
    router-id 192.0.2.21
    interface "int-1-bng-2-bng-1" {
        port 1/1/c1/6:1
        gre-termination true
        ipv4 {
            primary {
                address 192.168.7.2
                prefix-length 24
            }
        }
        ipv6 {
            address 2001:db8::702 {
                prefix-length 120
            }
        }
    }
    interface "int-1-bng-2-p-2" {
        port 1/1/c1/1:1
        ipv6 {
            address 2001:db8::601 {
                prefix-length 120
            }
        }
    }
    interface "int-1-bng-2-pe-2" {
        port 1/1/c1/8:1
        ipv6 {
            address 2001:db8::901 {
                prefix-length 120
            }
        }
    }
    interface "system" {
        ipv4 {
            primary {
                address 192.0.2.21
                prefix-length 32
            }
        }
        ipv6 {
            bfd {
                admin-state enable
                transmit-interval 100
                receive 100
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                multiplier 2
            }
            address 2001:db8::15 {
                prefix-length 128
            }
        }
    }
    interface "to-radius" {
        port 1/1/c1/10:114
        ipv4 {
            primary {
                address 192.168.114.21
                prefix-length 24
            }
        }
    }
    bgp {
        admin-state enable
        vpn-apply-export true
        vpn-apply-import true
        rapid-withdrawal true
        rapid-update {
            vpn-ipv4 true
            vpn-ipv6 true
            evpn true
        }
        extended-nh-encoding {
            vpn-ipv4 true
        }
        advertise-ipv6-next-hops {
            vpn-ipv6 true
            vpn-ipv4 true
        }
        group "evpn" {
            peer-as 64500
            local-address 2001:db8::15
            bfd-liveness true
            family {
                evpn true
            }
        }
        group "ipvpn" {
            peer-as 64500
            local-address 2001:db8::15
            bfd-liveness true
            family {
                vpn-ipv4 true
                vpn-ipv6 true
            }
        }
        neighbor "2001:db8::5" {
            group "evpn"
            export {
                policy ["pol-color-30"]
            }
        }
        neighbor "2001:db8::6" {
            group "ipvpn"
        }
        neighbor "2001:db8::14" {
            group "evpn"
        }
    }
    isis 0 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            admin-state enable
            advertise-passive-only false
            advertise-router-capability area
            ipv6-routing native
            level-capability 2
            traffic-engineering true
            area-address [49.0001]
            loopfree-alternate {
                remote-lfa {
                }
                ti-lfa {
                }
            }
            traffic-engineering-options {
                ipv6 true
                application-link-attributes {
                }
            }
            segment-routing-v6 {
                admin-state enable
                locator "bng-2-loc" {
                    level-capability 2
                    level 2 {
                        metric 10
                    }
                }
            }
            interface "int-1-bng-2-p-2" {
            }
            interface "system" {
            }
            level 2 {
                wide-metrics-only true
            }
        }
        isis 1 {
            admin-state enable
            advertise-passive-only false
            advertise-router-capability area
            ipv6-routing native
            level-capability 2
            traffic-engineering true
            area-address [49.0001]
            loopfree-alternate {
                remote-lfa {
                }
                ti-lfa {
                }
            }
            traffic-engineering-options {
                ipv6 true
                application-link-attributes {
                }
            }
            segment-routing-v6 {
                admin-state enable
                locator "bng-2-loc" {
                    level-capability 2
                    level 2 {
                        metric 10
                    }
                }
            }
            interface "int-1-bng-2-bng-1" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
            interface "int-1-bng-2-pe-2" {
            }
            interface "system" {
            }
            level 2 {
                wide-metrics-only true
            }
        }
        radius {
            server "free-radius-1" {
                address 192.168.114.2
                secret "HDqTwZYSvEu934VNhUQy/pubZxTKpSDzvHg=" hash2
                accept-coa true
            }
        }
        segment-routing {
            sr-policies {
                admin-state enable
                static-policy "to-pe-1-long-path" {
                    admin-state enable
                    color 10
                    endpoint 2001:db8::5
                    preference 150
                    head-end local
                    type srv6
                    segment-routing-v6 {
                        binding-sid 1 {
                            locator {
                                locator-name "bng-2-loc"
                                function end-b6-encaps-red
                            }
                        }
                    }
                    segment-list 1 {
                        admin-state enable
                        segment 1 {
                            srv6-sid 2001:db8:aaaa:a:0:1000::
                        }
                        segment 2 {
                            srv6-sid 2001:db8:aaaa:4:0:1000::
                        }
                    }
                }
            }
            segment-routing-v6 {
                origination-fpe [1]
                source-address 2001:db8::15
                locator "bng-2-loc" {
                    admin-state enable
                    block-length 48
                    termination-fpe [2]
                    prefix {
                        ip-prefix 2001:db8:aaaa:15::/64
                    }
                }
                base-routing-instance {
                    locator "bng-2-loc" {
                        function {
                            end 1 {
                                srh-mode usp
                            }
                            end-x-auto-allocate usp protection protected { }
                        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
            }
        }
    }
    static-routes {
        route 192.0.2.20/32 route-type unicast {
            next-hop "192.168.7.1" {
                admin-state enable
            }
        }
    }
}
service {
    oper-group "ES-1" {
        hold-time {
            up 0
        }
    }
    oper-group "network-port" {
        hold-time {
            up 0
        }
    }
    system {
        bgp {
            evpn {
                ethernet-segment "ES-1" {
                    admin-state enable
                    esi 0x01010101010101010101
                    orig-ip 2001:db8::15
                    route-next-hop 2001:db8::15
                    multi-homing-mode single-active
                    oper-group "ES-1"
                    df-election {
                        es-activation-timer 0
                        service-carving-mode manual
                        manual {
                            preference {
                                mode revertive
                                value 50
                            }
                        }
                    }
                    association {
                        pw-port 1 {
                            pw-port-headend true
                        }
                    }
                }
            }
        }
    }
    epipe "evpn-dual-homing" {
        admin-state enable
        service-id 11
        customer "1"
        segment-routing-v6 1 {
            locator "bng-2-loc" {
                function {
                    end-dx2 {
                    }
                }
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            bgp-evpn {
                evi 11
                local-attachment-circuit "bng" {
                    eth-tag 2
                }
                remote-attachment-circuit "access" {
                    eth-tag 1
                }
                segment-routing-v6 1 {
                    admin-state enable
                    default-route-tag 0xb
                    source-address 2001:db8::15
                    resolution fallback-tunnel-to-route-table
                    srv6 {
                        instance 1
                        default-locator "bng-2-loc"
                    }
                    route-next-hop {
                        ip-address 2001:db8::15
                    }
                }
            }
        }
        sdp 1 {
            admin-state enable
            signaling off
            far-end {
                ip-address 192.168.7.1
            }
        }
        vpls "capture-sap" {
            admin-state enable
            service-id 5
            customer "1"
            capture-sap pw-1:*.* {
                radius-auth-policy "radius-1"
                track-srrp 1
                trigger-packet {
                    dhcp true
                    dhcp6 true
                    pppoe true
                }
                ipoe-session {
                    admin-state enable
                    ipoe-session-policy "ipoe-session-policy-1"
                }
                pppoe {
                    policy "ppp-policy-1"
                }
            }
        }
        vprn "dual-homing" {
            admin-state enable
            service-id 10
            customer "1"
            segment-routing-v6 1 {
                locator "bng-2-loc" {
                    function {
                        end-dt4 {
                        }
                        end-dt6 {
                        }
                    }
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
            bgp-ipvpn {
                segment-routing-v6 1 {
                    admin-state enable
                    route-distinguisher "192.0.2.21:10"
                    source-address 2001:db8::15
                    vrf-target {
                        community "target:64500:10"
                    }
                    vrf-export {
                        policy ["srrp-aware-routing"]
                    }
                    srv6 {
                        instance 1
                        default-locator "bng-2-loc"
                    }
                }
            }
            bgp {
            }
            interface "loopback-1" {
                admin-state enable
                loopback true
                ipv4 {
                    local-dhcp-server "dhcpv4"
                    primary {
                        address 192.168.0.1
                        prefix-length 32
                    }
                    neighbor-discovery {
                        local-proxy-arp false
                        remote-proxy-arp false
                    }
                    dhcp {
                        admin-state enable
                    }
                }
                ipv6 {
                    local-dhcp-server "dhcpv6"
                    address 2001:db8::1 {
                        prefix-length 128
                    }
                }
            }
            redundant-interface "red-int-bng-2-bng-1" {
                admin-state enable
                spoke-sdp 1:1 {
                    admin-state enable
                    ingress {
                        vc-label 1000
                    }
                    egress {
                        vc-label 1000
                    }
                }
                ipv4 {
                    primary {
                        address 192.168.11.2
                        prefix-length 24
                        remote-ip 192.168.11.1
                    }
                }
            }
            dhcp-server {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                dhcpv4 "dhcpv4" {
                    admin-state enable
                    pool-selection {
                        use-gi-address {
                            scope pool
                        }
                        use-pool-from-client {
                        }
                    }
                    pool "dhcpv4-1" {
                        max-lease-time 1200
                        failover {
                            admin-state enable
                            peer 192.0.2.20 {
                                sync-tag "dhcp4"
                            }
                        }
                        subnet 10.10.0.0/24 {
                            address-range 10.10.0.10 end 10.10.0.100 {
                                failover-control-type access-driven
                            }
                        }
                    }
                    pool "pppoev4-1" {
                        max-lease-time 1200
                        subnet 10.10.1.0/24 {
                            address-range 10.10.1.10 end 10.10.1.100 {
                                failover-control-type access-driven
                            }
                        }
                    }
                }
                dhcpv6 "dhcpv6" {
                    admin-state enable
                    pool-selection {
                        use-pool-from-client {
                        }
                    }
                    pool "dhcpv6-1" {
                        delegated-prefix {
                            minimum 56
                        }
                        failover {
                            admin-state enable
                            peer 192.0.2.20 {
                                sync-tag "dhcp6"
                            }
                        }
                        prefix 2001:db8:bbbb::/56 {
                            failover-control-type access-driven
                            preferred-lifetime 900
                            valid-lifetime 1200
                            renew-time 600
                            rebind-time 1000
                            prefix-type {
                                wan-host true
                            }
                        }
                        prefix 2001:db8:bbbb:100::/56 {
                            failover-control-type access-driven
                            preferred-lifetime 900
                            valid-lifetime 1200
                            renew-time 600
                            rebind-time 1000
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        prefix-type {
                            pd true
                        }
                    }
                }
            }
        }
        subscriber-interface "sub-int-1" {
            admin-state enable
            ipv4 {
                address 10.10.0.2 {
                    prefix-length 24
                    gateway 10.10.0.254
                    track-srrp 1
                }
                address 10.10.1.2 {
                    prefix-length 24
                    gateway 10.10.1.254
                    track-srrp 1
                }
            }
            ipv6 {
                delegated-prefix-length variable
                prefix 2001:db8:bbbb::/56 {
                    track-srrp 1
                    host-type wan
                }
                prefix 2001:db8:bbbb:100::/56 {
                    track-srrp 1
                    host-type pd
                }
                link-local-address {
                    address fe80::b696:ffff:fe00:0
                }
            }
            group-interface "group-int-1" {
                admin-state enable
                radius-auth-policy "radius-1"
                oper-up-while-empty true
                redundant-interface "red-int-bng-2-bng-1"
                dynamic-routes-track-srrp {
                }
                ipv4 {
                    neighbor-discovery {
                        remote-proxy-arp true
                        populate false
                    }
                    dhcp {
                        admin-state enable
                        server [192.168.0.1]
                        trusted true
                        gi-address 10.10.0.2
                        match-circuit-id true
                        option-82 {
                            action keep
                            vendor-specific-option {
                                pool-name true
                            }
                        }
                        lease-populate {
                            max-leases 100
                        }
                        client-applications {
                            dhcp true
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        }
                    }
                }
                ipv6 {
                    auto-reply {
                        neighbor-solicitation true
                        router-solicitation true
                    }
                    dhcp6 {
                        pd-managed-route {
                        }
                        relay {
                            admin-state enable
                            server ["2001:db8::1"]
                            client-applications {
                                dhcp true
                                ppp true
                            }
                        }
                    }
                    router-advertisements {
                        admin-state enable
                        force-mcast ip-mac
                        options {
                            managed-configuration true
                            reachable-time 10000
                            retransmit-timer 3
                        }
                    }
                    router-solicit {
                        admin-state disable
                    }
                }
                ipoe-session {
                    admin-state enable
                    ipoe-session-policy "ipoe-session-policy-1"
                    sap-session-limit 100
                    force-auth {
                        cid-change false
                        rid-change false
                    }
                }
                pppoe {
                    admin-state enable
                    policy "ppp-policy-1"
                    session-limit 100
                    sap-session-limit 100
                }
                local-address-assignment {
                    admin-state enable
                    ipv4 {
                        server "dhcpv4"
                        client-applications {
                            ppp true
                        }
                    }
                }
                srrp 1 {
                    admin-state enable
                    keep-alive-interval 2
                    message-path pw-1:2.4094
                }
                sap pw-1:2.4094 {
                    monitor-oper-group "ES-1"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
            }
        }
    }
}
sfm 1 {
    sfm-type m-sfm6-7/12
}
subscriber-mgmt {
    ipoe-session-policy "ipoe-session-policy-1" {
    }
    sub-profile "sub-profile-1" {
    }
    sla-profile "sla-profile-1" {
    }
    sub-ident-policy "sub-ident-policy-1" {
        sla-profile-map {
            use-direct-map-as-default true
        }
        sub-profile-map {
            use-direct-map-as-default true
        }
    }
    ppp-policy "ppp-policy-1" {
        max-sessions-per-mac 100
        allow-same-circuit-id-for-dhcp true
        ncp-renegotiation ignore
    }
    radius-authentication-policy "radius-1" {
        password "ncd8qyrNUMhYfa2SfrUqHMDZ9IXn3sVSmYBzbw==" hash2
        pppoe-access-method pap-chap
        radius-server-policy "radius-server-1"
        user-name {
            format circuit-id
        }
        include-radius-attribute {
            circuit-id true
            nas-identifier true
        }
    }
    msap-policy "msap-policy-1" {
        sub-sla-mgmt {
            subscriber-limit 100
            sub-ident-policy "sub-ident-policy-1"
            defaults {
                subscriber-id {
                    sap-id
                }
            }
        }
        ies-vprn-only-sap-parameters {
            anti-spoof next-hop-ip-and-mac-addr
            ingress {
                qos {
                    queuing-type service
                }
            }
        }
    }
}
system {
    name "bng-2"
    management-interface {
        configuration-mode model-driven
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            cli {
                cli-engine [md-cli classic-cli]
            }
            yang-modules {
                nokia-submodules true
                nokia-combined-modules false
            }
            snmp {
                admin-state disable
            }
        }
        login-control {
            idle-timeout none
        }
        security {
            aaa {
                local-profiles {
                    profile "administrative" {
                        default-action permit-all
                        entry 10 {
                            match "configure system security"
                            action permit
                        }
                        entry 20 {
                            match "show system security"
                            action permit
                        }
                        entry 30 {
                            match "tools perform security"
                            action permit
                        }
                        entry 40 {
                            match "tools dump security"
                            action permit
                        }
                        entry 50 {
                            match "admin system security"
                            action permit
                        }
                        entry 100 {
                            match "configure li"
                            action deny
                        }
                        entry 110 {
                            match "show li"
                            action deny
                        }
                        entry 111 {
                            match "clear li"
                            action deny
                        }
                        entry 112 {
                            match "tools dump li"
                            action deny
                        }
                        netconf {
                            base-op-authorization {
                                action true
                                cancel-commit true
                                close-session true
                                commit true
                                copy-config true
                                create-subscription true
                                delete-config true
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    discard-changes true
                    edit-config true
                    get true
                    get-config true
                    get-data true
                    get-schema true
                    kill-session true
                    lock true
                    validate true
                }
            }
        }
        profile "default" {
            entry 10 {
                match "exec"
                action permit
            }
            entry 20 {
                match "exit"
                action permit
            }
            entry 30 {
                match "help"
                action permit
            }
            entry 40 {
                match "logout"
                action permit
            }
            entry 50 {
                match "password"
                action permit
            }
            entry 60 {
                match "show config"
                action deny
            }
            entry 65 {
                match "show li"
                action deny
            }
            entry 66 {
                match "clear li"
                action deny
            }
            entry 67 {
                match "tools dump li"
                action deny
            }
            entry 68 {
                match "state li"
                action deny
            }
            entry 70 {
                match "show"
                action permit
            }
            entry 75 {
                match "state"
                action permit
            }
            entry 80 {
                match "enable-admin"
                action permit
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    }
                    entry 90 {
                        match "enable"
                        action permit
                    }
                    entry 100 {
                        match "configure li"
                        action deny
                    }
                }
            }
        }
        ssh {
            server-cipher-list-v2 {
                cipher 190 {
                    name aes256-ctr
                }
                cipher 192 {
                    name aes192-ctr
                }
                cipher 194 {
                    name aes128-ctr
                }
                cipher 200 {
                    name aes128-cbc
                }
                cipher 205 {
                    name 3des-cbc
                }
                cipher 225 {
                    name aes192-cbc
                }
                cipher 230 {
                    name aes256-cbc
                }
            }
            client-cipher-list-v2 {
                cipher 190 {
                    name aes256-ctr
                }
                cipher 192 {
                    name aes192-ctr
                }
                cipher 194 {
                    name aes128-ctr
                }
                cipher 200 {
                    name aes128-cbc
                }
                cipher 205 {
                    name 3des-cbc
                }
                cipher 225 {
                    name aes192-cbc
                }
                cipher 230 {
                    name aes256-cbc
                }
            }
            server-mac-list-v2 {
                mac 200 {
                    name hmac-sha2-512
                }
                mac 210 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    name hmac-sha2-256
                }
                mac 215 {
                    name hmac-sha1
                }
                mac 220 {
                    name hmac-sha1-96
                }
                mac 225 {
                    name hmac-md5
                }
                mac 240 {
                    name hmac-md5-96
                }
            }
            client-mac-list-v2 {
                mac 200 {
                    name hmac-sha2-512
                }
                mac 210 {
                    name hmac-sha2-256
                }
                mac 215 {
                    name hmac-sha1
                }
                mac 220 {
                    name hmac-sha1-96
                }
                mac 225 {
                    name hmac-md5
                }
                mac 240 {
                    name hmac-md5-96
                }
            }
        }
        user-params {
            local-user {
                user "admin" {
                    password "$2y$10$TQrZlpBDra86.qoexZUzQeBXDY1FcdDhGWdD9lLxMuFyPVSm0OGy6"
                    access {
                        console true
                    }
                    console {
                        member ["administrative"]
                    }
                }
            }
        }
    }
    time {
        prefer-local-time true
        zone {
            standard {
                name cst
            }
        }
        dst-zone "CDT" {
            end {
                day sunday
                month november
                hours-minutes "02:00"
            }
            start {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    day sunday
                    month march
                    hours-minutes "02:00"
                }
            }
            sntp {
                admin-state enable
                server 135.227.160.253 {
                }
            }
        }
    }
}


# Finished 2022-11-01T10:41:32.6-05:00
```

p-1-config

```
# TiMOS-B-22.10.R1 both/x86_64 Nokia 7750 SR Copyright (c) 2000-2022 Nokia.
# All rights reserved. All use subject to applicable license agreements.
# Built on Sun Oct 30 14:49:55 PDT 2022 by builder in /builds/c/2210B/R1/panos/main/sros
# Configuration format version 22.10 revision 0

# Generated 2022-11-01T06:45:18.0Z by admin from 135.231.208.32
# Commit ID 1
#   Committed 2022-11-01T06:08:15.6Z by system (MD-CLI) from Console
#   Log       "System booted version B-22.10.R1."

configure {
    card 1 {
        card-type iom-1
        mda 1 {
            mda-type me6-100gb-qsfp28
        }
        mda 2 {
            mda-type me6-100gb-qsfp28
        }
        fp 1 {
        }
    }
    fwd-path-ext {
        fpe 1 {
            path {
                pxc 1
            }
            application {
                srv6 {
                    type origination
                }
            }
        }
        fpe 2 {
            path {
                pxc 2
            }
            application {
                srv6 {
                    type termination
                }
            }
        }
    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    log {
        filter "1001" {
            named-entry "10" {
                description "Collect only events of major severity or higher"
                action forward
                match {
                    severity {
                        gte major
                    }
                }
            }
        }
        log-id "100" {
            description "Default Serious Errors Log"
            filter "1001"
            source {
                main true
            }
            destination {
                memory {
                    max-entries 500
                }
            }
        }
        log-id "99" {
            description "Default System Log"
            source {
                main true
            }
            destination {
                memory {
                    max-entries 500
                }
            }
        }
    }
    port pxc-1.a {
        admin-state enable
    }
    port pxc-1.b {
        admin-state enable
    }
    port pxc-2.a {
        admin-state enable
    }
    port pxc-2.b {
        admin-state enable
    }
    port 1/1/c1 {
        admin-state enable
        connector {
            breakout c10-10g
        }
    }
    port 1/1/c1/1 {
        admin-state enable
        ethernet {
            mode hybrid
            encap-type dot1q
        }
    }
    port 1/1/c1/2 {
        admin-state enable
        ethernet {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            mode hybrid
        }
    }
    port 1/1/c1/3 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/4 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/5 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/6 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/7 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/8 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/9 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/10 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2 {
        admin-state enable
        connector {
            breakout c10-10g
        }
    }
    port 1/1/c2/1 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/2 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        admin-state disable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/3 {
        admin-state disable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/4 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/5 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/6 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/7 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/8 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/9 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/10 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3 {
        admin-state enable
        connector {
            breakout c10-10g
        }
    }
    port 1/1/c3/1 {
        admin-state enable
        ethernet {
            mode hybrid
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    }
    port 1/1/c3/2 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/3 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/4 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/5 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/6 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/7 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/8 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/9 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/10 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c4 {
        admin-state enable
        connector {
            breakout c4-10g
        }
    }
    port 1/1/c4/1 {
        admin-state enable
        ethernet {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            mode hybrid
            dot1x {
                tunneling true
            }
        }
    }
    port 1/1/c5 {
    }
    port 1/1/c6 {
    }
    port 1/2/c1 {
    }
    port 1/2/c2 {
    }
    port 1/2/c3 {
    }
    port 1/2/c4 {
    }
    port 1/2/c5 {
    }
    port 1/2/c6 {
    }
    port-xc {
        pxc 1 {
            admin-state enable
            port-id 1/1/c4/1
        }
        pxc 2 {
            admin-state enable
            port-id 1/1/c4/1
        }
    }
    router "Base" {
        autonomous-system 64500
        router-id 192.0.2.4
        interface "int-1-p-1-bng-1" {
            port 1/1/c2/1:1
            ipv6 {
                address 2001:db8::502 {
                    prefix-length 120
                }
            }
        }
        interface "int-1-p-1-p-2" {
            port 1/1/c1/9:1
            ipv6 {
                address 2001:db8::a01 {
                    prefix-length 120
                }
            }
        }
        interface "int-1-p-1-pe-1" {
            port 1/1/c1/1:1
            ipv6 {
                address 2001:db8::102 {
                    prefix-length 120
                }
            }
        }
        interface "system" {
            ipv4 {
                primary {
                    address 192.0.2.4
                    prefix-length 32
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
        }
        ipv6 {
            address 2001:db8::4 {
                prefix-length 128
            }
        }
    }
    interface "to-radius" {
        port 1/1/c3/10:114
        ipv4 {
            primary {
                address 192.168.114.4
                prefix-length 24
            }
        }
    }
    bgp {
    }
    isis 0 {
        admin-state enable
        advertise-passive-only false
        advertise-router-capability area
        ipv6-routing native
        level-capability 2
        traffic-engineering true
        area-address [49.0001]
        loopfree-alternate {
            remote-lfa {
            }
            ti-lfa {
            }
        }
        traffic-engineering-options {
            ipv6 true
            application-link-attributes {
            }
        }
        segment-routing-v6 {
            admin-state enable
            locator "p-1-loc" {
                level-capability 2
                level 2 {
                    metric 10
                }
            }
        }
        interface "int-1-p-1-bng-1" {
        }
        interface "int-1-p-1-p-2" {
        }
        interface "int-1-p-1-pe-1" {
        }
        interface "system" {
        }
        level 2 {
            wide-metrics-only true
        }
    }
    segment-routing {
        segment-routing-v6 {
            origination-fpe [1]
            source-address 2001:db8::4
            locator "p-1-loc" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    admin-state enable
                    block-length 48
                    function-length 20
                    termination-fpe [2]
                    prefix {
                        ip-prefix 2001:db8:aaaa:4::/64
                    }
                }
                base-routing-instance {
                    locator "p-1-loc" {
                        function {
                            end 1 {
                                srh-mode usp
                            }
                            end-x-auto-allocate usp protection protected { }
                        }
                    }
                }
            }
        }
    }
    system {
        name "p-1"
        management-interface {
            configuration-mode model-driven
            cli {
                cli-engine [md-cli classic-cli]
            }
            yang-modules {
                nokia-submodules true
                nokia-combined-modules false
            }
            snmp {
                admin-state disable
            }
        }
        bluetooth {
            advertising-timeout 30
        }
        login-control {
            idle-timeout none
        }
        security {
            aaa {
                local-profiles {
                    profile "administrative" {
                        default-action permit-all
                        entry 10 {
                            match "configure system security"
                            action permit
                        }
                        entry 20 {
                            match "show system security"
                            action permit
                        }
                        entry 30 {
                            match "tools perform security"
                            action permit
                        }
                        entry 40 {
                            match "tools dump security"
                            action permit
                        }
                        entry 50 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        match "admin system security"
                        action permit
                    }
                    entry 100 {
                        match "configure li"
                        action deny
                    }
                    entry 110 {
                        match "show li"
                        action deny
                    }
                    entry 111 {
                        match "clear li"
                        action deny
                    }
                    entry 112 {
                        match "tools dump li"
                        action deny
                    }
                    netconf {
                        base-op-authorization {
                            action true
                            cancel-commit true
                            close-session true
                            commit true
                            copy-config true
                            create-subscription true
                            delete-config true
                            discard-changes true
                            edit-config true
                            get true
                            get-config true
                            get-data true
                            get-schema true
                            kill-session true
                            lock true
                            validate true
                        }
                    }
                }
                profile "default" {
                    entry 10 {
                        match "exec"
                        action permit
                    }
                    entry 20 {
                        match "exit"
                        action permit
                    }
                    entry 30 {
                        match "help"
                        action permit
                    }
                    entry 40 {
                        match "logout"
                        action permit
                    }
                    entry 50 {
                        match "password"
                        action permit
                    }
                    entry 60 {
                        match "show config"
                        action deny
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    }
                    entry 65 {
                        match "show li"
                        action deny
                    }
                    entry 66 {
                        match "clear li"
                        action deny
                    }
                    entry 67 {
                        match "tools dump li"
                        action deny
                    }
                    entry 68 {
                        match "state li"
                        action deny
                    }
                    entry 70 {
                        match "show"
                        action permit
                    }
                    entry 75 {
                        match "state"
                        action permit
                    }
                    entry 80 {
                        match "enable-admin"
                        action permit
                    }
                    entry 90 {
                        match "enable"
                        action permit
                    }
                    entry 100 {
                        match "configure li"
                        action deny
                    }
                }
            }
        }
        ssh {
            server-cipher-list-v2 {
                cipher 190 {
                    name aes256-ctr
                }
                cipher 192 {
                    name aes192-ctr
                }
                cipher 194 {
                    name aes128-ctr
                }
                cipher 200 {
                    name aes128-cbc
                }
                cipher 205 {
                    name 3des-cbc
                }
                cipher 225 {
                    name aes192-cbc
                }
                cipher 230 {
                    name aes256-cbc
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                client-cipher-list-v2 {
                    cipher 190 {
                        name aes256-ctr
                    }
                    cipher 192 {
                        name aes192-ctr
                    }
                    cipher 194 {
                        name aes128-ctr
                    }
                    cipher 200 {
                        name aes128-cbc
                    }
                    cipher 205 {
                        name 3des-cbc
                    }
                    cipher 225 {
                        name aes192-cbc
                    }
                    cipher 230 {
                        name aes256-cbc
                    }
                }
                server-mac-list-v2 {
                    mac 200 {
                        name hmac-sha2-512
                    }
                    mac 210 {
                        name hmac-sha2-256
                    }
                    mac 215 {
                        name hmac-sha1
                    }
                    mac 220 {
                        name hmac-sha1-96
                    }
                    mac 225 {
                        name hmac-md5
                    }
                    mac 240 {
                        name hmac-md5-96
                    }
                }
                client-mac-list-v2 {
                    mac 200 {
                        name hmac-sha2-512
                    }
                    mac 210 {
                        name hmac-sha2-256
                    }
                    mac 215 {
                        name hmac-sha1
                    }
                    mac 220 {
                        name hmac-sha1-96
                    }
                    mac 225 {
                        name hmac-md5
                    }
                    mac 240 {
                        name hmac-md5-96
                    }
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            user-params {
                local-user {
                    user "admin" {
                        password "$2y$10$TQrZlpBDra86.qoexZUzQeBXDY1FcdDhGWdD9lLxMuFyPVSm0OGy6"
                        access {
                            console true
                        }
                        console {
                            member ["administrative"]
                        }
                    }
                }
            }
        }
    }
}


# Finished 2022-11-01T06:45:18.0Z
```

p-2-config

```
# TiMOS-B-22.10.R1 both/x86_64 Nokia 7750 SR Copyright (c) 2000-2022 Nokia.
# All rights reserved. All use subject to applicable license agreements.
# Built on Sun Oct 30 14:49:55 PDT 2022 by builder in /builds/c/2210B/R1/panos/main/sros
# Configuration format version 22.10 revision 0

# Generated 2022-11-01T08:50:39.8Z by admin from 135.231.208.32
# Commit ID 1
#   Committed 2022-11-01T08:13:40.0Z by system (MD-CLI) from Console
#   Log       "System booted version B-22.10.R1."

configure {
    card 1 {
        card-type iom-1
        mda 1 {
            mda-type me6-100gb-qsfp28
        }
        mda 2 {
            mda-type me6-100gb-qsfp28
        }
    }
    fwd-path-ext {
        fpe 1 {
            path {
                pxc 1
            }
            application {
                srv6 {
                    type origination
                }
            }
        }
        fpe 2 {
            path {
                pxc 2
            }
            application {
                srv6 {
                    type termination
                }
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    }
    log {
        filter "1001" {
            named-entry "10" {
                description "Collect only events of major severity or higher"
                action forward
                match {
                    severity {
                        gte major
                    }
                }
            }
        }
        log-id "100" {
            description "Default Serious Errors Log"
            filter "1001"
            source {
                main true
            }
            destination {
                memory {
                    max-entries 500
                }
            }
        }
        log-id "99" {
            description "Default System Log"
            source {
                main true
            }
            destination {
                memory {
                    max-entries 500
                }
            }
        }
    }
    port pxc-1.a {
        admin-state enable
    }
    port pxc-1.b {
        admin-state enable
    }
    port pxc-2.a {
        admin-state enable
    }
    port pxc-2.b {
        admin-state enable
    }
    port 1/1/c1 {
        admin-state enable
        connector {
            breakout c10-10g
        }
    }
    port 1/1/c1/1 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/2 {
        admin-state enable
        ethernet {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            mode hybrid
        }
    }
    port 1/1/c1/3 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/4 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/5 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/6 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/7 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/8 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/9 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/10 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2 {
        admin-state enable
        connector {
            breakout c10-10g
        }
    }
    port 1/1/c2/1 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/2 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        admin-state disable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/3 {
        admin-state disable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/4 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/5 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/6 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/7 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/8 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/9 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/10 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3 {
        admin-state enable
        connector {
            breakout c10-10g
        }
    }
    port 1/1/c3/1 {
        admin-state enable
        ethernet {
            mode hybrid
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    }
    port 1/1/c3/2 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/3 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/4 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/5 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/6 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/7 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/8 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/9 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/10 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c4 {
        admin-state enable
        connector {
            breakout c4-10g
        }
    }
    port 1/1/c4/1 {
        admin-state enable
        ethernet {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            mode hybrid
            dot1x {
                tunneling true
            }
        }
    }
    port-xc {
        pxc 1 {
            admin-state enable
            port-id 1/1/c4/1
        }
        pxc 2 {
            admin-state enable
            port-id 1/1/c4/1
        }
    }
    router "Base" {
        autonomous-system 64500
        router-id 192.0.2.10
        interface "int-1-p-2-bng-2" {
            port 1/1/c2/1:1
            ipv6 {
                address 2001:db8::602 {
                    prefix-length 120
                }
            }
        }
        interface "int-1-p-2-p-1" {
            port 1/1/c1/9:1
            ipv6 {
                address 2001:db8::a02 {
                    prefix-length 120
                }
            }
        }
        interface "int-1-p-2-pe-1" {
            port 1/1/c1/5:1
            ipv6 {
                address 2001:db8::202 {
                    prefix-length 120
                }
            }
        }
        interface "system" {
            ipv4 {
                primary {
                    address 192.0.2.10
                    prefix-length 32
                }
            }
            ipv6 {
                address 2001:db8::a {
                    prefix-length 128
                }
            }
        }
        interface "to-radius" {
            port 1/1/c3/10:114
            ipv4 {
                primary {
                    address 192.168.114.10
                    prefix-length 24
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
        isis 0 {
            admin-state enable
            advertise-passive-only false
            advertise-router-capability area
            ipv6-routing native
            level-capability 2
            traffic-engineering true
            area-address [49.0001]
            loopfree-alternate {
                remote-lfa {
                }
                ti-lfa {
                }
            }
            traffic-engineering-options {
                ipv6 true
                application-link-attributes {
                }
            }
            segment-routing-v6 {
                admin-state enable
                locator "p-2-loc" {
                    level-capability 2
                    level 2 {
                        metric 10
                    }
                }
            }
            interface "int-1-p-2-bng-2" {
            }
            interface "int-1-p-2-p-1" {
            }
            interface "int-1-p-2-pe-1" {
            }
            interface "system" {
            }
            level 2 {
                wide-metrics-only true
            }
        }
        segment-routing {
            segment-routing-v6 {
                origination-fpe [1]
                source-address 2001:db8::a
                locator "p-2-loc" {
                    admin-state enable
                    block-length 48
                    termination-fpe [2]
                    prefix {
                        ip-prefix 2001:db8:aaaa:a::/64
                    }
                }
                base-routing-instance {
                    locator "p-2-loc" {
                        function {
                            end 1 {
                                srh-mode usp
                            }
                            end-x-auto-allocate usp protection protected { }
                        }
                    }
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
    }
    system {
        name "p-2"
        management-interface {
            configuration-mode model-driven
            cli {
                cli-engine [md-cli classic-cli]
            }
            yang-modules {
                nokia-submodules true
                nokia-combined-modules false
            }
            snmp {
                admin-state disable
            }
        }
        login-control {
            idle-timeout none
        }
        security {
            aaa {
                local-profiles {
                    profile "administrative" {
                        default-action permit-all
                        entry 10 {
                            match "configure system security"
                            action permit
                        }
                        entry 20 {
                            match "show system security"
                            action permit
                        }
                        entry 30 {
                            match "tools perform security"
                            action permit
                        }
                        entry 40 {
                            match "tools dump security"
                            action permit
                        }
                        entry 50 {
                            match "admin system security"
                            action permit
                        }
                        entry 100 {
                            match "configure li"
                            action deny
                        }
                        entry 110 {
                            match "show li"
                            action deny
                        }
                        entry 111 {
                            match "clear li"
                            action deny
                        }
                        entry 112 {
                            match "tools dump li"
                            action deny
                        }
                        netconf {
                            base-op-authorization {
                                action true
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    cancel-commit true
                    close-session true
                    commit true
                    copy-config true
                    create-subscription true
                    delete-config true
                    discard-changes true
                    edit-config true
                    get true
                    get-config true
                    get-data true
                    get-schema true
                    kill-session true
                    lock true
                    validate true
                }
            }
        }
        profile "default" {
            entry 10 {
                match "exec"
                action permit
            }
            entry 20 {
                match "exit"
                action permit
            }
            entry 30 {
                match "help"
                action permit
            }
            entry 40 {
                match "logout"
                action permit
            }
            entry 50 {
                match "password"
                action permit
            }
            entry 60 {
                match "show config"
                action deny
            }
            entry 65 {
                match "show li"
                action deny
            }
            entry 66 {
                match "clear li"
                action deny
            }
            entry 67 {
                match "tools dump li"
                action deny
            }
            entry 68 {
                match "state li"
                action deny
            }
            entry 70 {
                match "show"
                action permit
            }
            entry 75 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        match "state"
                        action permit
                    }
                    entry 80 {
                        match "enable-admin"
                        action permit
                    }
                    entry 90 {
                        match "enable"
                        action permit
                    }
                    entry 100 {
                        match "configure li"
                        action deny
                    }
                }
            }
        }
        ssh {
            server-cipher-list-v2 {
                cipher 190 {
                    name aes256-ctr
                }
                cipher 192 {
                    name aes192-ctr
                }
                cipher 194 {
                    name aes128-ctr
                }
                cipher 200 {
                    name aes128-cbc
                }
                cipher 205 {
                    name 3des-cbc
                }
                cipher 225 {
                    name aes192-cbc
                }
                cipher 230 {
                    name aes256-cbc
                }
            }
            client-cipher-list-v2 {
                cipher 190 {
                    name aes256-ctr
                }
                cipher 192 {
                    name aes192-ctr
                }
                cipher 194 {
                    name aes128-ctr
                }
                cipher 200 {
                    name aes128-cbc
                }
                cipher 205 {
                    name 3des-cbc
                }
                cipher 225 {
                    name aes192-cbc
                }
                cipher 230 {
                    name aes256-cbc
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
                server-mac-list-v2 {
                    mac 200 {
                        name hmac-sha2-512
                    }
                    mac 210 {
                        name hmac-sha2-256
                    }
                    mac 215 {
                        name hmac-sha1
                    }
                    mac 220 {
                        name hmac-sha1-96
                    }
                    mac 225 {
                        name hmac-md5
                    }
                    mac 240 {
                        name hmac-md5-96
                    }
                }
                client-mac-list-v2 {
                    mac 200 {
                        name hmac-sha2-512
                    }
                    mac 210 {
                        name hmac-sha2-256
                    }
                    mac 215 {
                        name hmac-sha1
                    }
                    mac 220 {
                        name hmac-sha1-96
                    }
                    mac 225 {
                        name hmac-md5
                    }
                    mac 240 {
                        name hmac-md5-96
                    }
                }
            }
            user-params {
                local-user {
                    user "admin" {
                        password "$2y$10$TQrZlpBDra86.qoexZUzQeBXDY1FcdDhGWdD9lLxMuFyPVSm0OGy6"
                        access {
                            console true
                        }
                        console {
                            member ["administrative"]
                        }
                    }
                }
            }
        }
    }
}


# Finished 2022-11-01T08:50:39.8Z
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

pe-1-config

```
# TiMOS-B-22.7.R1 both/x86_64 Nokia 7750 SR Copyright (c) 2000-2022 Nokia.
# All rights reserved. All use subject to applicable license agreements.
# Built on Tue Aug 2 14:18:47 PDT 2022 by builder in /builds/c/227B/R1/panos/main/sros
# Configuration format version 22.7 revision 0

# Generated 2022-10-27T09:42:24.8Z by admin from 135.231.208.32
# Commit ID 1
#   Committed 2022-10-27T09:13:22.5Z by system (MD-CLI) from Console
#   Log       "System booted version B-22.7.R1."

configure {
    card 1 {
        card-type iom-1
        mda 1 {
            mda-type me6-100gb-qsfp28
        }
        mda 2 {
            mda-type me6-100gb-qsfp28
            xconnect {
                mac 1 {
                    loopback 1 {
                    }
                }
            }
        }
        fp 1 {
        }
    }
    connection-profile {
        vlan 1 {
            qtag-range 2 {
                end 2
            }
        }
        vlan 20 {
            qtag-range 20 {
                end 30
            }
        }
    }
    fwd-path-ext {
        fpe 1 {
            path {
                pxc 1
            }
            application {
                srv6 {
                    type origination
                }
            }
        }
        fpe 2 {
            path {
                pxc 2
            }
            application {
                srv6 {
                    type termination
                }
            }
        }
    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    log {
        filter "1001" {
            named-entry "10" {
                description "Collect only events of major severity or higher"
                action forward
                match {
                    severity {
                        gte major
                    }
                }
            }
        }
        log-id "100" {
            description "Default Serious Errors Log"
            filter "1001"
            source {
                main true
            }
            destination {
                memory {
                    max-entries 500
                }
            }
        }
        log-id "50" {
            source {
                debug true
            }
            destination {
                cli {
                    max-entries 1000
                }
            }
        }
        log-id "99" {
            description "Default System Log"
            source {
                main true
            }
            destination {
                memory {
                    max-entries 500
                }
            }
        }
    }
    policy-options {
        community "color-10" {
            member "color:00:10" { }
        }
        policy-statement "pol-color-10" {
            entry 10 {
                from {
                    family [evpn]
                    tag 11
                }
                action {
                    action-type accept
                    community {
                        add ["color-10"]
                    }
                }
            }
            entry 20 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                from {
                    family [evpn]
                    tag 12
                }
                action {
                    action-type accept
                    community {
                        add ["color-10"]
                    }
                }
            }
        }
    }
    port pxc-1.a {
        admin-state enable
    }
    port pxc-1.b {
        admin-state enable
    }
    port pxc-2.a {
        admin-state enable
    }
    port pxc-2.b {
        admin-state enable
    }
    port 1/1/c1 {
        admin-state enable
        connector {
            breakout c10-10g
        }
    }
    port 1/1/c1/1 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/2 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/3 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/4 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/5 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/6 {
        admin-state enable
        ethernet {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            mode hybrid
        }
    }
    port 1/1/c1/7 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/8 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/9 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/10 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2 {
        admin-state enable
        connector {
            breakout c10-10g
        }
    }
    port 1/1/c2/1 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/2 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/3 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/4 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/5 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/6 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/7 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/8 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/9 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/10 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3 {
        admin-state enable
        connector {
            breakout c10-10g
        }
    }
    port 1/1/c3/1 {
        admin-state enable
        ethernet {
            mode hybrid
            encap-type qinq
        }
    }
    port 1/1/c3/2 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/3 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/4 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/5 {
        admin-state enable
        ethernet {
            mode hybrid
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
    }
    port 1/1/c3/6 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/7 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/8 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/9 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/10 {
        admin-state enable
        ethernet {
            mode hybrid
            encap-type dot1q
        }
    }
    port 1/1/c4 {
        admin-state enable
        connector {
            breakout c4-10g
        }
    }
    port 1/1/c4/1 {
        admin-state enable
        ethernet {
            mode hybrid
            dot1x {
                tunneling true
            }
        }
    }
    port 1/1/c5 {
    }
    port 1/1/c6 {
    }
    port 1/2/c1 {
    }
    port 1/2/c2 {
    }
    port 1/2/c3 {
    }
    port 1/2/c4 {
    }
    port 1/2/c5 {
    }
    port 1/2/c6 {
    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    port 1/2/m1/1 {
    }
    port-xc {
        pxc 1 {
            admin-state enable
            port-id 1/1/c4/1
        }
        pxc 2 {
            admin-state enable
            port-id 1/1/c4/1
        }
    }
    router "Base" {
        autonomous-system 64500
        router-id 192.0.2.5
        interface "int-1-pe-1-p-1" {
            port 1/1/c1/1:1
            ipv6 {
                address 2001:db8::101 {
                    prefix-length 120
                }
            }
        }
        interface "int-1-pe-1-p-2" {
            port 1/1/c1/5:1
            ipv6 {
                address 2001:db8::201 {
                    prefix-length 120
                }
            }
        }
        interface "system" {
            ipv4 {
                primary {
                    address 192.0.2.5
                    prefix-length 32
                }
            }
            ipv6 {
                bfd {
                    admin-state enable
                    transmit-interval 100
                    receive 100
                    multiplier 2
                }
                address 2001:db8::5 {
                    prefix-length 128
                }
            }
        }
        interface "to-ixia" {
            port 1/1/c3/1:1.1
            ipv4 {
                primary {
                    address 172.16.100.1
                    prefix-length 24
                }
            }
            ipv6 {
                address 2001:db8::6401 {
                    prefix-length 120
                }
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        interface "to-radius" {
            port 1/1/c3/10:114
            ipv4 {
                primary {
                    address 192.168.114.5
                    prefix-length 24
                }
            }
        }
        bgp {
            admin-state enable
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            rapid-update {
                evpn true
            }
            advertise-ipv6-next-hops {
                evpn true
            }
            group "evpn" {
                peer-as 64500
                local-address 2001:db8::5
                bfd-liveness true
                family {
                    evpn true
                }
            }
            neighbor "2001:db8::14" {
                group "evpn"
                export {
                    policy ["pol-color-10"]
                }
            }
            neighbor "2001:db8::15" {
                group "evpn"
                export {
                    policy ["pol-color-10"]
                }
            }
        }
        isis 0 {
            admin-state enable
            advertise-passive-only false
            advertise-router-capability area
            ipv6-routing native
            level-capability 2
            traffic-engineering true
            area-address [49.0001]
            loopfree-alternate {
                remote-lfa {
                }
                ti-lfa {
                }
            }
            traffic-engineering-options {
                ipv6 true
                application-link-attributes {
                }
            }
            segment-routing-v6 {
                admin-state enable
                locator "pe-1-loc" {
                    level-capability 2
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                level 2 {
                    metric 10
                }
            }
        }
        interface "int-1-pe-1-p-1" {
        }
        interface "int-1-pe-1-p-2" {
        }
        interface "system" {
        }
        interface "to-ixia" {
        }
        level 2 {
            wide-metrics-only true
        }
    }
    segment-routing {
        sr-policies {
            admin-state enable
            static-policy "to-bng-1-long-path" {
                admin-state enable
                color 20
                endpoint 2001:db8::14
                preference 150
                head-end local
                type srv6
                segment-routing-v6 {
                    binding-sid 1 {
                        locator {
                            locator-name "pe-1-loc"
                            function end-b6-encaps-red
                        }
                    }
                }
                segment-list 1 {
                    admin-state enable
                    segment 1 {
                        srv6-sid 2001:db8:aaaa:a:0:1000::
                    }
                    segment 2 {
                        srv6-sid 2001:db8:aaaa:4:0:1000::
                    }
                }
            }
            static-policy "to-bng-2-long-path" {
                admin-state enable
                color 30
                endpoint 2001:db8::15
                preference 150
                head-end local
                type srv6
                segment-routing-v6 {
                    binding-sid 1 {
                        locator {
                            locator-name "pe-1-loc"
                            function end-b6-encaps-red
                        }
                    }
                }
                segment-list 1 {
                    admin-state enable
                    segment 1 {
                        srv6-sid 2001:db8:aaaa:4:0:1000::
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    }
                    segment 2 {
                        srv6-sid 2001:db8:aaaa:a:0:1000::
                    }
                }
            }
        }
        segment-routing-v6 {
            origination-fpe [1]
            source-address 2001:db8::5
            locator "pe-1-loc" {
                admin-state enable
                block-length 48
                termination-fpe [2]
                prefix {
                    ip-prefix 2001:db8:aaaa:5::/64
                }
                static-function {
                    max-entries 5
                }
            }
            base-routing-instance {
                locator "pe-1-loc" {
                    function {
                        end 1 {
                            srh-mode usp
                        }
                        end-x-auto-allocate usp protection protected { }
                    }
                }
            }
        }
    }
}
service {
    system {
        extended-default-qinq-sap-lookup true
    }
    epipe "dual-homing" {
        admin-state enable
        service-id 11
        customer "1"
        segment-routing-v6 1 {
            locator "pe-1-loc" {
                function {
                    end-dx2 {
                    }
                }
            }
        }
        sap 1/1/c3/1:*.* {
        }
        bgp-evpn {
            evi 11
            local-attachment-circuit "access" {
                eth-tag 1
            }
            remote-attachment-circuit "bng" {
                eth-tag 2
            }
            segment-routing-v6 1 {
                admin-state enable
                default-route-tag 0xb
                source-address 2001:db8::5
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    resolution fallback-tunnel-to-route-table
                    srv6 {
                        instance 1
                        default-locator "pe-1-loc"
                    }
                    route-next-hop {
                        system-ipv6
                    }
                }
            }
        }
    }
    system {
        name "pe-1"
        management-interface {
            configuration-mode model-driven
            cli {
                cli-engine [md-cli classic-cli]
            }
            yang-modules {
                nokia-submodules true
                nokia-combined-modules false
            }
            snmp {
                admin-state disable
            }
        }
        ip {
            allow-qinq-network-interface true
        }
        bluetooth {
            advertising-timeout 30
        }
        login-control {
            idle-timeout none
        }
        security {
            aaa {
                local-profiles {
                    profile "administrative" {
                        default-action permit-all
                        entry 10 {
                            match "configure system security"
                            action permit
                        }
                        entry 20 {
                            match "show system security"
                            action permit
                        }
                        entry 30 {
                            match "tools perform security"
                            action permit
                        }
                        entry 40 {
                            match "tools dump security"
                            action permit
                        }
                        entry 50 {
                            match "admin system security"
                            action permit
                        }
                        entry 100 {
                            match "configure li"
                            action deny
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    }
                    entry 110 {
                        match "show li"
                        action deny
                    }
                    entry 111 {
                        match "clear li"
                        action deny
                    }
                    entry 112 {
                        match "tools dump li"
                        action deny
                    }
                    netconf {
                        base-op-authorization {
                            action true
                            cancel-commit true
                            close-session true
                            commit true
                            copy-config true
                            create-subscription true
                            delete-config true
                            discard-changes true
                            edit-config true
                            get true
                            get-config true
                            get-data true
                            get-schema true
                            kill-session true
                            lock true
                            validate true
                        }
                    }
                }
                profile "default" {
                    entry 10 {
                        match "exec"
                        action permit
                    }
                    entry 20 {
                        match "exit"
                        action permit
                    }
                    entry 30 {
                        match "help"
                        action permit
                    }
                    entry 40 {
                        match "logout"
                        action permit
                    }
                    entry 50 {
                        match "password"
                        action permit
                    }
                    entry 60 {
                        match "show config"
                        action deny
                    }
                    entry 65 {
                        match "show li"
                        action deny
                    }
                    entry 66 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        match "clear li"
                        action deny
                    }
                    entry 67 {
                        match "tools dump li"
                        action deny
                    }
                    entry 68 {
                        match "state li"
                        action deny
                    }
                    entry 70 {
                        match "show"
                        action permit
                    }
                    entry 75 {
                        match "state"
                        action permit
                    }
                    entry 80 {
                        match "enable-admin"
                        action permit
                    }
                    entry 90 {
                        match "enable"
                        action permit
                    }
                    entry 100 {
                        match "configure li"
                        action deny
                    }
                    netconf {
                        base-op-authorization {
                            action true
                            cancel-commit true
                            close-session true
                            commit true
                            copy-config true
                            create-subscription true
                            delete-config true
                            discard-changes true
                            edit-config true
                            get true
                            get-config true
                            get-data true
                            get-schema true
                            validate true
                        }
                    }
                }
            }
        }
        ssh {
            server-cipher-list-v2 {
                cipher 190 {
                    name aes256-ctr
                }
                cipher 192 {
                    name aes192-ctr
                }
                cipher 194 {
                    name aes128-ctr
                }
                cipher 200 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        name aes128-cbc
                    }
                    cipher 205 {
                        name 3des-cbc
                    }
                    cipher 225 {
                        name aes192-cbc
                    }
                    cipher 230 {
                        name aes256-cbc
                    }
                }
                client-cipher-list-v2 {
                    cipher 190 {
                        name aes256-ctr
                    }
                    cipher 192 {
                        name aes192-ctr
                    }
                    cipher 194 {
                        name aes128-ctr
                    }
                    cipher 200 {
                        name aes128-cbc
                    }
                    cipher 205 {
                        name 3des-cbc
                    }
                    cipher 225 {
                        name aes192-cbc
                    }
                    cipher 230 {
                        name aes256-cbc
                    }
                }
                server-mac-list-v2 {
                    mac 200 {
                        name hmac-sha2-512
                    }
                    mac 210 {
                        name hmac-sha2-256
                    }
                    mac 215 {
                        name hmac-sha1
                    }
                    mac 220 {
                        name hmac-sha1-96
                    }
                    mac 225 {
                        name hmac-md5
                    }
                    mac 240 {
                        name hmac-md5-96
                    }
                }
                client-mac-list-v2 {
                    mac 200 {
                        name hmac-sha2-512
                    }
                    mac 210 {
                        name hmac-sha2-256
                    }
                    mac 215 {
                        name hmac-sha1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    }
                    mac 220 {
                        name hmac-sha1-96
                    }
                    mac 225 {
                        name hmac-md5
                    }
                    mac 240 {
                        name hmac-md5-96
                    }
                }
            }
            user-params {
                local-user {
                    user "admin" {
                        password "$2y$10$TQrZlpBDra86.qoexZUzQeBXDY1FcdDhGWdD9lLxMuFyPVSm0OGy6"
                        access {
                            console true
                        }
                        console {
                            member ["administrative"]
                        }
                    }
                }
            }
        }
    }
}


# Finished 2022-10-27T09:42:24.8Z
```

pe-2-config

```
# TiMOS-B-22.10.R1 both/x86_64 Nokia 7750 SR Copyright (c) 2000-2022 Nokia.
# All rights reserved. All use subject to applicable license agreements.
# Built on Sun Oct 30 14:49:55 PDT 2022 by builder in /builds/c/2210B/R1/panos/main/sros
# Configuration format version 22.10 revision 0

# Generated 2022-11-01T08:08:00.1Z by admin from 135.231.208.32
# Commit ID 1
#   Committed 2022-11-01T07:31:07.3Z by system (MD-CLI) from Console
#   Log       "System booted version B-22.10.R1."

configure {
    card 1 {
        card-type iom-1
        mda 1 {
            mda-type me6-100gb-qsfp28
        }
        mda 2 {
            mda-type me6-100gb-qsfp28
        }
        fp 1 {
        }
    }
    fwd-path-ext {
        fpe 1 {
            path {
                pxc 1
            }
            application {
                srv6 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    type origination
                }
            }
        }
        fpe 2 {
            path {
                pxc 2
            }
            application {
                srv6 {
                    type termination
                }
            }
        }
    }
    log {
        filter "1001" {
            named-entry "10" {
                description "Collect only events of major severity or higher"
                action forward
                match {
                    severity {
                        gte major
                    }
                }
            }
        }
        log-id "100" {
            description "Default Serious Errors Log"
            filter "1001"
            source {
                main true
            }
            destination {
                memory {
                    max-entries 500
                }
            }
        }
        log-id "99" {
            description "Default System Log"
            source {
                main true
            }
            destination {
                memory {
                    max-entries 500
                }
            }
        }
    }
    port pxc-1.a {
        admin-state enable
    }
    port pxc-1.b {
        admin-state enable
    }
    port pxc-2.a {
        admin-state enable
    }
    port pxc-2.b {
        admin-state enable
    }
    port 1/1/c1 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        admin-state enable
        connector {
            breakout c10-10g
        }
    }
    port 1/1/c1/1 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/2 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/3 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/4 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/5 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/6 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/7 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/8 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/9 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/10 {
        admin-state enable
        ethernet {
            mode hybrid
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    }
    port 1/1/c2 {
        admin-state enable
        connector {
            breakout c10-10g
        }
    }
    port 1/1/c2/1 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/2 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/3 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/4 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/5 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/6 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/7 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/8 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/9 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/10 {
        admin-state enable
        ethernet {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            mode hybrid
        }
    }
    port 1/1/c3 {
        admin-state enable
        connector {
            breakout c10-10g
        }
    }
    port 1/1/c3/1 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/2 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/3 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/4 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/5 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/6 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/7 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/8 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/9 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/10 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c4 {
        admin-state enable
        connector {
            breakout c4-10g
        }
    }
    port 1/1/c4/1 {
        admin-state enable
        ethernet {
            mode hybrid
            dot1x {
                tunneling true
            }
        }
    }
    port 1/1/c5 {
    }
    port 1/1/c6 {
    }
    port 1/2/c1 {
    }
    port 1/2/c2 {
    }
    port 1/2/c3 {
    }
    port 1/2/c4 {
    }
    port 1/2/c5 {
    }
    port 1/2/c6 {
    }
    port-xc {
        pxc 1 {
            admin-state enable
            port-id 1/1/c4/1
        }
        pxc 2 {
            admin-state enable
            port-id 1/1/c4/1
        }
    }
    router "Base" {
        autonomous-system 64500
        router-id 192.0.2.6
        interface "int-1-pe-2-bng-1" {
            port 1/1/c1/7:1
            ipv6 {
                address 2001:db8::802 {
                    prefix-length 120
                }
            }
        }
        interface "int-1-pe-2-bng-2" {
            port 1/1/c1/8:1
            ipv6 {
                address 2001:db8::902 {
                    prefix-length 120
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
        interface "system" {
            ipv4 {
                primary {
                    address 192.0.2.6
                    prefix-length 32
                }
            }
            ipv6 {
                bfd {
                    admin-state enable
                    transmit-interval 100
                    receive 100
                    multiplier 2
                }
                address 2001:db8::6 {
                    prefix-length 128
                }
            }
        }
        interface "to-ixia" {
            port 1/1/c3/1:1
            ipv4 {
                primary {
                    address 172.16.101.1
                    prefix-length 24
                }
            }
            ipv6 {
                address 2001:db8::6501 {
                    prefix-length 120
                }
            }
        }
        interface "to-radius" {
            port 1/1/c3/10:114
            ipv4 {
                primary {
                    address 192.168.114.6
                    prefix-length 24
                }
            }
        }
        bgp {
            admin-state enable
            bfd-liveness true
            rapid-withdrawal true
            rapid-update {
                vpn-ipv4 true
                vpn-ipv6 true
            }
            extended-nh-encoding {
                vpn-ipv4 true
            }
            advertise-ipv6-next-hops {
                vpn-ipv6 true
                vpn-ipv4 true
            }
            group "ipvpn" {
                peer-as 64500
                local-address 2001:db8::6
                bfd-liveness true
                family {
                    vpn-ipv4 true
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    vpn-ipv6 true
                }
            }
            neighbor "2001:db8::14" {
                group "ipvpn"
            }
            neighbor "2001:db8::15" {
                group "ipvpn"
            }
        }
        isis 1 {
            admin-state enable
            advertise-passive-only false
            advertise-router-capability area
            ipv6-routing native
            level-capability 2
            traffic-engineering true
            area-address [49.0001]
            loopfree-alternate {
                remote-lfa {
                }
                ti-lfa {
                }
            }
            traffic-engineering-options {
                ipv6 true
                application-link-attributes {
                }
            }
            segment-routing-v6 {
                admin-state enable
                locator "pe-2-loc" {
                    level-capability 2
                    level 2 {
                        metric 10
                    }
                }
            }
            interface "int-1-pe-2-bng-1" {
            }
            interface "int-1-pe-2-bng-2" {
            }
            interface "system" {
            }
            interface "to-ixia" {
            }
            level 2 {
                wide-metrics-only true
            }
        }
        segment-routing {
            segment-routing-v6 {
                origination-fpe [1]
                source-address 2001:db8::6
                locator "pe-2-loc" {
                    admin-state enable
                    block-length 48
                    termination-fpe [2]
                    prefix {
                        ip-prefix 2001:db8:aaaa:6::/64
                    }
                }
                base-routing-instance {
                    locator "pe-2-loc" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        function {
                            end 1 {
                                srh-mode usp
                            }
                            end-x-auto-allocate usp protection protected { }
                        }
                    }
                }
            }
        }
    }
    service {
        vprn "dual-homing" {
            admin-state enable
            service-id 10
            customer "1"
            segment-routing-v6 1 {
                locator "pe-2-loc" {
                    function {
                        end-dt4 {
                        }
                        end-dt6 {
                        }
                    }
                }
            }
            bgp-ipvpn {
                segment-routing-v6 1 {
                    admin-state enable
                    route-distinguisher "192.0.2.6:10"
                    source-address 2001:db8::6
                    vrf-target {
                        community "target:64500:10"
                    }
                    srv6 {
                        instance 1
                        default-locator "pe-2-loc"
                    }
                }
            }
            interface "vprn-10-to-ixia" {
                ipv4 {
                    primary {
                        address 172.16.102.1
                        prefix-length 24
                    }
                }
                sap 1/1/c3/1:2 {
                }
                ipv6 {
                    address 2001:db8::6601 {
                        prefix-length 120
                    }
                }
            }
            ipv6 {
                router-advertisement {
                    interface "vprn-10-to-ixia" {
                        admin-state enable
                        max-advertisement-interval 15
                        min-advertisement-interval 10
                    }
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
    }
    system {
        name "pe-2"
        management-interface {
            configuration-mode model-driven
            cli {
                cli-engine [md-cli classic-cli]
            }
            yang-modules {
                nokia-submodules true
                nokia-combined-modules false
            }
            snmp {
                admin-state disable
            }
        }
        bluetooth {
            advertising-timeout 30
        }
        login-control {
            idle-timeout none
        }
        security {
            aaa {
                local-profiles {
                    profile "administrative" {
                        default-action permit-all
                        entry 10 {
                            match "configure system security"
                            action permit
                        }
                        entry 20 {
                            match "show system security"
                            action permit
                        }
                        entry 30 {
                            match "tools perform security"
                            action permit
                        }
                        entry 40 {
                            match "tools dump security"
                            action permit
                        }
                        entry 50 {
                            match "admin system security"
                            action permit
                        }
                        entry 100 {
                            match "configure li"
                            action deny
                        }
                        entry 110 {
                            match "show li"
                            action deny
                        }
                        entry 111 {
                            match "clear li"
                            action deny
                        }
                        entry 112 {
                            match "tools dump li"
                            action deny
                        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        netconf {
                            base-op-authorization {
                                action true
                                cancel-commit true
                                close-session true
                                commit true
                                copy-config true
                                create-subscription true
                                delete-config true
                                discard-changes true
                                edit-config true
                                get true
                                get-config true
                                get-data true
                                get-schema true
                                kill-session true
                                lock true
                                validate true
                            }
                        }
                    }
                    profile "default" {
                        entry 10 {
                            match "exec"
                            action permit
                        }
                        entry 20 {
                            match "exit"
                            action permit
                        }
                        entry 30 {
                            match "help"
                            action permit
                        }
                        entry 40 {
                            match "logout"
                            action permit
                        }
                        entry 50 {
                            match "password"
                            action permit
                        }
                        entry 60 {
                            match "show config"
                            action deny
                        }
                        entry 65 {
                            match "show li"
                            action deny
                        }
                        entry 66 {
                            match "clear li"
                            action deny
                        }
                        entry 67 {
                            match "tools dump li"
                            action deny
                        }
                        entry 68 {
                            match "state li"
                            action deny
                        }
                        entry 70 {
                            match "show"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        action permit
                    }
                    entry 75 {
                        match "state"
                        action permit
                    }
                    entry 80 {
                        match "enable-admin"
                        action permit
                    }
                    entry 90 {
                        match "enable"
                        action permit
                    }
                    entry 100 {
                        match "configure li"
                        action deny
                    }
                }
            }
        }
        ssh {
            server-cipher-list-v2 {
                cipher 190 {
                    name aes256-ctr
                }
                cipher 192 {
                    name aes192-ctr
                }
                cipher 194 {
                    name aes128-ctr
                }
                cipher 200 {
                    name aes128-cbc
                }
                cipher 205 {
                    name 3des-cbc
                }
                cipher 225 {
                    name aes192-cbc
                }
                cipher 230 {
                    name aes256-cbc
                }
            }
            client-cipher-list-v2 {
                cipher 190 {
                    name aes256-ctr
                }
                cipher 192 {
                    name aes192-ctr
                }
                cipher 194 {
                    name aes128-ctr
                }
                cipher 200 {
                    name aes128-cbc
                }
                cipher 205 {
                    name 3des-cbc
                }
                cipher 225 {
                    name aes192-cbc
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    cipher 230 {
                        name aes256-cbc
                    }
                }
                server-mac-list-v2 {
                    mac 200 {
                        name hmac-sha2-512
                    }
                    mac 210 {
                        name hmac-sha2-256
                    }
                    mac 215 {
                        name hmac-sha1
                    }
                    mac 220 {
                        name hmac-sha1-96
                    }
                    mac 225 {
                        name hmac-md5
                    }
                    mac 240 {
                        name hmac-md5-96
                    }
                }
                client-mac-list-v2 {
                    mac 200 {
                        name hmac-sha2-512
                    }
                    mac 210 {
                        name hmac-sha2-256
                    }
                    mac 215 {
                        name hmac-sha1
                    }
                    mac 220 {
                        name hmac-sha1-96
                    }
                    mac 225 {
                        name hmac-md5
                    }
                    mac 240 {
                        name hmac-md5-96
                    }
                }
            }
            user-params {
                local-user {
                    user "admin" {
                        password "$2y$10$TQrZlpBDra86.qoexZUzQeBXDY1FcdDhGWdD9lLxMuFyPVSm0OGy6"
                        access {
                            console true
                        }
                        console {
                            member ["administrative"]
                        }
                    }
                }
            }
        }
    }
}


# Finished 2022-11-01T08:08:00.1Z
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

users.radius

```
#
#    Configuration file for the rlm_files module.
#    Please see rlm_files(5) manpage for more information.
#
#    This file contains authentication security and configuration
#    information for each user.  Accounting requests are NOT processed
#    through this file.  Instead, see 'accounting', in this directory.
#
#    The first field is the user's name and can be up to
#    253 characters in length.  This is followed (on the same line) with
#    the list of authentication requirements for that user.  This can
#    include password, comm server name, comm server port number, protocol
#    type (perhaps set by the "hints" file), and huntgroup name (set by
#    the "huntgroups" file).
#
#    If you are not sure why a particular reply is being sent by the
#    server, then run the server in debugging mode (radiusd -X), and
#    you will see which entries in this file are matched.
#
#    When an authentication request is received from the comm server,
#    these values are tested. Only the first match is used unless the
#    "Fall-Through" variable is set to "Yes".
#
#    A special user named "DEFAULT" matches on all usernames.
#    You can have several DEFAULT entries. All entries are processed
#    in the order they appear in this file. The first entry that
#    matches the login-request will stop processing unless you use
#    the Fall-Through variable.
#
#    Indented (with the tab character) lines following the first
#    line indicate the configuration values to be passed back to
#    the comm server to allow the initiation of a user session.
#    This can include things like the PPP configuration values
#    or the host to log the user onto.
#
#    You can include another `users' file with `$INCLUDE users.other'

#
#    For a list of RADIUS attributes, and links to their definitions,
#    see: http://www.freeradius.org/rfc/attributes.html
#
#    Entries below this point are examples included in the server for
#    educational purposes. They may be deleted from the deployed
#    configuration without impacting the operation of the server.
#

############################################################
# You should add test accounts to the TOP of this file! #
# See the example user "bob" above.                     #
############################################################


cid-1           Cleartext-Password := "cse-password"
                    Alc-Subsc-ID-Str = "ipoe-ds-1",
            Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
            Framed-IPv6-Pool = "dhcpv6-1",
            Alc-Delegated-IPv6-Pool = "dhcpv6-1",
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            Alc-Int-Dest-Id-Str = "vport-1",
                        Fall-Through = No

cid-2           Cleartext-Password := "cse-password"
                    Alc-Subsc-ID-Str = "ipoe-ds-2",
            Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
            Framed-IPv6-Pool = "dhcpv6-1",
            Alc-Delegated-IPv6-Pool = "dhcpv6-1",
            Alc-Int-Dest-Id-Str = "vport-1",
                        Fall-Through = No

cid-3           Cleartext-Password := "cse-password"
                    Alc-Subsc-ID-Str = "ipoe-ds-3",
            Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
            Framed-IPv6-Pool = "dhcpv6-1",
            Alc-Delegated-IPv6-Pool = "dhcpv6-1",
            Alc-Int-Dest-Id-Str = "vport-1",
                        Fall-Through = No

cid-4           Cleartext-Password := "cse-password"
                    Alc-Subsc-ID-Str = "ipoe-ds-4",
            Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
            Framed-IPv6-Pool = "dhcpv6-1",
            Alc-Delegated-IPv6-Pool = "dhcpv6-1",
            Alc-Int-Dest-Id-Str = "vport-1",
                        Fall-Through = No

cid-5           Cleartext-Password := "cse-password"
                    Alc-Subsc-ID-Str = "ipoe-ds-5",
            Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
            Framed-IPv6-Pool = "dhcpv6-1",
            Alc-Delegated-IPv6-Pool = "dhcpv6-1",
            Alc-Int-Dest-Id-Str = "vport-1",
                        Fall-Through = No

cid-6           Cleartext-Password := "cse-password"
                    Alc-Subsc-ID-Str = "ipoe-ds-6",
            Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
            Framed-IPv6-Pool = "dhcpv6-1",
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            Alc-Delegated-IPv6-Pool = "dhcpv6-1",
            Alc-Int-Dest-Id-Str = "vport-1",
                        Fall-Through = No

cid-7           Cleartext-Password := "cse-password"
                    Alc-Subsc-ID-Str = "ipoe-ds-7",
            Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
            Framed-IPv6-Pool = "dhcpv6-1",
            Alc-Delegated-IPv6-Pool = "dhcpv6-1",
            Alc-Int-Dest-Id-Str = "vport-1",
                        Fall-Through = No

cid-8           Cleartext-Password := "cse-password"
                    Alc-Subsc-ID-Str = "ipoe-ds-8",
            Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
            Framed-IPv6-Pool = "dhcpv6-1",
            Alc-Delegated-IPv6-Pool = "dhcpv6-1",
            Alc-Int-Dest-Id-Str = "vport-1",
                        Fall-Through = No

cid-9           Cleartext-Password := "cse-password"
                    Alc-Subsc-ID-Str = "ipoe-ds-9",
            Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
            Framed-IPv6-Pool = "dhcpv6-1",
            Alc-Delegated-IPv6-Pool = "dhcpv6-1",
            Alc-Int-Dest-Id-Str = "vport-1",
                        Fall-Through = No

cid-10          Cleartext-Password := "cse-password"
                    Alc-Subsc-ID-Str = "ipoe-ds-10",
            Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
            Framed-IPv6-Pool = "dhcpv6-1",
            Alc-Delegated-IPv6-Pool = "dhcpv6-1",
            Alc-Int-Dest-Id-Str = "vport-2",
                        Fall-Through = No

pppoe-1         Cleartext-Password := "cse-password"
                    Alc-Subsc-ID-Str = "pppoe-1",
            Alc-Subsc-Prof-Str = "sub-profile-1",
                    Alc-SLA-Prof-Str = "sla-profile-1",
                    Alc-MSAP-Interface = "group-int-1",
                    Alc-MSAP-Policy = "msap-policy-1",
                    Alc-MSAP-Serv-Id = "10",
                    Framed-Pool = "pppoev4-1",
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            Framed-IPv6-Pool = "dhcpv6-1",
            Alc-Delegated-IPv6-Pool = "dhcpv6-1",
            Alc-Int-Dest-Id-Str = "vport-1",
                    Fall-Through = No

pppoe-2         Cleartext-Password := "cse-password"
                    Alc-Subsc-ID-Str = "pppoe-2",
            Alc-Subsc-Prof-Str = "sub-profile-1",
                    Alc-SLA-Prof-Str = "sla-profile-1",
                    Alc-MSAP-Interface = "group-int-1",
                    Alc-MSAP-Policy = "msap-policy-1",
                    Alc-MSAP-Serv-Id = "10",
                    Framed-Pool = "pppoev4-1",
            Framed-IPv6-Pool = "dhcpv6-1",
            Alc-Delegated-IPv6-Pool = "dhcpv6-1",
            Alc-Int-Dest-Id-Str = "vport-1",
                    Fall-Through = No

pppoe-3         Cleartext-Password := "cse-password"
                    Alc-Subsc-ID-Str = "pppoe-3",
            Alc-Subsc-Prof-Str = "sub-profile-1",
                    Alc-SLA-Prof-Str = "sla-profile-1",
                    Alc-MSAP-Interface = "group-int-1",
                    Alc-MSAP-Policy = "msap-policy-1",
                    Alc-MSAP-Serv-Id = "10",
                    Framed-Pool = "pppoev4-1",
            Framed-IPv6-Pool = "dhcpv6-1",
            Alc-Delegated-IPv6-Pool = "dhcpv6-1",
            Alc-Int-Dest-Id-Str = "vport-1",
                    Fall-Through = No

pppoe-4         Cleartext-Password := "cse-password"
                    Alc-Subsc-ID-Str = "pppoe-4",
            Alc-Subsc-Prof-Str = "sub-profile-1",
                    Alc-SLA-Prof-Str = "sla-profile-1",
                    Alc-MSAP-Interface = "group-int-1",
                    Alc-MSAP-Policy = "msap-policy-1",
                    Alc-MSAP-Serv-Id = "10",
                    Framed-Pool = "pppoev4-1",
            Framed-IPv6-Pool = "dhcpv6-1",
            Alc-Delegated-IPv6-Pool = "dhcpv6-1",
            Alc-Int-Dest-Id-Str = "vport-1",
                    Fall-Through = No

pppoe-5         Cleartext-Password := "cse-password"
                    Alc-Subsc-ID-Str = "pppoe-5",
            Alc-Subsc-Prof-Str = "sub-profile-1",
                    Alc-SLA-Prof-Str = "sla-profile-1",
                    Alc-MSAP-Interface = "group-int-1",
                    Alc-MSAP-Policy = "msap-policy-1",
                    Alc-MSAP-Serv-Id = "10",
                    Framed-Pool = "pppoev4-1",
            Framed-IPv6-Pool = "dhcpv6-1",
            Alc-Delegated-IPv6-Pool = "dhcpv6-1",
            Alc-Int-Dest-Id-Str = "vport-1",
                    Fall-Through = No

pppoe-6         Cleartext-Password := "cse-password"
                    Alc-Subsc-ID-Str = "pppoe-6",
            Alc-Subsc-Prof-Str = "sub-profile-1",
                    Alc-SLA-Prof-Str = "sla-profile-1",
                    Alc-MSAP-Interface = "group-int-1",
                    Alc-MSAP-Policy = "msap-policy-1",
                    Alc-MSAP-Serv-Id = "10",
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    Framed-Pool = "pppoev4-1",
                Framed-IPv6-Pool = "dhcpv6-1",
                Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                Alc-Int-Dest-Id-Str = "vport-1",
                    Fall-Through = No

pppoe-7             Cleartext-Password := "cse-password"
                    Alc-Subsc-ID-Str = "pppoe-7",
                Alc-Subsc-Prof-Str = "sub-profile-1",
                    Alc-SLA-Prof-Str = "sla-profile-1",
                    Alc-MSAP-Interface = "group-int-1",
                    Alc-MSAP-Policy = "msap-policy-1",
                    Alc-MSAP-Serv-Id = "10",
                    Framed-Pool = "pppoev4-1",
                Framed-IPv6-Pool = "dhcpv6-1",
                Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                Alc-Int-Dest-Id-Str = "vport-1",
                    Fall-Through = No

pppoe-8             Cleartext-Password := "cse-password"
                    Alc-Subsc-ID-Str = "pppoe-8",
                Alc-Subsc-Prof-Str = "sub-profile-1",
                    Alc-SLA-Prof-Str = "sla-profile-1",
                    Alc-MSAP-Interface = "group-int-1",
                    Alc-MSAP-Policy = "msap-policy-1",
                    Alc-MSAP-Serv-Id = "10",
                    Framed-Pool = "pppoev4-1",
                Framed-IPv6-Pool = "dhcpv6-1",
                Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                Alc-Int-Dest-Id-Str = "vport-1",
                    Fall-Through = No

pppoe-9             Cleartext-Password := "cse-password"
                    Alc-Subsc-ID-Str = "pppoe-9",
                Alc-Subsc-Prof-Str = "sub-profile-1",
                    Alc-SLA-Prof-Str = "sla-profile-1",
                    Alc-MSAP-Interface = "group-int-1",
                    Alc-MSAP-Policy = "msap-policy-1",
                    Alc-MSAP-Serv-Id = "10",
                    Framed-Pool = "pppoev4-1",
                Framed-IPv6-Pool = "dhcpv6-1",
                Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                Alc-Int-Dest-Id-Str = "vport-1",
                    Fall-Through = No

pppoe-10            Cleartext-Password := "cse-password"
                    Alc-Subsc-ID-Str = "pppoe-10",
                Alc-Subsc-Prof-Str = "sub-profile-1",
                    Alc-SLA-Prof-Str = "sla-profile-1",
                    Alc-MSAP-Interface = "group-int-1",
                    Alc-MSAP-Policy = "msap-policy-1",
                    Alc-MSAP-Serv-Id = "10",
                    Framed-Pool = "pppoev4-1",
                Framed-IPv6-Pool = "dhcpv6-1",
                Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                Alc-Int-Dest-Id-Str = "vport-2",
                    Fall-Through = No
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Show commands

In this section, the following show commands are used:

```
# FPE
====
show  fwd-path-ext fpe 1 associations
show  fwd-path-ext fpe 1

# EVPN
====
show service id "evpn-dual-homing" base
show service id "evpn-dual-homing" bgp-evpn
show service system bgp-evpn ethernet-segment name "ES-1"
show service system bgp-evpn ethernet-segment name "ES-1" evi evi-1 11
show router bgp routes evpn eth-seg detail
show router bgp routes evpn auto-disc tag 1 detail
show  service  id "evpn-dual-homing" ethernet-segment "ES-1"
show  service id "evpn-dual-homing"  segment-routing-v6 detail
show service id "evpn-dual-homing" segment-routing-v6 instance 1 destinations
show service id "evpn-dual-homing" segment-routing-v6 instance 1 end-dx2
show  service id "evpn-dual-homing"  segment-routing-v6 instance 1 locator "bng-1-loc"

# IP-VPN
====
show service id "dual-homing" base
show service id "dual-homing" bgp-ipvpn segment-routing-v6
show service id "dual-homing" segment-routing-v6 instance 1 locator "bng-2-loc"
show  router 10 route-table
show  router 10 route-table ipv6
show router bgp routes vpn-ipv4 hunt
show router bgp routes vpn-ipv6 hunt

# SRv6
====
show router segment-routing-v6 base-routing-instance all
show router segment-routing-v6 local-sid context  "Base"
show router segment-routing-v6 local-sid context  "10"

# GRT router and tunnel table
====
show router route-table ipv6
show router tunnel-table ipv6
show router fp-tunnel-table 1 ipv6

# BGP
====
show router bgp  summary

# MCS
====
show redundancy multi-chassis sync peer 192.0.2.20 detail
tools dump redundancy multi-chassis sync-database

# subscribers:
====
show service active-subscribers summary
show service active-subscribers hierarchy
show service id 10 subscriber-hosts detail
show service id 10 ipoe session
show service id 10 pppoe session
show service id 10 dhcp lease-state
show service id 10 dhcp6 lease-state
show srrp 1 detail
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
# DHCP server:
====
show router 10 dhcp local-dhcp-server "dhcpv4" summary
show router 10 dhcp6 local-dhcp-server "dhcpv6" summary
show router 10 dhcp local-dhcp-server "dhcpv4" leases
show router 10 dhcp6 local-dhcp-server "dhcpv6" leases

#
===========
show port
show router interface
show router 10 interface
show  fwd-path-ext fpe 1
show  fwd-path-ext fpe 2
show  fwd-path-ext fpe 3
show service id "evpn-dual-homing" base
show service id "evpn-dual-homing" bgp-evpn
show service system bgp-evpn ethernet-segment name "ES-1"
show service system bgp-evpn ethernet-segment name "ES-1" evi evi-1 11
show router bgp routes evpn eth-seg detail
show router bgp routes evpn auto-disc tag 1 detail
show  service  id "evpn-dual-homing" ethernet-segment "ES-1"
show  service id "evpn-dual-homing"  segment-routing-v6 detail
show service id "evpn-dual-homing" segment-routing-v6 instance 1 destinations
show service id "evpn-dual-homing" segment-routing-v6 instance 1 end-dx2
show  service id "evpn-dual-homing" segment-routing-v6 instance 1 locator "bng-1-loc"
show service id "dual-homing" base
show service id "dual-homing" bgp-ipvpn segment-routing-v6
show service id "dual-homing" segment-routing-v6 instance 1 locator "bng-1-loc"
show router 10 route-table
show router 10 route-table ipv6
show router bgp routes vpn-ipv4 hunt
show router bgp routes vpn-ipv6 hunt
show router segment-routing-v6 base-routing-instance all
show router segment-routing-v6 local-sid context "Base"
show router segment-routing-v6 local-sid context "10"
show router route-table ipv6
show router tunnel-table ipv6
show router fp-tunnel-table 1 ipv6
show router  bgp summary
show router isis adjacency
show router isis database
show router isis 1 adjacency
show router isis 1 database
show redundancy multi-chassis sync peer 192.0.2.21 detail
tools dump redundancy multi-chassis sync-database
show  service  active-subscribers summary
show  service  active-subscribers hierarchy
show  service id 10 subscriber-hosts detail
show  service id 10 ipoe session
show  service id 10 pppoe session
show  service id 10 dhcp lease-state
show  service id 10 dhcp6 lease-state
show  srrp 1 detail
show  router 10 dhcp local-dhcp-server "dhcpv4" summary
show  router 10 dhcp6 local-dhcp-server "dhcpv6" summary
show  router 10 dhcp local-dhcp-server "dhcpv4" leases
show  router 10 dhcp6 local-dhcp-server "dhcpv6" leases

#
====
show port
show router interface
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
show router 10 interface
show fwd-path-ext fpe 1
show fwd-path-ext fpe 2
show fwd-path-ext fpe 3
show service id "evpn-dual-homing" base
show service id "evpn-dual-homing" bgp-evpn
show service system bgp-evpn ethernet-segment name "ES-1"
show service system bgp-evpn ethernet-segment name "ES-1" evi evi-1 11
show router bgp routes evpn eth-seg detail
show router bgp routes evpn auto-disc tag 1 detail
show service  id "evpn-dual-homing" ethernet-segment "ES-1"
show service id "evpn-dual-homing"  segment-routing-v6 detail
show service id "evpn-dual-homing" segment-routing-v6 instance 1 destinations
show service id "evpn-dual-homing" segment-routing-v6 instance 1 end-dx2
show service id "evpn-dual-homing"  segment-routing-v6 instance 1 locator "bng-2-loc"
show service id "dual-homing" base
show service id "dual-homing" bgp-ipvpn segment-routing-v6
show service id "dual-homing" segment-routing-v6 instance 1 locator "bng-2-loc"
show router 10 route-table
show router 10 route-table ipv6
show router bgp routes vpn-ipv4 hunt
show router bgp routes vpn-ipv6 hunt
show router segment-routing-v6 base-routing-instance all
show router segment-routing-v6 local-sid context  "Base"
show router segment-routing-v6 local-sid context  "10"
show router route-table ipv6
show router tunnel-table ipv6
show router fp-tunnel-table 1 ipv6
show router  bgp  summary
show router isis  adjacency
show router isis database
show router isis 1 adjacency
show router isis 1 database
show router isis 1 database
show redundancy multi-chassis sync peer 192.0.2.20 detail
tools dump redundancy multi-chassis sync-database
show service  active-subscribers summary
show service  active-subscribers hierarchy
show service id 10 subscriber-hosts detail
show service id 10 ipoe session
show service id 10 pppoe session
show service id 10 dhcp lease-state
show service id 10 dhcp6 lease-state
show srrp 1 detail
show router 10 dhcp local-dhcp-server "dhcpv4" summary
show router 10 dhcp6 local-dhcp-server "dhcpv6" summary
show router 10 dhcp local-dhcp-server "dhcpv4" leases
show router 10 dhcp6 local-dhcp-server "dhcpv6" leases

# on PE-1:
====
show port
show router interface
show fwd-path-ext fpe 1
show fwd-path-ext fpe 2
show service id "dual-homing" base
show service id "dual-homing" bgp-evpn
show router bgp routes evpn auto-disc tag 2 detail
show service id "dual-homing"  segment-routing-v6 detail
show service id "dual-homing" segment-routing-v6 instance 1 destinations
show service id "dual-homing" segment-routing-v6 instance 1 end-dx2
show service id "dual-homing"  segment-routing-v6 instance 1 locator "pe-1-loc"
show router segment-routing-v6 base-routing-instance all
show router segment-routing-v6 local-sid context  "Base"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
show router route-table ipv6
show router tunnel-table ipv6
show router fp-tunnel-table 1 ipv6
show router  bgp  summary
show router isis adjacency
show router isis database

# on PE-2:
=====
show port
show router interface
show router 10 interface
show fwd-path-ext fpe 1
show fwd-path-ext fpe 2
show service id "dual-homing" base
show service id "dual-homing" bgp-ipvpn segment-routing-v6
show service id "dual-homing" segment-routing-v6 instance 1 locator "pe-2-loc"
show router 10 route-table
show router 10 route-table ipv6
show router bgp routes vpn-ipv4 hunt
show router bgp routes vpn-ipv6 hunt
show router segment-routing-v6 base-routing-instance all
show router segment-routing-v6 local-sid context  "Base"
show router segment-routing-v6 local-sid context  "10"
show router route-table ipv6
show router tunnel-table ipv6
show router fp-tunnel-table 1 ipv6
show router  bgp  summary
show router isis 1 adjacency
show router isis 1 database

# on P-1:
====
show port
show router interface
show fwd-path-ext fpe 1
show fwd-path-ext fpe 2
show router segment-routing-v6 base-routing-instance all
show router segment-routing-v6 local-sid context  "Base"
show router route-table ipv6
show router tunnel-table ipv6
show router fp-tunnel-table 1 ipv6
show router isis adjacency
show router isis database

# on P-2:
====
show port
show router interface
show fwd-path-ext fpe 1
show fwd-path-ext fpe 2
show router segment-routing-v6 base-routing-instance all
show router segment-routing-v6 local-sid context  "Base"
show router route-table ipv6
show router tunnel-table ipv6
show router fp-tunnel-table 1 ipv6
show router isis adjacency
show router isis database
```

The output of several show commands is provided for the following nodes:

- bng-1
- p-1

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- pe-1

Show commands on bng-1

```
A:admin@bng-1# /show port

===============================================================================
Ports on Slot 1
===============================================================================
Port          Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id            State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
1/1/c1        Up         Link Up                        conn   100GBASE-LR4*
1/1/c1/1      Up    Yes  Up      8936 8936    - hybr dotq xgige
1/1/c1/2      Down  No   Down    8936 8936    - hybr dotq xgige
1/1/c1/3      Down  No   Down    8936 8936    - hybr dotq xgige
1/1/c1/4      Up    Yes  Up      8936 8936    - hybr dotq xgige
1/1/c1/5      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c1/6      Up    Yes  Up      8936 8936    - hybr dotq xgige
1/1/c1/7      Up    Yes  Up      8936 8936    - hybr dotq xgige
1/1/c1/8      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c1/9      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c1/10     Up    Yes  Up      8936 8936    - hybr dotq xgige
1/1/c2        Up         Link Up                        conn   100G CWDM4 M*
1/1/c2/1      Up    Yes  Up      8936 8936    - hybr dotq cgige
1/1/c3        Down       Down                           conn   100G CLR4   *
1/1/c4        Up         Link Up                        conn   100GBASE-LR4*
1/1/c4/1      Up    Yes  Up      8936 8936    - hybr dotq xgige
1/1/c4/2      Up    No   Down    8936 8936    - netw null xgige
1/1/c4/3      Up    No   Down    8936 8936    - netw null xgige
1/1/c4/4      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c5        Down       Down                           conn   100G CWDM4 M*
1/1/c6        Down       Down                           conn   100G CLR4   *
1/2/c1        Down       Down                           conn   100GBASE-LR4*
1/2/c2        Down       Down                           conn   100G CWDM4 M*
1/2/c3        Down       Down                           conn   100G CLR4   *
1/2/c4        Down       Down                           conn   100GBASE-LR4*
1/2/c5        Down       Down                           conn   100G CWDM4 M*
1/2/c6        Down       Down                           conn   100G CLR4   *

===============================================================================
Ports on Slot A
===============================================================================
Port          Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id            State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
A/1           Up    Yes  Up      1514 1514    - netw null faste  MDI
A/3           Down  No   Down    1514 1514    - netw null faste
A/4           Down  No   Down    1514 1514    - netw null faste

===============================================================================
Ports on Slot B
===============================================================================
Port          Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id            State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
B/1           Up    No   Ghost   1514 1514    - netw null faste
B/3           Down  No   Ghost   1514 1514    - netw null faste
B/4           Down  No   Ghost   1514 1514    - netw null faste

===============================================================================
Ports on Port Cross Connect 1
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Port          Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id            State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
pxc-1.a       Up    Yes  Up      8932 8932    - hybr dotq xgige
pxc-1.b       Up    Yes  Up      8932 8932    - hybr dotq xgige

===============================================================================
Ports on Port Cross Connect 2
===============================================================================
Port          Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id            State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
pxc-2.a       Up    Yes  Up      8932 8932    - hybr dotq xgige
pxc-2.b       Up    Yes  Up      8932 8932    - hybr dotq xgige

===============================================================================
Ports on Port Cross Connect 3
===============================================================================
Port          Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id            State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
pxc-3.a       Up    Yes  Up      8932 8932    - hybr dotq xgige
pxc-3.b       Up    Yes  Up      8932 8932    - hybr dotq xgige
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show router interface

===============================================================================
Interface Table (Router: Base)
===============================================================================
Interface-Name                   Adm       Opr(v4/v6)  Mode    Port/SapId
   IP-Address                                                   PfxState
-------------------------------------------------------------------------------
_tmnx_fpe_1.a                    Up        Up/Up       Network pxc-1.a:1
   fe80::100/64                                                 PREFERRED
_tmnx_fpe_1.b                    Up        Up/Up       Network pxc-1.b:1
   fe80::101/64                                                 PREFERRED
_tmnx_fpe_2.a                    Up        Up/Up       Network pxc-2.a:1
   fe80::200/64                                                 PREFERRED
_tmnx_fpe_2.b                    Up        Up/Up       Network pxc-2.b:1
   fe80::201/64                                                 PREFERRED
_tmnx_fpe_3.a                    Up        Up/Up       Network pxc-3.a:1
   fe80::300/64                                                 PREFERRED
_tmnx_fpe_3.b                    Up        Up/Up       Network pxc-3.b:1
   fe80::301/64                                                 PREFERRED
int-1-bng-1-bng-2                Up        Up/Up       Network 1/1/c1/6:1
   192.168.7.1/24                                               n/a
   2001:db8::701/120                                            PREFERRED
   fe80::b696:ffff:fe00:0/64                                    PREFERRED
int-1-bng-1-p-1                  Up        Down/Up     Network 1/1/c1/1:1
   2001:db8::501/120                                            PREFERRED
   fe80::b696:ffff:fe00:0/64                                    PREFERRED
int-1-bng-1-pe-2                 Up        Down/Up     Network 1/1/c1/7:1
   2001:db8::801/120                                            PREFERRED
   fe80::b696:ffff:fe00:0/64                                    PREFERRED
system                           Up        Up/Up       Network system
   192.0.2.20/32                                                n/a
   2001:db8::14/128                                             PREFERRED
to-radius                        Up        Up/Down     Network 1/1/c1/10:114
   192.168.114.20/24                                            n/a
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
Interfaces : 11
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show router 10 interface

===============================================================================
Interface Table (Service: 10)
===============================================================================
Interface-Name                   Adm       Opr(v4/v6)  Mode    Port/SapId
   IP-Address                                                   PfxState
-------------------------------------------------------------------------------
group-int-1                      Up        Up/Up       VPRN G* pw-1
loopback-1                       Up        Up/Up       VPRN    loopback
   192.168.0.1/32                                               n/a
   2001:db8::1/128                                              PREFERRED
   fe80::b696:ffff:fe00:0/64                                    PREFERRED
red-int-bng-1-bng-2              Up        Up/Up       VPRN R* spoke-1:1
   192.168.11.1/24                                              n/a
sub-int-1                        Up        Up/Up       VPRN S* subscriber
   10.10.0.1/24                                                 n/a
   10.10.1.1/24                                                 n/a
   2001:db8:bbbb::/56                                           PREFERRED
   2001:db8:bbbb:100::/56                                       PREFERRED
   fe80::b696:ffff:fe00:0/64                                    PREFERRED
-------------------------------------------------------------------------------
Interfaces : 4
===============================================================================
* indicates that the corresponding row element may have been truncated.

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  fwd-path-ext fpe 1

===============================================================================
FPE Id: 1
===============================================================================
Description        : srv6 origination
Multi-Path         : Disabled
Path               : pxc 1
Pw Port Extension  : Disabled                          Oper     : down
Sub Mgmt Extension : Disabled                          Oper     : N/A
Vxlan              : Disabled                          Oper     : down
Segment-Routing V6 : Enabled                           Oper     : up
SRv6 Type          : origination
If-A Qos Policy    : default
If-B MTU           : 9786 bytes                        Oper MTU : 8914 bytes
If-B Qos Policy    : default
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  fwd-path-ext fpe 2

===============================================================================
FPE Id: 2
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Description        : srv6 termination
Multi-Path         : Disabled
Path               : pxc 2
Pw Port Extension  : Disabled                          Oper     : down
Sub Mgmt Extension : Disabled                          Oper     : N/A
Vxlan              : Disabled                          Oper     : down
Segment-Routing V6 : Enabled                           Oper     : up
SRv6 Type          : termination
If-A Qos Policy    : default
If-B MTU           : 0 bytes                           Oper MTU : 8914 bytes
If-B Qos Policy    : default
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  fwd-path-ext fpe 3

===============================================================================
FPE Id: 3
===============================================================================
Description        : pw-port on a single port
Multi-Path         : Disabled
Path               : pxc 3
Pw Port Extension  : Enabled                           Oper     : up
Sub Mgmt Extension : Disabled                          Oper     : N/A
Vxlan              : Disabled                          Oper     : down
Segment-Routing V6 : Disabled
If-A Qos Policy    : default
If-B Qos Policy    : default
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show service id "evpn-dual-homing" base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 11                 Vpn Id            : 0
Service Type      : Epipe
MACSec enabled    : no
Name              : evpn-dual-homing
Description       : (Not Specified)
Customer Id       : 1                  Creation Origin   : manual
Last Status Change: 11/12/2022 16:22:36
Last Mgmt Change  : 11/12/2022 16:21:35
Test Service      : No
Admin State       : Up                 Oper State        : Up
MTU               : 1514
Vc Switching      : False
SAP Count         : 0                  SDP Bind Count    : 1
Per Svc Hashing   : Disabled           Lbl Eth/IP L4 TEID: Disabled
Ignore MTU Mismat*: Disabled
Vxlan Src Tep Ip  : N/A
Force QTag Fwd    : Disabled
Lcl Switch Svc St : sap
Oper Group        : <none>

-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                               Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sdp:17500:1 SB(fpe_3.b)                  Fpe          0       8910    Up   Up
===============================================================================
* indicates that the corresponding row element may have been truncated.

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show service id "evpn-dual-homing" bgp-evpn

===============================================================================
BGP EVPN Table
===============================================================================
EVI                : 11                Creation Origin    : manual

-------------------------------------------------------------------------------
Local AC Name              Eth Tag  Endpoint                   Ingress Label
-------------------------------------------------------------------------------
bng                        2                                   0
-------------------------------------------------------------------------------
Number of local ACs : 1

-------------------------------------------------------------------------------
Remote AC Name                     Eth Tag  Endpoint
-------------------------------------------------------------------------------
access                             1
-------------------------------------------------------------------------------
Number of Remote ACs : 1
===============================================================================

===============================================================================
Segment Routing v6 Instance 1 Service 11
===============================================================================
Admin State             : Enabled
Srv6 Instance           : 1
Default Locator         : bng-1-loc

Oper Group              : (Not Specified)
Default Route Tag       : 0xb
Source Address          : 2001:db8::14
ECMP                    : 1
Force Vlan VC Fwd       : disabled
Next Hop Type           : system-ipv6
Evi 3-byte Auto-RT      : disabled
Route Resolution        : fallback-tunnel-to-route-table
Force QinQ VC Fwd       : none
MH Mode                 : network
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show service system bgp-evpn ethernet-segment name "ES-1"

===============================================================================
Service Ethernet Segment
===============================================================================
Name                    : ES-1
Eth Seg Type            : None
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Admin State             : Enabled            Oper State         : Up
ESI                     : 01:01:01:01:01:01:01:01:01:01
Oper ESI                : 01:01:01:01:01:01:01:01:01:01
Auto-ESI Type           : None
AC DF Capability        : Include
Multi-homing            : singleActive       Oper Multi-homing  : singleActive
ES SHG Label            : None
Source BMAC LSB         : None
PW Port Id              : 1
PW Port Headend         : enabled
ES Activation Timer     : 0 secs
Oper Group              : ES-1
Svc Carving             : manual             Oper Svc Carving   : manual
Cfg Range Type          : lowest-pref

-------------------------------------------------------------------------------
DF Pref Election Information
-------------------------------------------------------------------------------
Preference     Preference     Last Admin Change        Oper Pref      Do No
Mode           Value                                   Value          Preempt
-------------------------------------------------------------------------------
revertive      150            11/12/2022 16:21:35      150            Disabled
-------------------------------------------------------------------------------
EVI Ranges: <none>
ISID Ranges: <none>
Vprn NextHop EVI Ranges : <none>
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show service system bgp-evpn ethernet-segment name "ES-1" evi evi-1 11

===============================================================================
EVI DF and Candidate List
===============================================================================
EVI           SvcId         Actv Timer Rem      DF  DF Last Change
-------------------------------------------------------------------------------
11            11            0                   yes 11/12/2022 16:21:35
===============================================================================

===============================================================================
DF Candidates                           Time Added           Oper Pref  Do Not
                                                               Value    Preempt
-------------------------------------------------------------------------------
2001:db8::14                            11/12/2022 16:23:33  150        Disabl*
2001:db8::15                            11/12/2022 16:22:36  50         Disabl*
-------------------------------------------------------------------------------
Number of entries: 2
===============================================================================
* indicates that the corresponding row element may have been truncated.

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show router bgp routes evpn eth-seg detail
===============================================================================
 BGP Router ID:192.0.2.20       AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Eth-Seg Routes
===============================================================================
Original Attributes

Network        : n/a
Nexthop        : 2001:db8::15
Path Id        : None
From           : 2001:db8::15
Res. Nexthop   : fe80::b697:ffff:fe00:0
Local Pref.    : 100                    Interface Name : int-1-bng-1-bng-2
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : 10
Connector      : None
Community      :
                 df-election::DF-Type:Preference/DP:0/DF-Preference:50/AC:1
                 target:01:01:01:01:01:01
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.21
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
EVPN type      : ETH-SEG
ESI            : 01:01:01:01:01:01:01:01:01:01
Originator IP  : 2001:db8::15
Route Dist.    : 192.0.2.21:0
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
Add Paths Send : Default
Last Modified  : 01d23h12m

Modified Attributes

Network        : n/a
Nexthop        : 2001:db8::15
Path Id        : None
From           : 2001:db8::15
Res. Nexthop   : fe80::b697:ffff:fe00:0
Local Pref.    : 100                    Interface Name : int-1-bng-1-bng-2
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : 10
Connector      : None
Community      :
                 df-election::DF-Type:Preference/DP:0/DF-Preference:50/AC:1
                 target:01:01:01:01:01:01
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.21
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
EVPN type      : ETH-SEG
ESI            : 01:01:01:01:01:01:01:01:01:01
Originator IP  : 2001:db8::15
Route Dist.    : 192.0.2.21:0
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Add Paths Send : Default
Last Modified  : 01d23h12m

-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Routes : 1
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show router bgp routes evpn auto-disc tag 1 detail
===============================================================================
 BGP Router ID:192.0.2.20       AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Auto-Disc Routes
===============================================================================
Original Attributes

Network        : n/a
Nexthop        : 2001:db8::5
Path Id        : None
From           : 2001:db8::5
Res. Nexthop   : fe80::b69e:ffff:fe00:0
Local Pref.    : 100                    Interface Name : int-1-bng-1-p-1
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : 30
Connector      : None
Community      : color:00:10 target:64500:11
                 l2-attribute:MTU: 1514 C: 0 P: 0 B: 0
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.5
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
EVPN type      : AUTO-DISC
ESI            : ESI-0
Tag            : 1
Route Dist.    : 192.0.2.5:11
MPLS Label     : 524288
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
Add Paths Send : Default
Last Modified  : 01d23h12m
SRv6 TLV Type  : SRv6 L2 Service TLV (6)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:aaaa:5::
Full Sid       : 2001:db8:aaaa:5:8000::
Behavior       : End.DX2 (21)
SRv6 SubSubTLV : SRv6 SID Structure (1)
Loc-Block-Len  : 48                     Loc-Node-Len   : 16
Func-Len       : 20                     Arg-Len        : 0
Tpose-Len      : 20                     Tpose-offset   : 64
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Modified Attributes

Network        : n/a
Nexthop        : 2001:db8::5
Path Id        : None
From           : 2001:db8::5
Res. Nexthop   : fe80::b69e:ffff:fe00:0
Local Pref.    : 100                    Interface Name : int-1-bng-1-p-1
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : 30
Connector      : None
Community      : color:00:10 target:64500:11
                 l2-attribute:MTU: 1514 C: 0 P: 0 B: 0
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.5
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
EVPN type      : AUTO-DISC
ESI            : ESI-0
Tag            : 1
Route Dist.    : 192.0.2.5:11
MPLS Label     : 524288
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
Add Paths Send : Default
Last Modified  : 01d23h12m
SRv6 TLV Type  : SRv6 L2 Service TLV (6)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:aaaa:5::
Full Sid       : 2001:db8:aaaa:5:8000::
Behavior       : End.DX2 (21)
SRv6 SubSubTLV : SRv6 SID Structure (1)
Loc-Block-Len  : 48                     Loc-Node-Len   : 16
Func-Len       : 20                     Arg-Len        : 0
Tpose-Len      : 20                     Tpose-offset   : 64

-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Routes : 1
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  service  id "evpn-dual-homing" ethernet-segment "ES-1"
No sap entries
No sdp entries
No vxlan instance entries

===============================================================================
SDP Ethernet-Segment Information
===============================================================================
Pw-Port               Eth-Seg                          Status
-------------------------------------------------------------------------------
1                     ES-1                             DF
===============================================================================

[/]
A:admin@bng-1#
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
[/]
A:admin@bng-1# /show  service id "evpn-dual-homing"  segment-routing-v6 detail

===============================================================================
Segment Routing v6 Instance 1 Service 11
===============================================================================
Locator
  Type         Function  SID                                      Status
-------------------------------------------------------------------------------
bng-1-loc
  End.DX2        *524286 2001:db8:aaaa:14:7fff:e000::            ok
===============================================================================
Legend: * - System allocated

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show service id "evpn-dual-homing" segment-routing-v6 instance 1 destinations

===============================================================================
TEP, SID
===============================================================================
Instance  TEP Address                       Segment Id
-------------------------------------------------------------------------------
1         2001:db8::5                       2001:db8:aaaa:5:8000::
-------------------------------------------------------------------------------
Number of TEP, SID: 1
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
Segment Routing v6 Ethernet Segment Dest
===============================================================================
Instance  Eth SegId                      Num. Macs     Last Change
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show service id "evpn-dual-homing" segment-routing-v6 instance 1 end-dx2

===============================================================================
Segment Routing v6 Instance 1 Service 11
===============================================================================
Locator
  Type         Function  SID                                      Status
-------------------------------------------------------------------------------
bng-1-loc
  End.DX2        *524286 2001:db8:aaaa:14:7fff:e000::            ok
===============================================================================
Legend: * - System allocated

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  service id "evpn-dual-homing"  segment-routing-v6 instance 1 locator
 "bng-1-loc"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Segment Routing v6 Instance 1 Service 11
===============================================================================
Locator
  Type         Function  SID                                        Status
-------------------------------------------------------------------------------
bng-1-loc
  End.DX2        *524286 2001:db8:aaaa:14:7fff:e000::            ok
===============================================================================
Legend: * - System allocated

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show service id "dual-homing" base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 10                  Vpn Id            : 0
Service Type      : VPRN
MACSec enabled    : no
Name              : dual-homing
Description       : (Not Specified)
Customer Id       : 1                   Creation Origin   : manual
Last Status Change: 11/12/2022 16:21:35
Last Mgmt Change  : 11/12/2022 16:21:35
Admin State       : Up                  Oper State        : Up

Router Oper State : Up
Route Dist.       : None                VPRN Type         : regular
Oper Route Dist   : 0:0
Oper RD Type      : none
AS Number         : None                Router Id         : 192.0.2.20
ECMP              : Enabled             ECMP Max Routes   : 1
Max IPv4 Routes   : No Limit
Local Rt Domain-Id: None                D-Path Lng Ignore : Disabled

Auto Bind Tunnel
Allow Flex-Alg-Fb : Disabled
Resolution        : disabled
Weighted ECMP     : Disabled            ECMP Max Routes   : 1
Strict Tnl Tag    : Disabled

Max IPv6 Routes   : No Limit
Ignore NH Metric  : Disabled
Hash Label        : Disabled
Entropy Label     : Disabled
Vrf Target        : None
Vrf Import        : None
Vrf Export        : None
MVPN Vrf Target   : None
MVPN Vrf Import   : None
MVPN Vrf Export   : None
Car. Sup C-VPN    : Disabled
Label mode        : vrf
BGP VPN Backup    : Disabled
BGP Export Inactv : Disabled
LOG all events    : Disabled

SAP Count         : 21                  SDP Bind Count    : 1
VSD Domain        : <none>
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                              Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:[pw-1:2.1]                          qinq         8932    8888    Up   Up
sap:[pw-1:2.2]                          qinq         8932    8888    Up   Up
sap:[pw-1:2.3]                          qinq         8932    8888    Up   Up
sap:[pw-1:2.4]                          qinq         8932    8888    Up   Up
sap:[pw-1:2.5]                          qinq         8932    8888    Up   Up
sap:[pw-1:2.6]                          qinq         8932    8888    Up   Up
sap:[pw-1:2.7]                          qinq         8932    8888    Up   Up
sap:[pw-1:2.8]                          qinq         8932    8888    Up   Up
sap:[pw-1:2.9]                          qinq         8932    8888    Up   Up
sap:[pw-1:2.10]                         qinq         8932    8888    Up   Up
sap:[pw-1:2.11]                         qinq         8932    8888    Up   Up
sap:[pw-1:2.12]                         qinq         8932    8888    Up   Up
sap:[pw-1:2.13]                         qinq         8932    8888    Up   Up
sap:[pw-1:2.14]                         qinq         8932    8888    Up   Up
sap:[pw-1:2.15]                         qinq         8932    8888    Up   Up
sap:[pw-1:2.16]                         qinq         8932    8888    Up   Up
sap:[pw-1:2.17]                         qinq         8932    8888    Up   Up
sap:[pw-1:2.18]                         qinq         8932    8888    Up   Up
sap:[pw-1:2.19]                         qinq         8932    8888    Up   Up
sap:[pw-1:2.20]                         qinq         8932    8888    Up   Up
sap:pw-1:2.4094                         qinq         8932    8888    Up   Up
sdp:1:1 S(192.168.7.2)                  None         0       8890    Up   Up
-------------------------------------------------------------------------------
[<sap-id>] indicates a Managed SAP
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show service id "dual-homing" bgp-ipvpn segment-routing-v6

===============================================================================
Service 10 BGP-IPVPN Segment-Routing-V6 Information
===============================================================================

Admin State       : Up
VRF Import        : None
VRF Export        : srrp-aware-routing
Route Dist.       : 192.0.2.20:10
Oper Route Dist   : 192.0.2.20:10
Oper RD Type      : configured
Route Target      : target:64500:10
Route Target Expor: None
Route Target Impor: None
Def Route Tag     : 0x0
Route Resolution  : route-table

Srv6 Instance     : 1
Default Locator   : bng-1-loc
Source Address    : 2001:db8::14
Domain-Id         : None

===============================================================================

[/]
A:admin@bng-1#

[/]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
A:admin@bng-1# /show service id "dual-homing" segment-routing-v6 instance 1 locator "bng-1-loc"

===============================================================================
Segment Routing v6 Instance 1 Service 10
===============================================================================
Locator
  Type         Function  SID                                          Status
-------------------------------------------------------------------------------
bng-1-loc
  End.DT4        *524285 2001:db8:aaaa:14:7fff:d000::                 ok
  End.DT6        *524284 2001:db8:aaaa:14:7fff:c000::                 ok
===============================================================================
Legend: * - System allocated

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  router 10 route-table

===============================================================================
Route Table (Service: 10)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
10.10.0.0/24                                  Local   Local     01d23h11m  0
       sub-int-1                                                    0
10.10.0.2/32                                  Remote  Sub Mgmt  01d23h12m  0
       [red-int-bng-1-bng-2]                                        0
10.10.0.10/32                                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                0
10.10.0.11/32                                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                0
10.10.0.12/32                                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                0
10.10.0.13/32                                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                0
10.10.0.14/32                                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                0
10.10.0.15/32                                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                0
10.10.0.16/32                                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                0
10.10.0.17/32                                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                0
10.10.0.18/32                                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                0
10.10.0.19/32                                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                0
10.10.1.0/24                                  Local   Local     01d23h11m  0
       sub-int-1                                                    0
10.10.1.2/32                                  Remote  Sub Mgmt  01d23h12m  0
       [red-int-bng-1-bng-2]                                        0
10.10.1.10/32                                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                0
10.10.1.11/32                                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                0
10.10.1.12/32                                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                0
10.10.1.13/32                                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                0
10.10.1.14/32                                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
10.10.1.15/32                                   Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                  0
10.10.1.16/32                                   Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                  0
10.10.1.17/32                                   Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                  0
10.10.1.18/32                                   Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                  0
10.10.1.19/32                                   Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                  0
172.16.102.0/24                                 Remote  BGP VPN   01d23h12m  170
       2001:db8:aaaa:6:8000:: (tunneled:SRV6)                         20
192.168.0.1/32                                  Local   Local     01d23h13m  0
       loopback-1                                                     0
192.168.11.0/24                                 Local   Local     01d23h13m  0
       red-int-bng-1-bng-2                                            0
-------------------------------------------------------------------------------
No. of Routes: 27
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  router 10 route-table ipv6

===============================================================================
IPv6 Route Table (Service: 10)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
2001:db8::1/128                                 Local   Local     01d23h13m  0
       loopback-1                                                     0
2001:db8::6600/120                              Remote  BGP VPN   01d23h12m  170
       2001:db8:aaaa:6:7fff:f000:: (tunneled:SRV6)                    20
2001:db8:bbbb::/56                              Local   Local     01d23h11m  0
       sub-int-1                                                      0
2001:db8:bbbb::1/128                            Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                  0
2001:db8:bbbb:1::/64                            Remote  Managed   01d23h11m  0
       2001:db8:bbbb::1                                               0
2001:db8:bbbb:2::1/128                          Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                  0
2001:db8:bbbb:3::/64                            Remote  Managed   01d23h11m  0
       2001:db8:bbbb:2::1                                             0
2001:db8:bbbb:4::1/128                          Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                  0
2001:db8:bbbb:5::/64                            Remote  Managed   01d23h11m  0
       2001:db8:bbbb:4::1                                             0
2001:db8:bbbb:6::1/128                          Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                  0
2001:db8:bbbb:7::/64                            Remote  Managed   01d23h11m  0
       2001:db8:bbbb:6::1                                             0
2001:db8:bbbb:8::1/128                          Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                  0
2001:db8:bbbb:9::/64                            Remote  Managed   01d23h11m  0
       2001:db8:bbbb:8::1                                             0
2001:db8:bbbb:a::1/128                          Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                  0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
2001:db8:bbbb:b::/64                          Remote  Managed   01d23h11m  0
       2001:db8:bbbb:a::1                                                  0
2001:db8:bbbb:c::1/128                        Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                       0
2001:db8:bbbb:d::/64                          Remote  Managed   01d23h11m  0
       2001:db8:bbbb:c::1                                                  0
2001:db8:bbbb:e::1/128                        Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                       0
2001:db8:bbbb:f::/64                          Remote  Managed   01d23h11m  0
       2001:db8:bbbb:e::1                                                  0
2001:db8:bbbb:10::1/128                       Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                       0
2001:db8:bbbb:11::/64                         Remote  Managed   01d23h11m  0
       2001:db8:bbbb:10::1                                                 0
2001:db8:bbbb:12::1/128                       Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                       0
2001:db8:bbbb:13::/64                         Remote  Managed   01d23h11m  0
       2001:db8:bbbb:12::1                                                 0
2001:db8:bbbb:14::1/128                       Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                       0
2001:db8:bbbb:15::/64                         Remote  Managed   01d23h11m  0
       2001:db8:bbbb:14::1                                                 0
2001:db8:bbbb:16::1/128                       Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                       0
2001:db8:bbbb:17::/64                         Remote  Managed   01d23h11m  0
       2001:db8:bbbb:16::1                                                 0
2001:db8:bbbb:18::1/128                       Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                       0
2001:db8:bbbb:19::/64                         Remote  Managed   01d23h11m  0
       2001:db8:bbbb:18::1                                                 0
2001:db8:bbbb:1a::1/128                       Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                       0
2001:db8:bbbb:1b::/64                         Remote  Managed   01d23h11m  0
       2001:db8:bbbb:1a::1                                                 0
2001:db8:bbbb:1c::1/128                       Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                       0
2001:db8:bbbb:1d::/64                         Remote  Managed   01d23h11m  0
       2001:db8:bbbb:1c::1                                                 0
2001:db8:bbbb:1e::1/128                       Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                       0
2001:db8:bbbb:1f::/64                         Remote  Managed   01d23h11m  0
       2001:db8:bbbb:1e::1                                                 0
2001:db8:bbbb:20::1/128                       Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                       0
2001:db8:bbbb:21::/64                         Remote  Managed   01d23h11m  0
       2001:db8:bbbb:20::1                                                 0
2001:db8:bbbb:22::1/128                       Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                       0
2001:db8:bbbb:23::/64                         Remote  Managed   01d23h11m  0
       2001:db8:bbbb:22::1                                                 0
2001:db8:bbbb:24::1/128                       Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                       0
2001:db8:bbbb:25::/64                         Remote  Managed   01d23h11m  0
       2001:db8:bbbb:24::1                                                 0
2001:db8:bbbb:26::1/128                       Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                       0
2001:db8:bbbb:27::/64                         Remote  Managed   01d23h11m  0
       2001:db8:bbbb:26::1                                                 0
2001:db8:bbbb:100::/56                        Local   Local     01d23h11m  0
       sub-int-1                                                           0
-------------------------------------------------------------------------------
No. of Routes: 44
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show router bgp routes vpn-ipv4 hunt
===============================================================================
 BGP Router ID:192.0.2.20       AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv4 Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : 172.16.102.0/24
Nexthop        : 2001:db8::6
Route Dist.    : 192.0.2.6:10           VPN Label      : 524288
Path Id        : None
From           : 2001:db8::6
Res. Nexthop   : n/a
Local Pref.    : 100                    Interface Name : int-1-bng-1-pe-2
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : 20
Connector      : None
Community      : target:64500:10
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.6
Fwd Class      : None                   Priority       : None
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
Add Paths Send : Default
Last Modified  : 01d23h12m
SRv6 TLV Type  : SRv6 L3 Service TLV (5)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:aaaa:6::
Full Sid       : 2001:db8:aaaa:6:8000::
Behavior       : End.DT4 (19)
SRv6 SubSubTLV : SRv6 SID Structure (1)
Loc-Block-Len  : 48                     Loc-Node-Len   : 16
Func-Len       : 20                     Arg-Len        : 0
Tpose-Len      : 20                     Tpose-offset   : 64
VPRN Imported  :  10

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
Network        : 10.10.0.0/24
Nexthop        : 2001:db8::14
Route Dist.    : 192.0.2.20:10          VPN Label      : 524285
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Path Id        : None
To             : 2001:db8::6
Res. Nexthop   : n/a
Local Pref.    : 150                    Interface Name : NotAvailable
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : n/a
Connector      : None
Community      : target:64500:10
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.6
Origin         : IGP
AS-Path        : No As-Path
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
SRv6 TLV Type  : SRv6 L3 Service TLV (5)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:aaaa:14::
Full Sid       : 2001:db8:aaaa:14:7fff:d000::
Behavior       : End.DT4 (19)
SRv6 SubSubTLV : SRv6 SID Structure (1)
Loc-Block-Len  : 48                     Loc-Node-Len   : 16
Func-Len       : 20                     Arg-Len        : 0
Tpose-Len      : 20                     Tpose-offset   : 64

Network        : 10.10.1.0/24
Nexthop        : 2001:db8::14
Route Dist.    : 192.0.2.20:10          VPN Label      : 524285
Path Id        : None
To             : 2001:db8::6
Res. Nexthop   : n/a
Local Pref.    : 150                    Interface Name : NotAvailable
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : n/a
Connector      : None
Community      : target:64500:10
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.6
Origin         : IGP
AS-Path        : No As-Path
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
SRv6 TLV Type  : SRv6 L3 Service TLV (5)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:aaaa:14::
Full Sid       : 2001:db8:aaaa:14:7fff:d000::
Behavior       : End.DT4 (19)
SRv6 SubSubTLV : SRv6 SID Structure (1)
Loc-Block-Len  : 48                     Loc-Node-Len   : 16
Func-Len       : 20                     Arg-Len        : 0
Tpose-Len      : 20                     Tpose-offset   : 64

-------------------------------------------------------------------------------
Routes : 3
===============================================================================

[/]
A:admin@bng-1#
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
[/]
A:admin@bng-1# /show router bgp routes vpn-ipv6 hunt
===============================================================================
 BGP Router ID:192.0.2.20       AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv6 Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : 2001:db8::6600/120
Nexthop        : 2001:db8::6
Route Dist.    : 192.0.2.6:10           VPN Label      : 524287
Path Id        : None
From           : 2001:db8::6
Res. Nexthop   : n/a
Local Pref.    : 100                    Interface Name : int-1-bng-1-pe-2
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : 20
Connector      : None
Community      : target:64500:10
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.6
Fwd Class      : None                   Priority       : None
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
Add Paths Send : Default
Last Modified  : 01d23h12m
SRv6 TLV Type  : SRv6 L3 Service TLV (5)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:aaaa:6::
Full Sid       : 2001:db8:aaaa:6:7fff:f000::
Behavior       : End.DT6 (18)
SRv6 SubSubTLV : SRv6 SID Structure (1)
Loc-Block-Len  : 48                     Loc-Node-Len   : 16
Func-Len       : 20                     Arg-Len        : 0
Tpose-Len      : 20                     Tpose-offset   : 64
VPRN Imported  :  10

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
Network        : 2001:db8:bbbb::/56
Nexthop        : 2001:db8::14
Route Dist.    : 192.0.2.20:10          VPN Label      : 524284
Path Id        : None
To             : 2001:db8::6
Res. Nexthop   : n/a
Local Pref.    : 150                    Interface Name : NotAvailable
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : n/a
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Connector      : None
Community      : target:64500:10
Cluster        : No Cluster Members
Originator Id  : None                  Peer Router Id : 192.0.2.6
Origin         : IGP
AS-Path        : No As-Path
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                     Dest Class     : 0
SRv6 TLV Type  : SRv6 L3 Service TLV (5)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:aaaa:14::
Full Sid       : 2001:db8:aaaa:14:7fff:c000::
Behavior       : End.DT6 (18)
SRv6 SubSubTLV : SRv6 SID Structure (1)
Loc-Block-Len  : 48                    Loc-Node-Len   : 16
Func-Len       : 20                    Arg-Len        : 0
Tpose-Len      : 20                    Tpose-offset   : 64

Network        : 2001:db8:bbbb:100::/56
Nexthop        : 2001:db8::14
Route Dist.    : 192.0.2.20:10         VPN Label      : 524284
Path Id        : None
To             : 2001:db8::6
Res. Nexthop   : n/a
Local Pref.    : 150                   Interface Name : NotAvailable
Aggregator AS  : None                  Aggregator     : None
Atomic Aggr.   : Not Atomic            MED            : None
AIGP Metric    : None                  IGP Cost       : n/a
Connector      : None
Community      : target:64500:10
Cluster        : No Cluster Members
Originator Id  : None                  Peer Router Id : 192.0.2.6
Origin         : IGP
AS-Path        : No As-Path
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                     Dest Class     : 0
SRv6 TLV Type  : SRv6 L3 Service TLV (5)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:aaaa:14::
Full Sid       : 2001:db8:aaaa:14:7fff:c000::
Behavior       : End.DT6 (18)
SRv6 SubSubTLV : SRv6 SID Structure (1)
Loc-Block-Len  : 48                    Loc-Node-Len   : 16
Func-Len       : 20                    Arg-Len        : 0
Tpose-Len      : 20                    Tpose-offset   : 64

-------------------------------------------------------------------------------
Routes : 3
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show router segment-routing-v6 base-routing-instance all

===============================================================================
Segment Routing v6 Base Routing Instance
===============================================================================
Locator
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
  Type       Function    SID                                        Status/InstId
    SRH-mode Protection  Interface
-------------------------------------------------------------------------------
bng-1-loc
  End                  1 2001:db8:aaaa:14:0:1000::                  ok
    USP
-------------------------------------------------------------------------------
  Auto-allocated End.X: USP Protected,
-------------------------------------------------------------------------------
  End.X          *524287 2001:db8:aaaa:14:7fff:f000::               1
    USP        Protected int-1-bng-1-bng-2
    ISIS Level: L2 Mac Address: b4:98:01:01:00:06 Nbr Sys Id: 1920.0000.2021
  End.X          *524288 2001:db8:aaaa:14:8000::                    0
    USP        Protected int-1-bng-1-p-1
    ISIS Level: L2 Mac Address: b4:9e:01:01:00:0b Nbr Sys Id: 1920.0000.2004
  End.X          *524289 2001:db8:aaaa:14:8000:1000::               1
    USP        Protected int-1-bng-1-pe-2
    ISIS Level: L2 Mac Address: b4:a0:01:01:00:07 Nbr Sys Id: 1920.0000.2006
-------------------------------------------------------------------------------
===============================================================================
Legend: * - System allocated

===============================================================================
Micro Segment Routing v6 Base Routing Instance
===============================================================================
Micro Segment Locator
  Type       Function    SID                                        Status/InstId
    SRH-mode Oper Func   Interface                                  Protection
-------------------------------------------------------------------------------
===============================================================================
Legend: * - System allocated

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show router segment-routing-v6 local-sid context  "Base"

===============================================================================
Segment Routing v6 Local SIDs
===============================================================================
SID                                                 Type           Function
  Locator
  Context
-------------------------------------------------------------------------------
2001:db8:aaaa:14:0:1000::                           End            1
  bng-1-loc
  Base
-------------------------------------------------------------------------------
SIDs : 1
-------------------------------------------------------------------------------
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show router segment-routing-v6 local-sid context  "10"

===============================================================================
Segment Routing v6 Local SIDs
===============================================================================
SID                                                 Type           Function
  Locator
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
  Context
-------------------------------------------------------------------------------
2001:db8:aaaa:14:7fff:c000::                       End.DT6        524284
  bng-1-loc
  SvcId: 10 Name: dual-homing
2001:db8:aaaa:14:7fff:d000::                       End.DT4        524285
  bng-1-loc
  SvcId: 10 Name: dual-homing
-------------------------------------------------------------------------------
SIDs : 2
-------------------------------------------------------------------------------
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show router route-table ipv6

===============================================================================
IPv6 Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                             Type    Proto     Age        Pref
      Next Hop[Interface Name]                                     Metric
-------------------------------------------------------------------------------
2001:db8::4/128                                Remote  ISIS      01d23h12m  18
       fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                        10
2001:db8::5/128                                Remote  ISIS      01d23h12m  18
       fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                        20
2001:db8::6/128 [L]                            Remote  ISIS(1)   01d23h12m  18
       fe80::b6a0:ffff:fe00:0-"int-1-bng-1-pe-2"                       10
2001:db8::a/128                                Remote  ISIS      01d23h12m  18
       fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                        20
2001:db8::14/128                               Local   Local     01d23h13m  0
       system                                                          0
2001:db8::15/128 [L]                           Remote  ISIS(1)   01d23h12m  18
       fe80::b697:ffff:fe00:0-"int-1-bng-1-bng-2"                      10
2001:db8::100/120                              Remote  ISIS      01d23h12m  18
       fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                        20
2001:db8::200/120                              Remote  ISIS      01d23h12m  18
       fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                        30
2001:db8::500/120                              Local   Local     01d23h13m  0
       int-1-bng-1-p-1                                                 0
2001:db8::600/120                              Remote  ISIS      01d23h12m  18
       fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                        30
2001:db8::700/120                              Local   Local     01d23h13m  0
       int-1-bng-1-bng-2                                               0
2001:db8::800/120                              Local   Local     01d23h13m  0
       int-1-bng-1-pe-2                                                0
2001:db8::900/120 [L]                          Remote  ISIS(1)   01d23h12m  18
       fe80::b6a0:ffff:fe00:0-"int-1-bng-1-pe-2"                       20
2001:db8::a00/120                              Remote  ISIS      01d23h12m  18
       fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                        20
2001:db8::6400/120                             Remote  ISIS      01d23h12m  18
       fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                        30
2001:db8::6500/120 [L]                         Remote  ISIS(1)   01d23h12m  18
       fe80::b6a0:ffff:fe00:0-"int-1-bng-1-pe-2"                       20
2001:db8:aaaa:4::/64                           Remote  ISIS      01d23h12m  18
       2001:db8:aaaa:4::/64 (tunneled:SRV6-ISIS)                       20
2001:db8:aaaa:5::/64                           Remote  ISIS      01d23h12m  18
       2001:db8:aaaa:5::/64 (tunneled:SRV6-ISIS)                       30
2001:db8:aaaa:6::/64                           Remote  ISIS(1)   01d23h12m  18
       2001:db8:aaaa:6::/64 (tunneled:SRV6-ISIS)                       20
2001:db8:aaaa:a::/64                           Remote  ISIS      01d23h12m  18
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
       2001:db8:aaaa:a::/64 (tunneled:SRV6-ISIS)                    30
2001:db8:aaaa:14::/64                       Local   SRV6      01d23h13m  3
       fe80::201-"_tmnx_fpe_2.a"                                    0
2001:db8:aaaa:14:0:1000::/128               Local   SRV6      01d23h13m  3
       Black Hole                                                   0
2001:db8:aaaa:14:0:2000::/128               Local   SRV6-Pol* 01d23h13m  14
       2001:db8::5 (tunneled:SRV6-Policy:917506)                    1
2001:db8:aaaa:14:7fff:f000::/128            Local   ISIS(1)   01d23h12m  18
       2001:db8:aaaa:14:7fff:f000:: (tunneled:SRV6-ISIS)            10
2001:db8:aaaa:14:8000::/128                 Local   ISIS      01d23h12m  18
       2001:db8:aaaa:14:8000:: (tunneled:SRV6-ISIS)                 10
2001:db8:aaaa:14:8000:1000::/128            Local   ISIS(1)   01d23h12m  18
       2001:db8:aaaa:14:8000:1000:: (tunneled:SRV6-ISIS)            10
2001:db8:aaaa:15::/64                       Remote  ISIS(1)   01d23h12m  18
       2001:db8:aaaa:15::/64 (tunneled:SRV6-ISIS)                   20
-------------------------------------------------------------------------------
No. of Routes: 27
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
* indicates that the corresponding row element may have been truncated.

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show router tunnel-table ipv6

===============================================================================
IPv6 Tunnel Table (Router: Base)
===============================================================================
Destination                                 Owner     Encap TunnelId  Pref
Nexthop                                     Color           Metric
-------------------------------------------------------------------------------
2001:db8::5/128                             srv6-pol  SRV6  917506    14
  fpe_1.a                                     10              0
2001:db8:aaaa:4::/64                        srv6-isis SRV6  524296    0
  fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                    20
2001:db8:aaaa:5::/64                        srv6-isis SRV6  524299    0
  fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                    30
2001:db8:aaaa:6::/64 [L]                    srv6-isis SRV6  524295    0
  fe80::b6a0:ffff:fe00:0-"int-1-bng-1-pe-2"                   20
2001:db8:aaaa:a::/64                        srv6-isis SRV6  524297    0
  fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                    30
2001:db8:aaaa:14:7fff:f000::/128 [L]        srv6-isis SRV6  524291    0
  fe80::b697:ffff:fe00:0-"int-1-bng-1-bng-2"                  10
2001:db8:aaaa:14:8000::/128                 srv6-isis SRV6  524292    0
  fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                    10
2001:db8:aaaa:14:8000:1000::/128 [L]        srv6-isis SRV6  524293    0
  fe80::b6a0:ffff:fe00:0-"int-1-bng-1-pe-2"                   10
2001:db8:aaaa:15::/64 [L]                   srv6-isis SRV6  524294    0
  fe80::b697:ffff:fe00:0-"int-1-bng-1-bng-2"                  20
-------------------------------------------------------------------------------
Flags: B = BGP or MPLS backup hop available
       L = Loop-Free Alternate (LFA) hop available
       E = Inactive best-external BGP route
       k = RIB-API or Forwarding Policy backup hop
===============================================================================

[/]
A:admin@bng-1#
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
[/]
A:admin@bng-1# /show router fp-tunnel-table 1 ipv6

===============================================================================
IPv6 Tunnel Table Display

Legend:
label stack is ordered from bottom-most to top-most
B - FRR Backup
===============================================================================
Destination                                Protocol         Tunnel-ID
  Lbl/SID
    NextHop                                                 Intf/Tunnel
  Lbl/SID (backup)
    NextHop   (backup)
-------------------------------------------------------------------------------
2001:db8:aaaa:4::/64                       SRV6              524296
  -
    fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                1/1/c1/1:1
2001:db8:aaaa:5::/64                       SRV6              524299
  -
    fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                1/1/c1/1:1
2001:db8:aaaa:6::/64                       SRV6              524295
  -
    fe80::b6a0:ffff:fe00:0-"int-1-bng-1-pe-2"               1/1/c1/7:1
  -
    fe80::b697:ffff:fe00:0-"int-1-bng-1-bng-2"(B)           1/1/c1/6:1
2001:db8:aaaa:a::/64                       SRV6              524297
  -
    fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                1/1/c1/1:1
2001:db8:aaaa:15::/64                      SRV6              524294
  -
    fe80::b697:ffff:fe00:0-"int-1-bng-1-bng-2"              1/1/c1/6:1
  -
    fe80::b6a0:ffff:fe00:0-"int-1-bng-1-pe-2"(B)            1/1/c1/7:1
2001:db8::5/128                            SRV6-Policy       -
  2001:db8:aaaa:a:0:1000::/2001:db8:aaaa:4:0:1000::
    0.140.1.1                                               pxc-1.b:1
2001:db8:aaaa:14:7fff:f000::/128           SRV6              524291
  -
    fe80::b697:ffff:fe00:0-"int-1-bng-1-bng-2"              1/1/c1/6:1
  2001:db8:aaaa:15:0:1000::
    fe80::b6a0:ffff:fe00:0-"int-1-bng-1-pe-2"(B)            1/1/c1/7:1
2001:db8:aaaa:14:8000::/128                SRV6              524292
  -
    fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                1/1/c1/1:1
2001:db8:aaaa:14:8000:1000::/128           SRV6              524293
  -
    fe80::b6a0:ffff:fe00:0-"int-1-bng-1-pe-2"               1/1/c1/7:1
  2001:db8:aaaa:6:0:1000::
    fe80::b697:ffff:fe00:0-"int-1-bng-1-bng-2"(B)           1/1/c1/6:1
-------------------------------------------------------------------------------
Total Entries : 9
-------------------------------------------------------------------------------
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show router  bgp  summary
===============================================================================
 BGP Router ID:192.0.2.20       AS:64500       Local AS:64500
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
BGP Admin State         : Up          BGP Oper State              : Up
Total Peer Groups       : 2           Total Peers                 : 3
Total VPN Peer Groups   : 0           Total VPN Peers             : 0
Current Internal Groups : 3           Max Internal Groups         : 3
Total BGP Paths         : 59          Total Path Memory           : 21280

Total IPv4 Remote Rts   : 0           Total IPv4 Rem. Active Rts  : 0
Total IPv6 Remote Rts   : 0           Total IPv6 Rem. Active Rts  : 0
Total IPv4 Backup Rts   : 0           Total IPv6 Backup Rts       : 0
Total LblIpv4 Rem Rts   : 0           Total LblIpv4 Rem. Act Rts  : 0
Total LblIpv6 Rem Rts   : 0           Total LblIpv6 Rem. Act Rts  : 0
Total LblIpv4 Bkp Rts   : 0           Total LblIpv6 Bkp Rts       : 0
Total Supressed Rts     : 0           Total Hist. Rts             : 0
Total Decay Rts         : 0

Total VPN-IPv4 Rem. Rts : 1           Total VPN-IPv4 Rem. Act. Rts: 1
Total VPN-IPv6 Rem. Rts : 1           Total VPN-IPv6 Rem. Act. Rts: 1
Total VPN-IPv4 Bkup Rts : 0           Total VPN-IPv6 Bkup Rts     : 0
Total VPN Local Rts     : 7           Total VPN Supp. Rts         : 0
Total VPN Hist. Rts     : 0           Total VPN Decay Rts         : 0

Total MVPN-IPv4 Rem Rts : 0           Total MVPN-IPv4 Rem Act Rts : 0
Total MVPN-IPv6 Rem Rts : 0           Total MVPN-IPv6 Rem Act Rts : 0
Total MDT-SAFI Rem Rts  : 0           Total MDT-SAFI Rem Act Rts  : 0
Total McIPv4 Remote Rts : 0           Total McIPv4 Rem. Active Rts: 0
Total McIPv6 Remote Rts : 0           Total McIPv6 Rem. Active Rts: 0
Total McVpnIPv4 Rem Rts : 0           Total McVpnIPv4 Rem Act Rts : 0
Total McVpnIPv6 Rem Rts : 0           Total McVpnIPv6 Rem Act Rts : 0

Total EVPN Rem Rts      : 4           Total EVPN Rem Act Rts      : 4
Total L2-VPN Rem. Rts   : 0           Total L2VPN Rem. Act. Rts   : 0
Total MSPW Rem Rts      : 0           Total MSPW Rem Act Rts      : 0
Total RouteTgt Rem Rts  : 0           Total RouteTgt Rem Act Rts  : 0
Total FlowIpv4 Rem Rts  : 0           Total FlowIpv4 Rem Act Rts  : 0
Total FlowIpv6 Rem Rts  : 0           Total FlowIpv6 Rem Act Rts  : 0
Total FlowVpnv4 Rem Rts : 0           Total FlowVpnv4 Rem Act Rts : 0
Total FlowVpnv6 Rem Rts : 0           Total FlowVpnv6 Rem Act Rts : 0
Total Link State Rem Rts: 0           Total Link State Rem Act Rts: 0
Total SrPlcyIpv4 Rem Rts: 0           Total SrPlcyIpv4 Rem Act Rts: 0
Total SrPlcyIpv6 Rem Rts: 0           Total SrPlcyIpv6 Rem Act Rts: 0

===============================================================================
BGP Summary
===============================================================================
Legend : D - Dynamic Neighbor
===============================================================================
Neighbor
Description
                   AS PktRcvd InQ  Up/Down   State|Rcv/Act/Sent (Addr Family)
                      PktSent OutQ
-------------------------------------------------------------------------------
2001:db8::5
                64500    5669    0 01d23h12m 1/1/3 (Evpn)
                         5672    0
2001:db8::6
                64500    5670    0 01d23h12m 1/1/2 (VpnIPv4)
                         5672    0           1/1/2 (VpnIPv6)
2001:db8::15
                64500    5673    0 01d23h12m 3/3/3 (Evpn)
                         5672    0
-------------------------------------------------------------------------------

[/]
A:admin@bng-1#
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
[/]
A:admin@bng-1# /show router isis  adjacency

===============================================================================
Rtr Base ISIS Instance 0 Adjacency
===============================================================================
System ID               Usage State Hold Interface                    MT-ID
-------------------------------------------------------------------------------
p-1                     L2    Up    7    int-1-bng-1-p-1              0
-------------------------------------------------------------------------------
Adjacencies : 1
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show router isis database

===============================================================================
Rtr Base ISIS Instance 0 Database
===============================================================================
LSP ID                                Sequence  Checksum Lifetime Attributes
-------------------------------------------------------------------------------

Displaying Level 1 database
-------------------------------------------------------------------------------
Level (1) LSP Count : 0

Displaying Level 2 database
-------------------------------------------------------------------------------
p-1.00-00                             0x13d     0x8ffb   762      L1L2
p-1.01-00                             0x134     0x4d03   902      L1L2
pe-1.00-00                            0x139     0x659e   863      L1L2
pe-1.01-00                            0x133     0xe585   1063     L1L2
p-2.00-00                             0x137     0xfd42   1085     L1L2
p-2.01-00                             0x135     0x6dc8   938      L1L2
p-2.02-00                             0x135     0xe56c   1042     L1L2
p-2.03-00                             0x132     0xfa58   764      L1L2
bng-1.00-00                           0x13d     0x160f   748      L1L2
bng-2.00-00                           0x137     0x777b   675      L1L2
Level (2) LSP Count : 10
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show router isis 1 adjacency

===============================================================================
Rtr Base ISIS Instance 1 Adjacency
===============================================================================
System ID               Usage State Hold Interface                    MT-ID
-------------------------------------------------------------------------------
bng-2                   L2    Up    27   int-1-bng-1-bng-2            0
pe-2                    L2    Up    7    int-1-bng-1-pe-2             0
-------------------------------------------------------------------------------
Adjacencies : 2
===============================================================================

[/]
A:admin@bng-1#
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
[/]
A:admin@bng-1# /show router isis 1 database

===============================================================================
Rtr Base ISIS Instance 1 Database
===============================================================================
LSP ID                                Sequence  Checksum Lifetime Attributes
-------------------------------------------------------------------------------

Displaying Level 1 database
-------------------------------------------------------------------------------
Level (1) LSP Count : 0

Displaying Level 2 database
-------------------------------------------------------------------------------
pe-2.00-00                            0x13a     0x105f   963      L1L2
pe-2.01-00                            0x133     0x51fb   708      L1L2
pe-2.02-00                            0x133     0x60ea   1055     L1L2
bng-1.00-00                           0x13c     0x29a7   747      L1L2
bng-1.01-00                           0x136     0x7b99   670      L1L2
bng-2.00-00                           0x138     0x8827   1129     L1L2
Level (2) LSP Count : 6
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show redundancy multi-chassis sync peer 192.0.2.21 detail

===============================================================================
Multi-chassis Peer Table
===============================================================================
Peer
-------------------------------------------------------------------------------
Peer IP Address         : 192.0.2.21
Description             : (Not Specified)
Authentication          : Disabled
Source IP Address       : 192.0.2.20
Admin State             : Enabled
Warm standby            : No
Remote warm standby     : No
Sub-mgmt options        :
  DHCP lease threshold  : Inactive
    Local / Remote      : -- / --
-------------------------------------------------------------------------------
Sync-status
-------------------------------------------------------------------------------
Client Applications     : SUBMGMT-IPOE SUBMGMT-PPPOE SRRP DHCPServer
Sync Admin State        : Up
Sync Oper State         : Up
Sync Oper Flags         :
DB Sync State           : inSync
Num Entries             : 100
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
Rem Num Entries         : 100
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
MCS Application Stats
===============================================================================
Application             : igmp
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : igmpSnooping
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : subMgmtIpoe
Num Entries             : 10
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 10
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : srrp
Num Entries             : 26
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 26
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : mcRing
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : mldSnooping
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : dhcpServer
Num Entries             : 54
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 54
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : subHostTrk
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : subMgmtPppoe
Num Entries             : 10
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 10
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : ipsec
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : mld
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : python
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : l2tp
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : diameterProxy
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : pimSnpgSap
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : pimSnpgSdp
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : diameterNode
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : nat
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
Ports synced on peer 192.0.2.21
===============================================================================
Port/Encap                    Tag
-------------------------------------------------------------------------------
pw-1                          pw-port-1
===============================================================================

===============================================================================
SDPs synced on peer 192.0.2.21
===============================================================================
SDP/Vc-Id                     Tag
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
DHCP Server instances synced on peer 192.0.2.21
===============================================================================
Router-Name                     Server-Name
  Tag
-------------------------------------------------------------------------------
No instances found
===============================================================================

===============================================================================
Python cache instances synced on peer 192.0.2.21
===============================================================================
Python-Policy                   Tag
-------------------------------------------------------------------------------
No instances found
===============================================================================
No L2TP instances found.

===============================================================================
Track SRRP instances
===============================================================================
SRRP                    : 1
-------------------------------------------------------------------------------
L2TP tunnel ID start    : 0
L2TP tunnel ID end      : 0

===============================================================================

===============================================================================
Diameter proxy instances synced on peer 192.0.2.21
===============================================================================
Diameter-Peer-Policy            Tag
-------------------------------------------------------------------------------
No instances found
===============================================================================

===============================================================================
Diameter node instances synced on peer 192.0.2.21
===============================================================================
Diameter Node                                 Tag
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
No. of Diameter Nodes: 0
===============================================================================

===============================================================================
Nat groups synced on peer 192.0.2.21
===============================================================================
Nat group                                     Tag
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
No. of Nat groups: 0
===============================================================================
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /tools dump redundancy multi-chassis sync-database

The following totals are for:
 peer ip ALL, port/lag/sdp ALL, sync-tag ALL, application ALL
Valid Entries:                  100
Locally Deleted Entries:        0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Locally Deleted Alarmed Entries: 0
Pending Global Delete Entries:   0
Omcr Alarmed Entries:            0
Omcr Standby Entries:            0
Associated Shared Records (ALL): 0
Associated Shared Records (LD):  0

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  service  active-subscribers summary

===============================================================================
Active Subscriber table summary
===============================================================================
Total Count       : 20
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  service  active-subscribers hierarchy

===============================================================================
Active Subscribers Hierarchy
===============================================================================
-- ipoe-ds-1
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.11] - sla:sla-profile-1
       |
       +-- IPOE-session - mac:00:13:01:00:00:01 - svc:10
           |
           |-- 10.10.0.12 - DHCP
           |
           +-- 2001:db8:bbbb::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:1::/64 - DHCP6-PD-MR

-- ipoe-ds-10
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.20] - sla:sla-profile-1
       |
       +-- IPOE-session - mac:00:13:01:00:00:0a - svc:10
           |
           |-- 10.10.0.13 - DHCP
           |
           +-- 2001:db8:bbbb:10::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:11::/64 - DHCP6-PD-MR

-- ipoe-ds-2
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.12] - sla:sla-profile-1
       |
       +-- IPOE-session - mac:00:13:01:00:00:02 - svc:10
           |
           |-- 10.10.0.10 - DHCP
           |
           +-- 2001:db8:bbbb:2::1/128 - DHCP6
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
               |
               +-- 2001:db8:bbbb:3::/64 - DHCP6-PD-MR

-- ipoe-ds-3
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.13] - sla:sla-profile-1
       |
       +-- IPOE-session - mac:00:13:01:00:00:03 - svc:10
           |
           |-- 10.10.0.14 - DHCP
           |
           +-- 2001:db8:bbbb:c::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:d::/64 - DHCP6-PD-MR

-- ipoe-ds-4
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.14] - sla:sla-profile-1
       |
       +-- IPOE-session - mac:00:13:01:00:00:04 - svc:10
           |
           |-- 10.10.0.17 - DHCP
           |
           +-- 2001:db8:bbbb:8::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:9::/64 - DHCP6-PD-MR

-- ipoe-ds-5
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.15] - sla:sla-profile-1
       |
       +-- IPOE-session - mac:00:13:01:00:00:05 - svc:10
           |
           |-- 10.10.0.11 - DHCP
           |
           +-- 2001:db8:bbbb:4::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:5::/64 - DHCP6-PD-MR

-- ipoe-ds-6
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.16] - sla:sla-profile-1
       |
       +-- IPOE-session - mac:00:13:01:00:00:06 - svc:10
           |
           |-- 10.10.0.19 - DHCP
           |
           +-- 2001:db8:bbbb:e::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:f::/64 - DHCP6-PD-MR

-- ipoe-ds-7
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.17] - sla:sla-profile-1
       |
       +-- IPOE-session - mac:00:13:01:00:00:07 - svc:10
           |
           |-- 10.10.0.18 - DHCP
           |
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
           +-- 2001:db8:bbbb:a::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:b::/64 - DHCP6-PD-MR

-- ipoe-ds-8
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.18] - sla:sla-profile-1
       |
       +-- IPOE-session - mac:00:13:01:00:00:08 - svc:10
           |
           |-- 10.10.0.16 - DHCP
           |
           +-- 2001:db8:bbbb:6::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:7::/64 - DHCP6-PD-MR

-- ipoe-ds-9
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.19] - sla:sla-profile-1
       |
       +-- IPOE-session - mac:00:13:01:00:00:09 - svc:10
           |
           |-- 10.10.0.15 - DHCP
           |
           +-- 2001:db8:bbbb:12::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:13::/64 - DHCP6-PD-MR

-- pppoe-1
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.1] - sla:sla-profile-1
       |
       +-- PPP-session - mac:00:11:01:00:00:01 - sid:1 - svc:10
           |
           |-- 10.10.1.10 - IPCP
           |
           +-- 2001:db8:bbbb:14::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:15::/64 - DHCP6-PD-MR

-- pppoe-10
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.10] - sla:sla-profile-1
       |
       +-- PPP-session - mac:00:11:01:00:00:0a - sid:1 - svc:10
           |
           |-- 10.10.1.19 - IPCP
           |
           +-- 2001:db8:bbbb:26::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:27::/64 - DHCP6-PD-MR

-- pppoe-2
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.2] - sla:sla-profile-1
       |
       +-- PPP-session - mac:00:11:01:00:00:02 - sid:1 - svc:10
           |
           |-- 10.10.1.11 - IPCP
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
           |
           +-- 2001:db8:bbbb:18::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:19::/64 - DHCP6-PD-MR

-- pppoe-3
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.3] - sla:sla-profile-1
       |
       +-- PPP-session - mac:00:11:01:00:00:03 - sid:1 - svc:10
           |
           |-- 10.10.1.12 - IPCP
           |
           +-- 2001:db8:bbbb:16::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:17::/64 - DHCP6-PD-MR

-- pppoe-4
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.4] - sla:sla-profile-1
       |
       +-- PPP-session - mac:00:11:01:00:00:04 - sid:1 - svc:10
           |
           |-- 10.10.1.13 - IPCP
           |
           +-- 2001:db8:bbbb:1a::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:1b::/64 - DHCP6-PD-MR

-- pppoe-5
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.5] - sla:sla-profile-1
       |
       +-- PPP-session - mac:00:11:01:00:00:05 - sid:1 - svc:10
           |
           |-- 10.10.1.14 - IPCP
           |
           +-- 2001:db8:bbbb:1c::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:1d::/64 - DHCP6-PD-MR

-- pppoe-6
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.6] - sla:sla-profile-1
       |
       +-- PPP-session - mac:00:11:01:00:00:06 - sid:1 - svc:10
           |
           |-- 10.10.1.15 - IPCP
           |
           +-- 2001:db8:bbbb:1e::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:1f::/64 - DHCP6-PD-MR

-- pppoe-7
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.7] - sla:sla-profile-1
       |
       +-- PPP-session - mac:00:11:01:00:00:07 - sid:1 - svc:10
           |
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
           |-- 10.10.1.16 - IPCP
           |
           +-- 2001:db8:bbbb:20::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:21::/64 - DHCP6-PD-MR

-- pppoe-8
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.8] - sla:sla-profile-1
       |
       +-- PPP-session - mac:00:11:01:00:00:08 - sid:1 - svc:10
           |
           |-- 10.10.1.17 - IPCP
           |
           +-- 2001:db8:bbbb:22::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:23::/64 - DHCP6-PD-MR

-- pppoe-9
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.9] - sla:sla-profile-1
       |
       +-- PPP-session - mac:00:11:01:00:00:09 - sid:1 - svc:10
           |
           |-- 10.10.1.18 - IPCP
           |
           +-- 2001:db8:bbbb:24::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:25::/64 - DHCP6-PD-MR

-------------------------------------------------------------------------------
Number of active subscribers : 20
Flags: (N) = the host or the managed route is in non-forwarding state
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  service id 10 subscriber-hosts detail

===============================================================================
Subscriber Host table
===============================================================================
Sap
  IP Address
    MAC Address               PPPoE-SID       Origin          Fwding State
      Subscriber
-------------------------------------------------------------------------------
[pw-1:2.1]
  10.10.1.10
    00:11:01:00:00:01         1               IPCP            Fwding
      pppoe-1
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000029636FE951
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Acct-Q-Inst-Session-Id: B496FF0000002A636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.1]
  2001:db8:bbbb:14::1/128
    00:11:01:00:00:01          1               PPP-DHCP6       Fwding
      pppoe-1
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000047636FE951
Acct-Q-Inst-Session-Id: B496FF0000002A636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.2]
  10.10.1.11
    00:11:01:00:00:02          1               IPCP            Fwding
      pppoe-2
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF0000002C636FE951
Acct-Q-Inst-Session-Id: B496FF0000002D636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.2]
  2001:db8:bbbb:18::1/128
    00:11:01:00:00:02          1               PPP-DHCP6       Fwding
      pppoe-2
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000049636FE951
Acct-Q-Inst-Session-Id: B496FF0000002D636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.3]
  10.10.1.12
    00:11:01:00:00:03          1               IPCP            Fwding
      pppoe-3
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF0000002F636FE951
Acct-Q-Inst-Session-Id: B496FF00000030636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.3]
  2001:db8:bbbb:16::1/128
    00:11:01:00:00:03          1               PPP-DHCP6       Fwding
      pppoe-3
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000046636FE951
Acct-Q-Inst-Session-Id: B496FF00000030636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.4]
  10.10.1.13
    00:11:01:00:00:04          1               IPCP            Fwding
      pppoe-4
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000032636FE951
Acct-Q-Inst-Session-Id: B496FF00000033636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.4]
  2001:db8:bbbb:1a::1/128
    00:11:01:00:00:04          1               PPP-DHCP6       Fwding
      pppoe-4
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000048636FE951
Acct-Q-Inst-Session-Id: B496FF00000033636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.5]
  10.10.1.14
    00:11:01:00:00:05          1               IPCP            Fwding
      pppoe-5
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000035636FE951
Acct-Q-Inst-Session-Id: B496FF00000036636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.5]
  2001:db8:bbbb:1c::1/128
    00:11:01:00:00:05          1               PPP-DHCP6       Fwding
      pppoe-5
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF0000004A636FE951
Acct-Q-Inst-Session-Id: B496FF00000036636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.6]
  10.10.1.15
    00:11:01:00:00:06          1               IPCP            Fwding
      pppoe-6
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000038636FE951
Acct-Q-Inst-Session-Id: B496FF00000039636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.6]
  2001:db8:bbbb:1e::1/128
    00:11:01:00:00:06          1               PPP-DHCP6       Fwding
      pppoe-6
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF0000004B636FE951
Acct-Q-Inst-Session-Id: B496FF00000039636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.7]
  10.10.1.16
    00:11:01:00:00:07          1               IPCP            Fwding
      pppoe-7
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF0000003B636FE951
Acct-Q-Inst-Session-Id: B496FF0000003C636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.7]
  2001:db8:bbbb:20::1/128
    00:11:01:00:00:07          1               PPP-DHCP6       Fwding
      pppoe-7
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF0000004C636FE951
Acct-Q-Inst-Session-Id: B496FF0000003C636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.8]
  10.10.1.17
    00:11:01:00:00:08          1               IPCP            Fwding
      pppoe-8
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF0000003E636FE951
Acct-Q-Inst-Session-Id: B496FF0000003F636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.8]
  2001:db8:bbbb:22::1/128
    00:11:01:00:00:08          1               PPP-DHCP6       Fwding
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
      pppoe-8
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF0000004D636FE951
Acct-Q-Inst-Session-Id: B496FF0000003F636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.9]
  10.10.1.18
    00:11:01:00:00:09          1               IPCP            Fwding
      pppoe-9
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000041636FE951
Acct-Q-Inst-Session-Id: B496FF00000042636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.9]
  2001:db8:bbbb:24::1/128
    00:11:01:00:00:09          1               PPP-DHCP6       Fwding
      pppoe-9
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF0000004E636FE951
Acct-Q-Inst-Session-Id: B496FF00000042636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.10]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
  10.10.1.19
    00:11:01:00:00:0a          1               IPCP            Fwding
      pppoe-10
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000044636FE951
Acct-Q-Inst-Session-Id: B496FF00000045636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.10]
  2001:db8:bbbb:26::1/128
    00:11:01:00:00:0a          1               PPP-DHCP6       Fwding
      pppoe-10
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF0000004F636FE951
Acct-Q-Inst-Session-Id: B496FF00000045636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.11]
  10.10.0.12
    00:13:01:00:00:01          N/A             DHCP            Fwding
      ipoe-ds-1
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000013636FE94C
Acct-Q-Inst-Session-Id: B496FF00000010636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
[pw-1:2.11]
  2001:db8:bbbb::1/128
    00:13:01:00:00:01          N/A             IPoE-DHCP6      Fwding
      ipoe-ds-1
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF0000000F636FE94C
Acct-Q-Inst-Session-Id: B496FF00000010636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.12]
  10.10.0.10
    00:13:01:00:00:02          N/A             DHCP            Fwding
      ipoe-ds-2
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF0000000A636FE94C
Acct-Q-Inst-Session-Id: B496FF00000005636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.12]
  2001:db8:bbbb:2::1/128
    00:13:01:00:00:02          N/A             IPoE-DHCP6      Fwding
      ipoe-ds-2
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000004636FE94C
Acct-Q-Inst-Session-Id: B496FF00000005636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.13]
  10.10.0.14
    00:13:01:00:00:03          N/A             DHCP            Fwding
      ipoe-ds-3
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF0000001B636FE94C
Acct-Q-Inst-Session-Id: B496FF0000001C636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.13]
  2001:db8:bbbb:c::1/128
    00:13:01:00:00:03          N/A             IPoE-DHCP6      Fwding
      ipoe-ds-3
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF0000001F636FE94C
Acct-Q-Inst-Session-Id: B496FF0000001C636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.14]
  10.10.0.17
    00:13:01:00:00:04          N/A             DHCP            Fwding
      ipoe-ds-4
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000025636FE951
Acct-Q-Inst-Session-Id: B496FF00000012636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.14]
  2001:db8:bbbb:8::1/128
    00:13:01:00:00:04          N/A             IPoE-DHCP6      Fwding
      ipoe-ds-4
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000011636FE94C
Acct-Q-Inst-Session-Id: B496FF00000012636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.15]
  10.10.0.11
    00:13:01:00:00:05          N/A             DHCP            Fwding
      ipoe-ds-5
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF0000000B636FE94C
Acct-Q-Inst-Session-Id: B496FF00000007636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.15]
  2001:db8:bbbb:4::1/128
    00:13:01:00:00:05          N/A             IPoE-DHCP6      Fwding
      ipoe-ds-5
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000006636FE94C
Acct-Q-Inst-Session-Id: B496FF00000007636FE94C
Address Origin        : Dynamic
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.16]
  10.10.0.19
    00:13:01:00:00:06          N/A             DHCP            Fwding
      ipoe-ds-6
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000027636FE951
Acct-Q-Inst-Session-Id: B496FF0000001E636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.16]
  2001:db8:bbbb:e::1/128
    00:13:01:00:00:06          N/A             IPoE-DHCP6      Fwding
      ipoe-ds-6
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF0000001D636FE94C
Acct-Q-Inst-Session-Id: B496FF0000001E636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.17]
  10.10.0.18
    00:13:01:00:00:07          N/A             DHCP            Fwding
      ipoe-ds-7
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000026636FE951
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Acct-Q-Inst-Session-Id: B496FF00000015636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.17]
  2001:db8:bbbb:a::1/128
    00:13:01:00:00:07          N/A             IPoE-DHCP6      Fwding
      ipoe-ds-7
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000014636FE94C
Acct-Q-Inst-Session-Id: B496FF00000015636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.18]
  10.10.0.16
    00:13:01:00:00:08          N/A             DHCP            Fwding
      ipoe-ds-8
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000024636FE950
Acct-Q-Inst-Session-Id: B496FF00000009636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.18]
  2001:db8:bbbb:6::1/128
    00:13:01:00:00:08          N/A             IPoE-DHCP6      Fwding
      ipoe-ds-8
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000008636FE94C
Acct-Q-Inst-Session-Id: B496FF00000009636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.19]
  10.10.0.15
    00:13:01:00:00:09          N/A             DHCP            Fwding
      ipoe-ds-9
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000020636FE94C
Acct-Q-Inst-Session-Id: B496FF00000021636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.19]
  2001:db8:bbbb:12::1/128
    00:13:01:00:00:09          N/A             IPoE-DHCP6      Fwding
      ipoe-ds-9
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000023636FE94E
Acct-Q-Inst-Session-Id: B496FF00000021636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.20]
  10.10.0.13
    00:13:01:00:00:0a          N/A             DHCP            Fwding
      ipoe-ds-10
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000016636FE94C
Acct-Q-Inst-Session-Id: B496FF00000017636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.20]
  2001:db8:bbbb:10::1/128
    00:13:01:00:00:0a          N/A             IPoE-DHCP6      Fwding
      ipoe-ds-10
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000022636FE94E
Acct-Q-Inst-Session-Id: B496FF00000017636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
Number of subscriber hosts : 40
===============================================================================

[/]
A:admin@bng-1# /show  service id 10 ipoe session

===============================================================================
IPoE sessions for svc-id 10
===============================================================================
Sap Id                          Mac Address         Up Time         MC-Stdby
    Subscriber-Id
        [CircuitID] | [RemoteID]
-------------------------------------------------------------------------------
[pw-1:2.11]                     00:13:01:00:00:01   2d 02:51:27
    ipoe-ds-1
[pw-1:2.12]                     00:13:01:00:00:02   2d 02:51:27
    ipoe-ds-2
[pw-1:2.13]                     00:13:01:00:00:03   2d 02:51:26
    ipoe-ds-3
[pw-1:2.14]                     00:13:01:00:00:04   2d 02:51:26
    ipoe-ds-4
[pw-1:2.15]                     00:13:01:00:00:05   2d 02:51:27
    ipoe-ds-5
[pw-1:2.16]                     00:13:01:00:00:06   2d 02:51:26
    ipoe-ds-6
[pw-1:2.17]                     00:13:01:00:00:07   2d 02:51:26
    ipoe-ds-7
[pw-1:2.18]                     00:13:01:00:00:08   2d 02:51:26
    ipoe-ds-8
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
[pw-1:2.19]                      00:13:01:00:00:09   2d 02:51:26
    ipoe-ds-9
[pw-1:2.20]                      00:13:01:00:00:0a   2d 02:51:26
    ipoe-ds-10
-------------------------------------------------------------------------------
CID | RID displayed when included in session-key
Number of sessions : 10
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  service id 10 pppoe session

===============================================================================
PPPoE sessions for svc-id 10
===============================================================================
Sap Id              Mac Address       Sid    Up Time         Type
    IP/L2TP-Id/Interface-Id                                    MC-Stdby
-------------------------------------------------------------------------------
[pw-1:2.1]          00:11:01:00:00:01 1      2d 02:51:21     local
    10.10.1.10
    00:01:00:02:00:01:00:01
[pw-1:2.2]          00:11:01:00:00:02 1      2d 02:51:21     local
    10.10.1.11
    00:02:00:02:00:02:00:01
[pw-1:2.3]          00:11:01:00:00:03 1      2d 02:51:21     local
    10.10.1.12
    00:03:00:02:00:03:00:01
[pw-1:2.4]          00:11:01:00:00:04 1      2d 02:51:21     local
    10.10.1.13
    00:04:00:02:00:04:00:01
[pw-1:2.5]          00:11:01:00:00:05 1      2d 02:51:21     local
    10.10.1.14
    00:05:00:02:00:05:00:01
[pw-1:2.6]          00:11:01:00:00:06 1      2d 02:51:21     local
    10.10.1.15
    00:06:00:02:00:06:00:01
[pw-1:2.7]          00:11:01:00:00:07 1      2d 02:51:21     local
    10.10.1.16
    00:07:00:02:00:07:00:01
[pw-1:2.8]          00:11:01:00:00:08 1      2d 02:51:21     local
    10.10.1.17
    00:08:00:02:00:08:00:01
[pw-1:2.9]          00:11:01:00:00:09 1      2d 02:51:21     local
    10.10.1.18
    00:09:00:02:00:09:00:01
[pw-1:2.10]         00:11:01:00:00:0a 1      2d 02:51:21     local
    10.10.1.19
    00:0A:00:02:00:0A:00:01
-------------------------------------------------------------------------------
Number of sessions   : 10
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  service id 10 dhcp lease-state

===============================================================================
DHCP lease state table, service 10
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
IP Address      Mac Address       Sap/Sdp Id          Remaining  Lease    MC
                                                      LeaseTime  Origin   Stdby
-------------------------------------------------------------------------------
10.10.0.10      00:13:01:00:00:02 [pw-1:2.12]         00h13m42s  DHCP
10.10.0.11      00:13:01:00:00:05 [pw-1:2.15]         00h13m42s  DHCP
10.10.0.12      00:13:01:00:00:01 [pw-1:2.11]         00h13m42s  DHCP
10.10.0.13      00:13:01:00:00:0a [pw-1:2.20]         00h13m42s  DHCP
10.10.0.14      00:13:01:00:00:03 [pw-1:2.13]         00h13m42s  DHCP
10.10.0.15      00:13:01:00:00:09 [pw-1:2.19]         00h13m42s  DHCP
10.10.0.16      00:13:01:00:00:08 [pw-1:2.18]         00h13m43s  DHCP
10.10.0.17      00:13:01:00:00:04 [pw-1:2.14]         00h13m44s  DHCP
10.10.0.18      00:13:01:00:00:07 [pw-1:2.17]         00h13m44s  DHCP
10.10.0.19      00:13:01:00:00:06 [pw-1:2.16]         00h13m44s  DHCP
-------------------------------------------------------------------------------
Number of lease states : 10
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  service id 10 dhcp6 lease-state

===============================================================================
DHCP lease state table, service 10
===============================================================================
IP Address      Mac Address       Sap/Sdp Id          Remaining  Lease    MC
                                                      LeaseTime  Origin   Stdby
-------------------------------------------------------------------------------
2001:db8:bbbb::1/128
                00:13:01:00:00:01 [pw-1:2.11]         23h43m37s  DHCP
2001:db8:bbbb:1::/64
                00:13:01:00:00:01 [pw-1:2.11]         23h43m37s  DHCP
2001:db8:bbbb:2::1/128
                00:13:01:00:00:02 [pw-1:2.12]         23h43m36s  DHCP
2001:db8:bbbb:3::/64
                00:13:01:00:00:02 [pw-1:2.12]         23h43m36s  DHCP
2001:db8:bbbb:4::1/128
                00:13:01:00:00:05 [pw-1:2.15]         23h43m36s  DHCP
2001:db8:bbbb:5::/64
                00:13:01:00:00:05 [pw-1:2.15]         23h43m36s  DHCP
2001:db8:bbbb:6::1/128
                00:13:01:00:00:08 [pw-1:2.18]         23h43m36s  DHCP
2001:db8:bbbb:7::/64
                00:13:01:00:00:08 [pw-1:2.18]         23h43m36s  DHCP
2001:db8:bbbb:8::1/128
                00:13:01:00:00:04 [pw-1:2.14]         23h43m37s  DHCP
2001:db8:bbbb:9::/64
                00:13:01:00:00:04 [pw-1:2.14]         23h43m37s  DHCP
2001:db8:bbbb:a::1/128
                00:13:01:00:00:07 [pw-1:2.17]         23h43m37s  DHCP
2001:db8:bbbb:b::/64
                00:13:01:00:00:07 [pw-1:2.17]         23h43m37s  DHCP
2001:db8:bbbb:c::1/128
                00:13:01:00:00:03 [pw-1:2.13]         23h43m37s  DHCP
2001:db8:bbbb:d::/64
                00:13:01:00:00:03 [pw-1:2.13]         23h43m37s  DHCP
2001:db8:bbbb:e::1/128
                00:13:01:00:00:06 [pw-1:2.16]         23h43m37s  DHCP
2001:db8:bbbb:f::/64
                00:13:01:00:00:06 [pw-1:2.16]         23h43m37s  DHCP
2001:db8:bbbb:10::1/128
                00:13:01:00:00:0a [pw-1:2.20]         23h38m40s  DHCP
2001:db8:bbbb:11::/64
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                00:13:01:00:00:0a [pw-1:2.20]         23h38m40s  DHCP
2001:db8:bbbb:12::1/128
                00:13:01:00:00:09 [pw-1:2.19]         23h38m40s  DHCP
2001:db8:bbbb:13::/64
                00:13:01:00:00:09 [pw-1:2.19]         23h38m40s  DHCP
2001:db8:bbbb:14::1/128
                00:11:01:00:00:01 [pw-1:2.1]          23h43m41s  DHCP
2001:db8:bbbb:15::/64
                00:11:01:00:00:01 [pw-1:2.1]          23h43m41s  DHCP
2001:db8:bbbb:16::1/128
                00:11:01:00:00:03 [pw-1:2.3]          23h43m41s  DHCP
2001:db8:bbbb:17::/64
                00:11:01:00:00:03 [pw-1:2.3]          23h43m41s  DHCP
2001:db8:bbbb:18::1/128
                00:11:01:00:00:02 [pw-1:2.2]          23h43m41s  DHCP
2001:db8:bbbb:19::/64
                00:11:01:00:00:02 [pw-1:2.2]          23h43m41s  DHCP
2001:db8:bbbb:1a::1/128
                00:11:01:00:00:04 [pw-1:2.4]          23h43m41s  DHCP
2001:db8:bbbb:1b::/64
                00:11:01:00:00:04 [pw-1:2.4]          23h43m41s  DHCP
2001:db8:bbbb:1c::1/128
                00:11:01:00:00:05 [pw-1:2.5]          23h43m42s  DHCP
2001:db8:bbbb:1d::/64
                00:11:01:00:00:05 [pw-1:2.5]          23h43m42s  DHCP
2001:db8:bbbb:1e::1/128
                00:11:01:00:00:06 [pw-1:2.6]          23h43m41s  DHCP
2001:db8:bbbb:1f::/64
                00:11:01:00:00:06 [pw-1:2.6]          23h43m41s  DHCP
2001:db8:bbbb:20::1/128
                00:11:01:00:00:07 [pw-1:2.7]          23h43m41s  DHCP
2001:db8:bbbb:21::/64
                00:11:01:00:00:07 [pw-1:2.7]          23h43m41s  DHCP
2001:db8:bbbb:22::1/128
                00:11:01:00:00:08 [pw-1:2.8]          23h43m42s  DHCP
2001:db8:bbbb:23::/64
                00:11:01:00:00:08 [pw-1:2.8]          23h43m42s  DHCP
2001:db8:bbbb:24::1/128
                00:11:01:00:00:09 [pw-1:2.9]          23h43m41s  DHCP
2001:db8:bbbb:25::/64
                00:11:01:00:00:09 [pw-1:2.9]          23h43m41s  DHCP
2001:db8:bbbb:26::1/128
                00:11:01:00:00:0a [pw-1:2.10]         23h43m41s  DHCP
2001:db8:bbbb:27::/64
                00:11:01:00:00:0a [pw-1:2.10]         23h43m41s  DHCP
-------------------------------------------------------------------------------
Number of lease states : 40
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  srrp 1 detail

===============================================================================
SRRP Instance 1
===============================================================================
Description         : (Not Specified)
Admin State         : Up                Oper State       : master
Preempt             : yes               One GARP per SAP : no
Monitor Oper Group  : None
System IP           : 192.0.2.20
Service ID          : VPRN 10
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Group If            : group-int-1        MAC Address      : b4:96:ff:00:00:00
Grp If Description  : N/A
Grp If Admin State  : Up                 Grp If Oper State: Up
Subscriber If       : sub-int-1
Sub If Admin State  : Up                 Sub If Oper State: Up
Address             : 10.10.0.1/24       Gateway IP       : 10.10.0.254
Address             : 10.10.1.1/24       Gateway IP       : 10.10.1.254
Redundant If        : red-int-bng-1-bng*
Red If Admin State  : Up                 Red If Oper State: Up
Address             : 192.168.11.1/24
Red Spoke-sdp       : 1:1
Msg Path SAP        : pw-1:2.4094        Passive          : no
Admin Gateway MAC   :                    Oper Gateway MAC : 00:00:5e:00:01:01
Standby-Forwarding  : Disabled
Config Priority     : 100                In-use Priority  : 100
Master Priority     : 100
Keep-alive Interval : 2 deci-seconds     Master Since     : 11/12/2022 16:23:34
Fib Population Mode : all
VRRP Policy 1       : None               VRRP Policy 2    : None

-------------------------------------------------------------------------------
Statistics
-------------------------------------------------------------------------------
Become Master       : 1                  Master Changes   : 1
Become Bkup Routing : 0                  Become Bkup Shunt: 1
Become Non-Master   : 0
Adv Sent            : 849384             Adv Received     : 0
Pri 0 Pkts Sent     : 0                  Pri 0 Pkts Rcvd  : 0
Preempt Events      : 0                  Preempted Events : 0
Mesg Intvl Discards : 0                  Mesg Intvl Errors: 0

===============================================================================
* indicates that the corresponding row element may have been truncated.

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  router 10 dhcp local-dhcp-server "dhcpv4" summary
===============================================================================
DHCP server dhcpv4  router 10
===============================================================================
Admin State            : inService
Operational State      : inService
Persistency State      : shutdown
User Data Base         : N/A
Use gateway IP address : enabled (scope pool)
Use pool from client   : enabled
Send force-renewals    : disabled
Creation Origin        : manual
Lease Hold Time        : 0h0m0s
Lease Hold Time For    : N/A
User-ident             : mac-circuit-id

Failover Admin State   : outOfService
Failover Oper State    : shutdown
Failover Persist Key   : N/A
Administrative MCLT    : 0h10m0s
Operational MCLT       : 0h10m0s
Startup wait time      : 0h2m0s
Partner down delay     : 23h59m59s
  Ignore MCLT          : disabled

-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Pool name : dhcpv4-1
-------------------------------------------------------------------------------
Failover Admin State   : inService
Failover Oper State    : normal
Failover Persist Key   : N/A
Administrative MCLT    : 0h10m0s
Operational MCLT       : 0h10m0s
Startup wait time      : 0h2m0s
Partner down delay     : 23h59m59s
  Ignore MCLT          : disabled
-------------------------------------------------------------------------------
Subnet                 Free     %    Stable   Declined Offered  Rem-pend Drain
-------------------------------------------------------------------------------
10.10.0.0/24       (A) 81       89%  10       0        0        0        N

Totals for pool        81       89%  10       0        0        0
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
Pool name : pppoev4-1
-------------------------------------------------------------------------------
Failover Admin State   : outOfService
Failover Oper State    : shutdown
Failover Persist Key   : N/A
Administrative MCLT    : 0h10m0s
Operational MCLT       : 0h10m0s
Startup wait time      : 0h2m0s
Partner down delay     : 23h59m59s
  Ignore MCLT          : disabled
-------------------------------------------------------------------------------
Subnet                 Free     %    Stable   Declined Offered  Rem-pend Drain
-------------------------------------------------------------------------------
10.10.1.0/24           81       89%  10       0        0        0        N

Totals for pool        81       89%  10       0        0        0
-------------------------------------------------------------------------------

Totals for server      162      89%  20       0        0        0

-------------------------------------------------------------------------------
Interface associations
Interface                       Admin
-------------------------------------------------------------------------------
loopback-1                      Up

-------------------------------------------------------------------------------
Local Address Assignment associations
Group interface                 Admin
-------------------------------------------------------------------------------
group-int-1                     Up
No associated firewall domains found.
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  router 10 dhcp6 local-dhcp-server "dhcpv6" summary
===============================================================================
DHCP server dhcpv6  router 10
===============================================================================
Admin State            : inService
Operational State      : inService
Persistency State      : shutdown
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Use Link Address       : disabled
Use pool from client   : enabled
Creation Origin        : manual
Lease Hold Time        : 0h0m0s
Lease Hold Time For    : N/A
User-ident             : duid
Interface-id-mapping   : disabled
Ignore-rapid-commit    : disabled
Allow-lease-query      : disabled
Auto-provisioned       : false

Failover Admin State   : outOfService
Failover Oper State    : shutdown
Failover Persist Key   : N/A
Administrative MCLT    : 0h10m0s
Operational MCLT       : 0h10m0s
Startup wait time      : 0h2m0s
Partner down delay     : 23h59m59s
  Ignore MCLT          : disabled

-------------------------------------------------------------------------------
Pool name : dhcpv6-1
-------------------------------------------------------------------------------
Failover Admin State   : inService
Failover Oper State    : normal
Failover Persist Key   : N/A
Administrative MCLT    : 0h10m0s
Operational MCLT       : 0h10m0s
Startup wait time      : 0h2m0s
Partner down delay     : 23h59m59s
  Ignore MCLT          : disabled
-------------------------------------------------------------------------------
Prefix
                                    Stable   Declined Advert   Rem-pend Drain
-------------------------------------------------------------------------------
2001:db8:bbbb::/56
                               (A)  40       0        0        0        N
2001:db8:bbbb:100::/56
                               (A)  0        0        0        0        N

Totals for pool                     40       0        0        0
-------------------------------------------------------------------------------

Totals for server                   40       0        0        0

-------------------------------------------------------------------------------
Interface associations
Interface                       Admin
-------------------------------------------------------------------------------
loopback-1                      Up

-------------------------------------------------------------------------------
Local Address Assignment associations
Group interface                 Admin
-------------------------------------------------------------------------------
No associated firewall domains found.
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  router 10 dhcp local-dhcp-server "dhcpv4" leases
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Leases for DHCP server dhcpv4 router 10
===============================================================================
IP Address      Lease State       Mac Address       Remaining   Clnt  Fail
  PPP user name/Opt82 Circuit Id                    LifeTime    Type  Ctrl
  User-db/Sticky-lease Hostname
-------------------------------------------------------------------------------
10.10.0.10      stable            00:13:01:00:00:02 0h23m42s    dhcp  local
  cid-2
10.10.0.11      stable            00:13:01:00:00:05 0h23m42s    dhcp  local
  cid-5
10.10.0.12      stable            00:13:01:00:00:01 0h23m42s    dhcp  local
  cid-1
10.10.0.13      stable            00:13:01:00:00:0a 0h23m42s    dhcp  local
  cid-10
10.10.0.14      stable            00:13:01:00:00:03 0h23m42s    dhcp  local
  cid-3
10.10.0.15      stable            00:13:01:00:00:09 0h23m42s    dhcp  local
  cid-9
10.10.0.16      stable            00:13:01:00:00:08 0h23m43s    dhcp  local
  cid-8
10.10.0.17      stable            00:13:01:00:00:04 0h23m44s    dhcp  local
  cid-4
10.10.0.18      stable            00:13:01:00:00:07 0h23m44s    dhcp  local
  cid-7
10.10.0.19      stable            00:13:01:00:00:06 0h23m44s    dhcp  local
  cid-6
10.10.1.10      internal          N/A               N/A         ppp   N/A

10.10.1.11      internal          N/A               N/A         ppp   N/A

10.10.1.12      internal          N/A               N/A         ppp   N/A

10.10.1.13      internal          N/A               N/A         ppp   N/A

10.10.1.14      internal          N/A               N/A         ppp   N/A

10.10.1.15      internal          N/A               N/A         ppp   N/A

10.10.1.16      internal          N/A               N/A         ppp   N/A

10.10.1.17      internal          N/A               N/A         ppp   N/A

10.10.1.18      internal          N/A               N/A         ppp   N/A

10.10.1.19      internal          N/A               N/A         ppp   N/A

-------------------------------------------------------------------------------
20 leases found
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  router 10 dhcp6 local-dhcp-server "dhcpv6" leases

===============================================================================
Leases for DHCPv6 server dhcpv6
===============================================================================
IP Address/Prefix                          Lease State       Remaining   Fail
  Link-local Address                                         LifeTime    Ctrl
-------------------------------------------------------------------------------
2001:db8:bbbb::1/64
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
  fe80::213:1ff:fe00:1                    stable           1d0h13m     local
2001:db8:bbbb:1::/64
  fe80::213:1ff:fe00:1                    stable           1d0h13m     local
2001:db8:bbbb:2::1/64
  fe80::213:1ff:fe00:2                    stable           1d0h13m     local
2001:db8:bbbb:3::/64
  fe80::213:1ff:fe00:2                    stable           1d0h13m     local
2001:db8:bbbb:4::1/64
  fe80::213:1ff:fe00:5                    stable           1d0h13m     local
2001:db8:bbbb:5::/64
  fe80::213:1ff:fe00:5                    stable           1d0h13m     local
2001:db8:bbbb:6::1/64
  fe80::213:1ff:fe00:8                    stable           1d0h13m     local
2001:db8:bbbb:7::/64
  fe80::213:1ff:fe00:8                    stable           1d0h13m     local
2001:db8:bbbb:8::1/64
  fe80::213:1ff:fe00:4                    stable           1d0h13m     local
2001:db8:bbbb:9::/64
  fe80::213:1ff:fe00:4                    stable           1d0h13m     local
2001:db8:bbbb:a::1/64
  fe80::213:1ff:fe00:7                    stable           1d0h13m     local
2001:db8:bbbb:b::/64
  fe80::213:1ff:fe00:7                    stable           1d0h13m     local
2001:db8:bbbb:c::1/64
  fe80::213:1ff:fe00:3                    stable           1d0h13m     local
2001:db8:bbbb:d::/64
  fe80::213:1ff:fe00:3                    stable           1d0h13m     local
2001:db8:bbbb:e::1/64
  fe80::213:1ff:fe00:6                    stable           1d0h13m     local
2001:db8:bbbb:f::/64
  fe80::213:1ff:fe00:6                    stable           1d0h13m     local
2001:db8:bbbb:10::1/64
  fe80::213:1ff:fe00:a                    stable           1d0h8m      local
2001:db8:bbbb:11::/64
  fe80::213:1ff:fe00:a                    stable           1d0h8m      local
2001:db8:bbbb:12::1/64
  fe80::213:1ff:fe00:9                    stable           1d0h8m      local
2001:db8:bbbb:13::/64
  fe80::213:1ff:fe00:9                    stable           1d0h8m      local
2001:db8:bbbb:14::1/64
  fe80::1:2:1:1                           stable           1d0h13m     local
2001:db8:bbbb:15::/64
  fe80::1:2:1:1                           stable           1d0h13m     local
2001:db8:bbbb:16::1/64
  fe80::3:2:3:1                           stable           1d0h13m     local
2001:db8:bbbb:17::/64
  fe80::3:2:3:1                           stable           1d0h13m     local
2001:db8:bbbb:18::1/64
  fe80::2:2:2:1                           stable           1d0h13m     local
2001:db8:bbbb:19::/64
  fe80::2:2:2:1                           stable           1d0h13m     local
2001:db8:bbbb:1a::1/64
  fe80::4:2:4:1                           stable           1d0h13m     local
2001:db8:bbbb:1b::/64
  fe80::4:2:4:1                           stable           1d0h13m     local
2001:db8:bbbb:1c::1/64
  fe80::5:2:5:1                           stable           1d0h13m     local
2001:db8:bbbb:1d::/64
  fe80::5:2:5:1                           stable           1d0h13m     local
2001:db8:bbbb:1e::1/64
  fe80::6:2:6:1                           stable           1d0h13m     local
2001:db8:bbbb:1f::/64
  fe80::6:2:6:1                           stable           1d0h13m     local
2001:db8:bbbb:20::1/64
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
  fe80::7:2:7:1                           stable           1d0h13m     local
2001:db8:bbbb:21::/64
  fe80::7:2:7:1                           stable           1d0h13m     local
2001:db8:bbbb:22::1/64
  fe80::8:2:8:1                           stable           1d0h13m     local
2001:db8:bbbb:23::/64
  fe80::8:2:8:1                           stable           1d0h13m     local
2001:db8:bbbb:24::1/64
  fe80::9:2:9:1                           stable           1d0h13m     local
2001:db8:bbbb:25::/64
  fe80::9:2:9:1                           stable           1d0h13m     local
2001:db8:bbbb:26::1/64
  fe80::a:2:a:1                           stable           1d0h13m     local
2001:db8:bbbb:27::/64
  fe80::a:2:a:1                           stable           1d0h13m     local
-------------------------------------------------------------------------------
40 leases found
===============================================================================A:admin@bng-1# /show
 port

===============================================================================
Ports on Slot 1
===============================================================================
Port          Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id            State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
1/1/c1        Up         Link Up                        conn   100GBASE-LR4*
1/1/c1/1      Up    Yes  Up      8936 8936    - hybr dotq xgige
1/1/c1/2      Down  No   Down    8936 8936    - hybr dotq xgige
1/1/c1/3      Down  No   Down    8936 8936    - hybr dotq xgige
1/1/c1/4      Up    Yes  Up      8936 8936    - hybr dotq xgige
1/1/c1/5      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c1/6      Up    Yes  Up      8936 8936    - hybr dotq xgige
1/1/c1/7      Up    Yes  Up      8936 8936    - hybr dotq xgige
1/1/c1/8      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c1/9      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c1/10     Up    Yes  Up      8936 8936    - hybr dotq xgige
1/1/c2        Up         Link Up                        conn   100G CWDM4 M*
1/1/c2/1      Up    Yes  Up      8936 8936    - hybr dotq cgige
1/1/c3        Down       Down                           conn   100G CLR4   *
1/1/c4        Up         Link Up                        conn   100GBASE-LR4*
1/1/c4/1      Up    Yes  Up      8936 8936    - hybr dotq xgige
1/1/c4/2      Up    No   Down    8936 8936    - netw null xgige
1/1/c4/3      Up    No   Down    8936 8936    - netw null xgige
1/1/c4/4      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c5        Down       Down                           conn   100G CWDM4 M*
1/1/c6        Down       Down                           conn   100G CLR4   *
1/2/c1        Down       Down                           conn   100GBASE-LR4*
1/2/c2        Down       Down                           conn   100G CWDM4 M*
1/2/c3        Down       Down                           conn   100G CLR4   *
1/2/c4        Down       Down                           conn   100GBASE-LR4*
1/2/c5        Down       Down                           conn   100G CWDM4 M*
1/2/c6        Down       Down                           conn   100G CLR4   *

===============================================================================
Ports on Slot A
===============================================================================
Port          Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id            State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
A/1           Up    Yes  Up      1514 1514    - netw null faste  MDI
A/3           Down  No   Down    1514 1514    - netw null faste
A/4           Down  No   Down    1514 1514    - netw null faste
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Ports on Slot B
===============================================================================
Port          Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id            State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
B/1           Up    No   Ghost   1514 1514    - netw null faste
B/3           Down  No   Ghost   1514 1514    - netw null faste
B/4           Down  No   Ghost   1514 1514    - netw null faste

===============================================================================
Ports on Port Cross Connect 1
===============================================================================
Port          Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id            State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
pxc-1.a       Up    Yes  Up      8932 8932    - hybr dotq xgige
pxc-1.b       Up    Yes  Up      8932 8932    - hybr dotq xgige

===============================================================================
Ports on Port Cross Connect 2
===============================================================================
Port          Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id            State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
pxc-2.a       Up    Yes  Up      8932 8932    - hybr dotq xgige
pxc-2.b       Up    Yes  Up      8932 8932    - hybr dotq xgige

===============================================================================
Ports on Port Cross Connect 3
===============================================================================
Port          Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id            State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
pxc-3.a       Up    Yes  Up      8932 8932    - hybr dotq xgige
pxc-3.b       Up    Yes  Up      8932 8932    - hybr dotq xgige
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show router interface

===============================================================================
Interface Table (Router: Base)
===============================================================================
Interface-Name                   Adm       Opr(v4/v6)  Mode    Port/SapId
   IP-Address                                                   PfxState
-------------------------------------------------------------------------------
_tmnx_fpe_1.a                    Up        Up/Up       Network pxc-1.a:1
   fe80::100/64                                                 PREFERRED
_tmnx_fpe_1.b                    Up        Up/Up       Network pxc-1.b:1
   fe80::101/64                                                 PREFERRED
_tmnx_fpe_2.a                    Up        Up/Up       Network pxc-2.a:1
   fe80::200/64                                                 PREFERRED
_tmnx_fpe_2.b                    Up        Up/Up       Network pxc-2.b:1
   fe80::201/64                                                 PREFERRED
_tmnx_fpe_3.a                    Up        Up/Up       Network pxc-3.a:1
   fe80::300/64                                                 PREFERRED
_tmnx_fpe_3.b                    Up        Up/Up       Network pxc-3.b:1
   fe80::301/64                                                 PREFERRED
int-1-bng-1-bng-2                Up        Up/Up       Network 1/1/c1/6:1
   192.168.7.1/24                                               n/a
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
   2001:db8::701/120                                          PREFERRED
   fe80::b696:ffff:fe00:0/64                                  PREFERRED
int-1-bng-1-p-1                  Up        Down/Up     Network 1/1/c1/1:1
   2001:db8::501/120                                          PREFERRED
   fe80::b696:ffff:fe00:0/64                                  PREFERRED
int-1-bng-1-pe-2                 Up        Down/Up     Network 1/1/c1/7:1
   2001:db8::801/120                                          PREFERRED
   fe80::b696:ffff:fe00:0/64                                  PREFERRED
system                           Up        Up/Up       Network system
   192.0.2.20/32                                              n/a
   2001:db8::14/128                                           PREFERRED
to-radius                        Up        Up/Down     Network 1/1/c1/10:114
   192.168.114.20/24                                          n/a
-------------------------------------------------------------------------------
Interfaces : 11
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show router 10 interface

===============================================================================
Interface Table (Service: 10)
===============================================================================
Interface-Name                   Adm       Opr(v4/v6)  Mode    Port/SapId
   IP-Address                                                   PfxState
-------------------------------------------------------------------------------
group-int-1                      Up        Up/Up       VPRN G* pw-1
loopback-1                       Up        Up/Up       VPRN    loopback
   192.168.0.1/32                                             n/a
   2001:db8::1/128                                            PREFERRED
   fe80::b696:ffff:fe00:0/64                                  PREFERRED
red-int-bng-1-bng-2              Up        Up/Up       VPRN R* spoke-1:1
   192.168.11.1/24                                            n/a
sub-int-1                        Up        Up/Up       VPRN S* subscriber
   10.10.0.1/24                                               n/a
   10.10.1.1/24                                               n/a
   2001:db8:bbbb::/56                                         PREFERRED
   2001:db8:bbbb:100::/56                                     PREFERRED
   fe80::b696:ffff:fe00:0/64                                  PREFERRED
-------------------------------------------------------------------------------
Interfaces : 4
===============================================================================
* indicates that the corresponding row element may have been truncated.

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  fwd-path-ext fpe 1

===============================================================================
FPE Id: 1
===============================================================================
Description        : srv6 origination
Multi-Path         : Disabled
Path               : pxc 1
Pw Port Extension  : Disabled                          Oper     : down
Sub Mgmt Extension : Disabled                          Oper     : N/A
Vxlan              : Disabled                          Oper     : down
Segment-Routing V6 : Enabled                           Oper     : up
SRv6 Type          : origination
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
If-A Qos Policy    : default
If-B MTU           : 9786 bytes                        Oper MTU : 8914 bytes
If-B Qos Policy    : default
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  fwd-path-ext fpe 2

===============================================================================
FPE Id: 2
===============================================================================
Description        : srv6 termination
Multi-Path         : Disabled
Path               : pxc 2
Pw Port Extension  : Disabled                          Oper     : down
Sub Mgmt Extension : Disabled                          Oper     : N/A
Vxlan              : Disabled                          Oper     : down
Segment-Routing V6 : Enabled                           Oper     : up
SRv6 Type          : termination
If-A Qos Policy    : default
If-B MTU           : 0 bytes                           Oper MTU : 8914 bytes
If-B Qos Policy    : default
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  fwd-path-ext fpe 3

===============================================================================
FPE Id: 3
===============================================================================
Description        : pw-port on a single port
Multi-Path         : Disabled
Path               : pxc 3
Pw Port Extension  : Enabled                           Oper     : up
Sub Mgmt Extension : Disabled                          Oper     : N/A
Vxlan              : Disabled                          Oper     : down
Segment-Routing V6 : Disabled
If-A Qos Policy    : default
If-B Qos Policy    : default
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show service id "evpn-dual-homing" base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 11                 Vpn Id            : 0
Service Type      : Epipe
MACSec enabled    : no
Name              : evpn-dual-homing
Description       : (Not Specified)
Customer Id       : 1                  Creation Origin   : manual
Last Status Change: 11/12/2022 16:22:36
Last Mgmt Change  : 11/12/2022 16:21:35
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Test Service      : No
Admin State       : Up                 Oper State        : Up
MTU               : 1514
Vc Switching      : False
SAP Count         : 0                  SDP Bind Count    : 1
Per Svc Hashing   : Disabled           Lbl Eth/IP L4 TEID: Disabled
Ignore MTU Mismat*: Disabled
Vxlan Src Tep Ip  : N/A
Force QTag Fwd    : Disabled
Lcl Switch Svc St : sap
Oper Group        : <none>

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                            Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sdp:17500:1 SB(fpe_3.b)                Fpe          0       8910    Up   Up
===============================================================================
* indicates that the corresponding row element may have been truncated.

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show service id "evpn-dual-homing" bgp-evpn

===============================================================================
BGP EVPN Table
===============================================================================
EVI               : 11                 Creation Origin    : manual

-------------------------------------------------------------------------------
Local AC Name              Eth Tag  Endpoint                   Ingress Label
-------------------------------------------------------------------------------
bng                        2                                   0
-------------------------------------------------------------------------------
Number of local ACs : 1

-------------------------------------------------------------------------------
Remote AC Name                     Eth Tag  Endpoint
-------------------------------------------------------------------------------
access                             1
-------------------------------------------------------------------------------
Number of Remote ACs : 1
===============================================================================

===============================================================================
Segment Routing v6 Instance 1 Service 11
===============================================================================
Admin State             : Enabled
Srv6 Instance           : 1
Default Locator         : bng-1-loc

Oper Group              : (Not Specified)
Default Route Tag       : 0xb
Source Address          : 2001:db8::14
ECMP                    : 1
Force Vlan VC Fwd       : disabled
Next Hop Type           : system-ipv6
Evi 3-byte Auto-RT      : disabled
Route Resolution        : fallback-tunnel-to-route-table
Force QinQ VC Fwd       : none
MH Mode                 : network
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show service system bgp-evpn ethernet-segment name "ES-1"

===============================================================================
Service Ethernet Segment
===============================================================================
Name                    : ES-1
Eth Seg Type            : None
Admin State             : Enabled            Oper State         : Up
ESI                     : 01:01:01:01:01:01:01:01:01:01
Oper ESI                : 01:01:01:01:01:01:01:01:01:01
Auto-ESI Type           : None
AC DF Capability        : Include
Multi-homing            : singleActive       Oper Multi-homing  : singleActive
ES SHG Label            : None
Source BMAC LSB         : None
PW Port Id              : 1
PW Port Headend         : enabled
ES Activation Timer     : 0 secs
Oper Group              : ES-1
Svc Carving             : manual             Oper Svc Carving   : manual
Cfg Range Type          : lowest-pref

-------------------------------------------------------------------------------
DF Pref Election Information
-------------------------------------------------------------------------------
Preference     Preference     Last Admin Change        Oper Pref      Do No
Mode           Value                                   Value          Preempt
-------------------------------------------------------------------------------
revertive      150            11/12/2022 16:21:35      150            Disabled
-------------------------------------------------------------------------------
EVI Ranges: <none>
ISID Ranges: <none>
Vprn NextHop EVI Ranges : <none>
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show service system bgp-evpn ethernet-segment name "ES-1" evi evi-1 11

===============================================================================
EVI DF and Candidate List
===============================================================================
EVI           SvcId         Actv Timer Rem      DF  DF Last Change
-------------------------------------------------------------------------------
11            11            0                   yes 11/12/2022 16:21:35
===============================================================================

===============================================================================
DF Candidates                           Time Added           Oper Pref  Do Not
                                                               Value    Preempt
-------------------------------------------------------------------------------
2001:db8::14                            11/12/2022 16:23:33  150        Disabl*
2001:db8::15                            11/12/2022 16:22:36  50         Disabl*
-------------------------------------------------------------------------------
Number of entries: 2
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
* indicates that the corresponding row element may have been truncated.

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show router bgp routes evpn eth-seg detail
===============================================================================
 BGP Router ID:192.0.2.20       AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Eth-Seg Routes
===============================================================================
Original Attributes

Network        : n/a
Nexthop        : 2001:db8::15
Path Id        : None
From           : 2001:db8::15
Res. Nexthop   : fe80::b697:ffff:fe00:0
Local Pref.    : 100                    Interface Name : int-1-bng-1-bng-2
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : 10
Connector      : None
Community      :
                 df-election::DF-Type:Preference/DP:0/DF-Preference:50/AC:1
                 target:01:01:01:01:01:01
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.21
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
EVPN type      : ETH-SEG
ESI            : 01:01:01:01:01:01:01:01:01:01
Originator IP  : 2001:db8::15
Route Dist.    : 192.0.2.21:0
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
Add Paths Send : Default
Last Modified  : 01d23h12m

Modified Attributes

Network        : n/a
Nexthop        : 2001:db8::15
Path Id        : None
From           : 2001:db8::15
Res. Nexthop   : fe80::b697:ffff:fe00:0
Local Pref.    : 100                    Interface Name : int-1-bng-1-bng-2
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : 10
Connector      : None
Community      :
                 df-election::DF-Type:Preference/DP:0/DF-Preference:50/AC:1
                 target:01:01:01:01:01:01
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Cluster        : No Cluster Members
Originator Id  : None                    Peer Router Id : 192.0.2.21
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
EVPN type      : ETH-SEG
ESI            : 01:01:01:01:01:01:01:01:01:01
Originator IP  : 2001:db8::15
Route Dist.    : 192.0.2.21:0
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                       Dest Class     : 0
Add Paths Send : Default
Last Modified  : 01d23h12m

-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Routes : 1
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show router bgp routes evpn auto-disc tag 1 detail
===============================================================================
 BGP Router ID:192.0.2.20       AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Auto-Disc Routes
===============================================================================
Original Attributes

Network        : n/a
Nexthop        : 2001:db8::5
Path Id        : None
From           : 2001:db8::5
Res. Nexthop   : fe80::b69e:ffff:fe00:0
Local Pref.    : 100                    Interface Name : int-1-bng-1-p-1
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : 30
Connector      : None
Community      : color:00:10 target:64500:11
                 l2-attribute:MTU: 1514 C: 0 P: 0 B: 0
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.5
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
EVPN type      : AUTO-DISC
ESI            : ESI-0
Tag            : 1
Route Dist.    : 192.0.2.5:11
MPLS Label     : 524288
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Source Class   : 0                      Dest Class     : 0
Add Paths Send : Default
Last Modified  : 01d23h12m
SRv6 TLV Type  : SRv6 L2 Service TLV (6)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:aaaa:5::
Full Sid       : 2001:db8:aaaa:5:8000::
Behavior       : End.DX2 (21)
SRv6 SubSubTLV : SRv6 SID Structure (1)
Loc-Block-Len  : 48                     Loc-Node-Len   : 16
Func-Len       : 20                     Arg-Len        : 0
Tpose-Len      : 20                     Tpose-offset   : 64

Modified Attributes

Network        : n/a
Nexthop        : 2001:db8::5
Path Id        : None
From           : 2001:db8::5
Res. Nexthop   : fe80::b69e:ffff:fe00:0
Local Pref.    : 100                    Interface Name : int-1-bng-1-p-1
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : 30
Connector      : None
Community      : color:00:10 target:64500:11
                 l2-attribute:MTU: 1514 C: 0 P: 0 B: 0
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.5
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
EVPN type      : AUTO-DISC
ESI            : ESI-0
Tag            : 1
Route Dist.    : 192.0.2.5:11
MPLS Label     : 524288
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
Add Paths Send : Default
Last Modified  : 01d23h12m
SRv6 TLV Type  : SRv6 L2 Service TLV (6)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:aaaa:5::
Full Sid       : 2001:db8:aaaa:5:8000::
Behavior       : End.DX2 (21)
SRv6 SubSubTLV : SRv6 SID Structure (1)
Loc-Block-Len  : 48                     Loc-Node-Len   : 16
Func-Len       : 20                     Arg-Len        : 0
Tpose-Len      : 20                     Tpose-offset   : 64

-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Routes : 1
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  service  id "evpn-dual-homing" ethernet-segment "ES-1"
No sap entries
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
No sdp entries
No vxlan instance entries

===============================================================================
SDP Ethernet-Segment Information
===============================================================================
Pw-Port               Eth-Seg                           Status
-------------------------------------------------------------------------------
1                     ES-1                              DF
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  service id "evpn-dual-homing"  segment-routing-v6 detail

===============================================================================
Segment Routing v6 Instance 1 Service 11
===============================================================================
Locator
  Type         Function  SID                                      Status
-------------------------------------------------------------------------------
bng-1-loc
  End.DX2        *524286 2001:db8:aaaa:14:7fff:e000::            ok
===============================================================================
Legend: * - System allocated

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show service id "evpn-dual-homing" segment-routing-v6 instance 1 destinations

===============================================================================
TEP, SID
===============================================================================
Instance  TEP Address                       Segment Id
-------------------------------------------------------------------------------
1         2001:db8::5                       2001:db8:aaaa:5:8000::
-------------------------------------------------------------------------------
Number of TEP, SID: 1
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
Segment Routing v6 Ethernet Segment Dest
===============================================================================
Instance  Eth SegId                     Num. Macs     Last Change
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show service id "evpn-dual-homing" segment-routing-v6 instance 1 end-dx2

===============================================================================
Segment Routing v6 Instance 1 Service 11
===============================================================================
Locator
  Type         Function  SID                                      Status
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
bng-1-loc
  End.DX2        *524286 2001:db8:aaaa:14:7fff:e000::            ok
===============================================================================
Legend: * - System allocated

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  service id "evpn-dual-homing"  segment-routing-v6 instance 1 locator
 "bng-1-loc"

===============================================================================
Segment Routing v6 Instance 1 Service 11
===============================================================================
Locator
  Type         Function  SID                                          Status
-------------------------------------------------------------------------------
bng-1-loc
  End.DX2        *524286 2001:db8:aaaa:14:7fff:e000::            ok
===============================================================================
Legend: * - System allocated

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show service id "dual-homing" base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 10                  Vpn Id            : 0
Service Type      : VPRN
MACSec enabled    : no
Name              : dual-homing
Description       : (Not Specified)
Customer Id       : 1                   Creation Origin   : manual
Last Status Change: 11/12/2022 16:21:35
Last Mgmt Change  : 11/12/2022 16:21:35
Admin State       : Up                  Oper State        : Up

Router Oper State : Up
Route Dist.       : None                VPRN Type         : regular
Oper Route Dist   : 0:0
Oper RD Type      : none
AS Number         : None                Router Id         : 192.0.2.20
ECMP              : Enabled             ECMP Max Routes   : 1
Max IPv4 Routes   : No Limit
Local Rt Domain-Id: None                D-Path Lng Ignore : Disabled

Auto Bind Tunnel
Allow Flex-Alg-Fb : Disabled
Resolution        : disabled
Weighted ECMP     : Disabled            ECMP Max Routes   : 1
Strict Tnl Tag    : Disabled

Max IPv6 Routes   : No Limit
Ignore NH Metric  : Disabled
Hash Label        : Disabled
Entropy Label     : Disabled
Vrf Target        : None
Vrf Import        : None
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Vrf Export        : None
MVPN Vrf Target   : None
MVPN Vrf Import   : None
MVPN Vrf Export   : None
Car. Sup C-VPN    : Disabled
Label mode        : vrf
BGP VPN Backup    : Disabled
BGP Export Inactv : Disabled
LOG all events    : Disabled

SAP Count         : 21                  SDP Bind Count    : 1
VSD Domain        : <none>

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                              Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:[pw-1:2.1]                          qinq         8932    8888    Up   Up
sap:[pw-1:2.2]                          qinq         8932    8888    Up   Up
sap:[pw-1:2.3]                          qinq         8932    8888    Up   Up
sap:[pw-1:2.4]                          qinq         8932    8888    Up   Up
sap:[pw-1:2.5]                          qinq         8932    8888    Up   Up
sap:[pw-1:2.6]                          qinq         8932    8888    Up   Up
sap:[pw-1:2.7]                          qinq         8932    8888    Up   Up
sap:[pw-1:2.8]                          qinq         8932    8888    Up   Up
sap:[pw-1:2.9]                          qinq         8932    8888    Up   Up
sap:[pw-1:2.10]                         qinq         8932    8888    Up   Up
sap:[pw-1:2.11]                         qinq         8932    8888    Up   Up
sap:[pw-1:2.12]                         qinq         8932    8888    Up   Up
sap:[pw-1:2.13]                         qinq         8932    8888    Up   Up
sap:[pw-1:2.14]                         qinq         8932    8888    Up   Up
sap:[pw-1:2.15]                         qinq         8932    8888    Up   Up
sap:[pw-1:2.16]                         qinq         8932    8888    Up   Up
sap:[pw-1:2.17]                         qinq         8932    8888    Up   Up
sap:[pw-1:2.18]                         qinq         8932    8888    Up   Up
sap:[pw-1:2.19]                         qinq         8932    8888    Up   Up
sap:[pw-1:2.20]                         qinq         8932    8888    Up   Up
sap:pw-1:2.4094                         qinq         8932    8888    Up   Up
sdp:1:1 S(192.168.7.2)                  None         0       8890    Up   Up
-------------------------------------------------------------------------------
[<sap-id>] indicates a Managed SAP
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show service id "dual-homing" bgp-ipvpn segment-routing-v6

===============================================================================
Service 10 BGP-IPVPN Segment-Routing-V6 Information
===============================================================================

Admin State       : Up
VRF Import        : None
VRF Export        : srrp-aware-routing
Route Dist.       : 192.0.2.20:10
Oper Route Dist   : 192.0.2.20:10
Oper RD Type      : configured
Route Target      : target:64500:10
Route Target Expor: None
Route Target Impor: None
Def Route Tag     : 0x0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Route Resolution  : route-table

Srv6 Instance     : 1
Default Locator   : bng-1-loc
Source Address    : 2001:db8::14
Domain-Id         : None

===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show service id "dual-homing" segment-routing-v6 instance 1 locator "bng-1-loc"

===============================================================================
Segment Routing v6 Instance 1 Service 10
===============================================================================
Locator
  Type         Function  SID                                        Status
-------------------------------------------------------------------------------
bng-1-loc
  End.DT4        *524285 2001:db8:aaaa:14:7fff:d000::               ok
  End.DT6        *524284 2001:db8:aaaa:14:7fff:c000::               ok
===============================================================================
Legend: * - System allocated

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  router 10 route-table

===============================================================================
Route Table (Service: 10)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
10.10.0.0/24                                  Local   Local     01d23h11m  0
       sub-int-1                                                  0
10.10.0.2/32                                  Remote  Sub Mgmt  01d23h12m  0
       [red-int-bng-1-bng-2]                                      0
10.10.0.10/32                                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                              0
10.10.0.11/32                                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                              0
10.10.0.12/32                                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                              0
10.10.0.13/32                                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                              0
10.10.0.14/32                                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                              0
10.10.0.15/32                                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                              0
10.10.0.16/32                                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                              0
10.10.0.17/32                                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                              0
10.10.0.18/32                                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                              0
10.10.0.19/32                                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                              0
10.10.1.0/24                                  Local   Local     01d23h11m  0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
       sub-int-1                                                        0
10.10.1.2/32                                 Remote  Sub Mgmt  01d23h12m  0
       [red-int-bng-1-bng-2]                                            0
10.10.1.10/32                                Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                    0
10.10.1.11/32                                Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                    0
10.10.1.12/32                                Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                    0
10.10.1.13/32                                Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                    0
10.10.1.14/32                                Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                    0
10.10.1.15/32                                Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                    0
10.10.1.16/32                                Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                    0
10.10.1.17/32                                Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                    0
10.10.1.18/32                                Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                    0
10.10.1.19/32                                Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                    0
172.16.102.0/24                              Remote  BGP VPN   01d23h12m  170
       2001:db8:aaaa:6:8000:: (tunneled:SRV6)                           20
192.168.0.1/32                               Local   Local     01d23h13m  0
       loopback-1                                                       0
192.168.11.0/24                              Local   Local     01d23h13m  0
       red-int-bng-1-bng-2                                              0
-------------------------------------------------------------------------------
No. of Routes: 27
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  router 10 route-table ipv6

===============================================================================
IPv6 Route Table (Service: 10)
===============================================================================
Dest Prefix[Flags]                           Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
2001:db8::1/128                              Local   Local     01d23h13m  0
       loopback-1                                                       0
2001:db8::6600/120                           Remote  BGP VPN   01d23h12m  170
       2001:db8:aaaa:6:7fff:f000:: (tunneled:SRV6)                      20
2001:db8:bbbb::/56                           Local   Local     01d23h11m  0
       sub-int-1                                                        0
2001:db8:bbbb::1/128                         Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                    0
2001:db8:bbbb:1::/64                         Remote  Managed   01d23h11m  0
       2001:db8:bbbb::1                                                 0
2001:db8:bbbb:2::1/128                       Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                    0
2001:db8:bbbb:3::/64                         Remote  Managed   01d23h11m  0
       2001:db8:bbbb:2::1                                               0
2001:db8:bbbb:4::1/128                       Remote  Sub Mgmt  01d23h11m  0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
       [group-int-1]                                                0
2001:db8:bbbb:5::/64                    Remote  Managed   01d23h11m  0
       2001:db8:bbbb:4::1                                           0
2001:db8:bbbb:6::1/128                  Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                0
2001:db8:bbbb:7::/64                    Remote  Managed   01d23h11m  0
       2001:db8:bbbb:6::1                                           0
2001:db8:bbbb:8::1/128                  Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                0
2001:db8:bbbb:9::/64                    Remote  Managed   01d23h11m  0
       2001:db8:bbbb:8::1                                           0
2001:db8:bbbb:a::1/128                  Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                0
2001:db8:bbbb:b::/64                    Remote  Managed   01d23h11m  0
       2001:db8:bbbb:a::1                                           0
2001:db8:bbbb:c::1/128                  Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                0
2001:db8:bbbb:d::/64                    Remote  Managed   01d23h11m  0
       2001:db8:bbbb:c::1                                           0
2001:db8:bbbb:e::1/128                  Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                0
2001:db8:bbbb:f::/64                    Remote  Managed   01d23h11m  0
       2001:db8:bbbb:e::1                                           0
2001:db8:bbbb:10::1/128                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                0
2001:db8:bbbb:11::/64                   Remote  Managed   01d23h11m  0
       2001:db8:bbbb:10::1                                          0
2001:db8:bbbb:12::1/128                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                0
2001:db8:bbbb:13::/64                   Remote  Managed   01d23h11m  0
       2001:db8:bbbb:12::1                                          0
2001:db8:bbbb:14::1/128                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                0
2001:db8:bbbb:15::/64                   Remote  Managed   01d23h11m  0
       2001:db8:bbbb:14::1                                          0
2001:db8:bbbb:16::1/128                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                0
2001:db8:bbbb:17::/64                   Remote  Managed   01d23h11m  0
       2001:db8:bbbb:16::1                                          0
2001:db8:bbbb:18::1/128                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                0
2001:db8:bbbb:19::/64                   Remote  Managed   01d23h11m  0
       2001:db8:bbbb:18::1                                          0
2001:db8:bbbb:1a::1/128                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                0
2001:db8:bbbb:1b::/64                   Remote  Managed   01d23h11m  0
       2001:db8:bbbb:1a::1                                          0
2001:db8:bbbb:1c::1/128                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                0
2001:db8:bbbb:1d::/64                   Remote  Managed   01d23h11m  0
       2001:db8:bbbb:1c::1                                          0
2001:db8:bbbb:1e::1/128                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                0
2001:db8:bbbb:1f::/64                   Remote  Managed   01d23h11m  0
       2001:db8:bbbb:1e::1                                          0
2001:db8:bbbb:20::1/128                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                0
2001:db8:bbbb:21::/64                   Remote  Managed   01d23h11m  0
       2001:db8:bbbb:20::1                                          0
2001:db8:bbbb:22::1/128                 Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                0
2001:db8:bbbb:23::/64                   Remote  Managed   01d23h11m  0
       2001:db8:bbbb:22::1                                          0
2001:db8:bbbb:24::1/128                 Remote  Sub Mgmt  01d23h11m  0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
       [group-int-1]                                                  0
2001:db8:bbbb:25::/64                        Remote  Managed   01d23h11m  0
       2001:db8:bbbb:24::1                                            0
2001:db8:bbbb:26::1/128                      Remote  Sub Mgmt  01d23h11m  0
       [group-int-1]                                                  0
2001:db8:bbbb:27::/64                        Remote  Managed   01d23h11m  0
       2001:db8:bbbb:26::1                                            0
2001:db8:bbbb:100::/56                       Local   Local     01d23h11m  0
       sub-int-1                                                      0
-------------------------------------------------------------------------------
No. of Routes: 44
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show router bgp routes vpn-ipv4 hunt
===============================================================================
 BGP Router ID:192.0.2.20       AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv4 Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : 172.16.102.0/24
Nexthop        : 2001:db8::6
Route Dist.    : 192.0.2.6:10           VPN Label      : 524288
Path Id        : None
From           : 2001:db8::6
Res. Nexthop   : n/a
Local Pref.    : 100                    Interface Name : int-1-bng-1-pe-2
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : 20
Connector      : None
Community      : target:64500:10
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.6
Fwd Class      : None                   Priority       : None
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
Add Paths Send : Default
Last Modified  : 01d23h12m
SRv6 TLV Type  : SRv6 L3 Service TLV (5)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:aaaa:6::
Full Sid       : 2001:db8:aaaa:6:8000::
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Behavior       : End.DT4 (19)
SRv6 SubSubTLV : SRv6 SID Structure (1)
Loc-Block-Len  : 48                    Loc-Node-Len   : 16
Func-Len       : 20                    Arg-Len        : 0
Tpose-Len      : 20                    Tpose-offset   : 64
VPRN Imported  :  10

-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
Network        : 10.10.0.0/24
Nexthop        : 2001:db8::14
Route Dist.    : 192.0.2.20:10         VPN Label      : 524285
Path Id        : None
To             : 2001:db8::6
Res. Nexthop   : n/a
Local Pref.    : 150                   Interface Name : NotAvailable
Aggregator AS  : None                  Aggregator     : None
Atomic Aggr.   : Not Atomic            MED            : None
AIGP Metric    : None                  IGP Cost       : n/a
Connector      : None
Community      : target:64500:10
Cluster        : No Cluster Members
Originator Id  : None                  Peer Router Id : 192.0.2.6
Origin         : IGP
AS-Path        : No As-Path
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                     Dest Class     : 0
SRv6 TLV Type  : SRv6 L3 Service TLV (5)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:aaaa:14::
Full Sid       : 2001:db8:aaaa:14:7fff:d000::
Behavior       : End.DT4 (19)
SRv6 SubSubTLV : SRv6 SID Structure (1)
Loc-Block-Len  : 48                    Loc-Node-Len   : 16
Func-Len       : 20                    Arg-Len        : 0
Tpose-Len      : 20                    Tpose-offset   : 64

Network        : 10.10.1.0/24
Nexthop        : 2001:db8::14
Route Dist.    : 192.0.2.20:10         VPN Label      : 524285
Path Id        : None
To             : 2001:db8::6
Res. Nexthop   : n/a
Local Pref.    : 150                   Interface Name : NotAvailable
Aggregator AS  : None                  Aggregator     : None
Atomic Aggr.   : Not Atomic            MED            : None
AIGP Metric    : None                  IGP Cost       : n/a
Connector      : None
Community      : target:64500:10
Cluster        : No Cluster Members
Originator Id  : None                  Peer Router Id : 192.0.2.6
Origin         : IGP
AS-Path        : No As-Path
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                     Dest Class     : 0
SRv6 TLV Type  : SRv6 L3 Service TLV (5)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:aaaa:14::
Full Sid       : 2001:db8:aaaa:14:7fff:d000::
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Behavior       : End.DT4 (19)
SRv6 SubSubTLV : SRv6 SID Structure (1)
Loc-Block-Len  : 48                    Loc-Node-Len   : 16
Func-Len       : 20                    Arg-Len        : 0
Tpose-Len      : 20                    Tpose-offset   : 64

-------------------------------------------------------------------------------
Routes : 3
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show router bgp routes vpn-ipv6 hunt
===============================================================================
 BGP Router ID:192.0.2.20       AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv6 Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : 2001:db8::6600/120
Nexthop        : 2001:db8::6
Route Dist.    : 192.0.2.6:10          VPN Label      : 524287
Path Id        : None
From           : 2001:db8::6
Res. Nexthop   : n/a
Local Pref.    : 100                   Interface Name : int-1-bng-1-pe-2
Aggregator AS  : None                  Aggregator     : None
Atomic Aggr.   : Not Atomic            MED            : None
AIGP Metric    : None                  IGP Cost       : 20
Connector      : None
Community      : target:64500:10
Cluster        : No Cluster Members
Originator Id  : None                  Peer Router Id : 192.0.2.6
Fwd Class      : None                  Priority       : None
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                     Dest Class     : 0
Add Paths Send : Default
Last Modified  : 01d23h12m
SRv6 TLV Type  : SRv6 L3 Service TLV (5)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:aaaa:6::
Full Sid       : 2001:db8:aaaa:6:7fff:f000::
Behavior       : End.DT6 (18)
SRv6 SubSubTLV : SRv6 SID Structure (1)
Loc-Block-Len  : 48                    Loc-Node-Len   : 16
Func-Len       : 20                    Arg-Len        : 0
Tpose-Len      : 20                    Tpose-offset   : 64
VPRN Imported  :  10
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
Network        : 2001:db8:bbbb::/56
Nexthop        : 2001:db8::14
Route Dist.    : 192.0.2.20:10          VPN Label      : 524284
Path Id        : None
To             : 2001:db8::6
Res. Nexthop   : n/a
Local Pref.    : 150                    Interface Name : NotAvailable
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : n/a
Connector      : None
Community      : target:64500:10
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.6
Origin         : IGP
AS-Path        : No As-Path
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
SRv6 TLV Type  : SRv6 L3 Service TLV (5)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:aaaa:14::
Full Sid       : 2001:db8:aaaa:14:7fff:c000::
Behavior       : End.DT6 (18)
SRv6 SubSubTLV : SRv6 SID Structure (1)
Loc-Block-Len  : 48                     Loc-Node-Len   : 16
Func-Len       : 20                     Arg-Len        : 0
Tpose-Len      : 20                     Tpose-offset   : 64

Network        : 2001:db8:bbbb:100::/56
Nexthop        : 2001:db8::14
Route Dist.    : 192.0.2.20:10          VPN Label      : 524284
Path Id        : None
To             : 2001:db8::6
Res. Nexthop   : n/a
Local Pref.    : 150                    Interface Name : NotAvailable
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : n/a
Connector      : None
Community      : target:64500:10
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.6
Origin         : IGP
AS-Path        : No As-Path
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
SRv6 TLV Type  : SRv6 L3 Service TLV (5)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:aaaa:14::
Full Sid       : 2001:db8:aaaa:14:7fff:c000::
Behavior       : End.DT6 (18)
SRv6 SubSubTLV : SRv6 SID Structure (1)
Loc-Block-Len  : 48                     Loc-Node-Len   : 16
Func-Len       : 20                     Arg-Len        : 0
Tpose-Len      : 20                     Tpose-offset   : 64

-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Routes : 3
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show router segment-routing-v6 base-routing-instance all

===============================================================================
Segment Routing v6 Base Routing Instance
===============================================================================
Locator
  Type       Function    SID                                     Status/InstId
    SRH-mode Protection  Interface
-------------------------------------------------------------------------------
bng-1-loc
  End                  1 2001:db8:aaaa:14:0:1000::               ok
    USP
-------------------------------------------------------------------------------
  Auto-allocated End.X: USP Protected,
-------------------------------------------------------------------------------
  End.X          *524287 2001:db8:aaaa:14:7fff:f000::            1
    USP        Protected int-1-bng-1-bng-2
    ISIS Level: L2 Mac Address: b4:98:01:01:00:06 Nbr Sys Id: 1920.0000.2021
  End.X          *524288 2001:db8:aaaa:14:8000::                 0
    USP        Protected int-1-bng-1-p-1
    ISIS Level: L2 Mac Address: b4:9e:01:01:00:0b Nbr Sys Id: 1920.0000.2004
  End.X          *524289 2001:db8:aaaa:14:8000:1000::            1
    USP        Protected int-1-bng-1-pe-2
    ISIS Level: L2 Mac Address: b4:a0:01:01:00:07 Nbr Sys Id: 1920.0000.2006
-------------------------------------------------------------------------------
===============================================================================
Legend: * - System allocated

===============================================================================
Micro Segment Routing v6 Base Routing Instance
===============================================================================
Micro Segment Locator
  Type       Function    SID                                     Status/InstId
    SRH-mode Oper Func   Interface                               Protection
-------------------------------------------------------------------------------
===============================================================================
Legend: * - System allocated

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show router segment-routing-v6 local-sid context  "Base"

===============================================================================
Segment Routing v6 Local SIDs
===============================================================================
SID                                                 Type           Function
  Locator
  Context
-------------------------------------------------------------------------------
2001:db8:aaaa:14:0:1000::                           End            1
  bng-1-loc
  Base
-------------------------------------------------------------------------------
SIDs : 1
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show router segment-routing-v6 local-sid context  "10"

===============================================================================
Segment Routing v6 Local SIDs
===============================================================================
SID                                                 Type           Function
  Locator
  Context
-------------------------------------------------------------------------------
2001:db8:aaaa:14:7fff:c000::                        End.DT6        524284
  bng-1-loc
  SvcId: 10 Name: dual-homing
2001:db8:aaaa:14:7fff:d000::                        End.DT4        524285
  bng-1-loc
  SvcId: 10 Name: dual-homing
-------------------------------------------------------------------------------
SIDs : 2
-------------------------------------------------------------------------------
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show router route-table ipv6

===============================================================================
IPv6 Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
2001:db8::4/128                                 Remote  ISIS      01d23h12m  18
       fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                       10
2001:db8::5/128                                 Remote  ISIS      01d23h12m  18
       fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                       20
2001:db8::6/128 [L]                             Remote  ISIS(1)   01d23h12m  18
       fe80::b6a0:ffff:fe00:0-"int-1-bng-1-pe-2"                      10
2001:db8::a/128                                 Remote  ISIS      01d23h12m  18
       fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                       20
2001:db8::14/128                                Local   Local     01d23h13m  0
       system                                                         0
2001:db8::15/128 [L]                            Remote  ISIS(1)   01d23h12m  18
       fe80::b697:ffff:fe00:0-"int-1-bng-1-bng-2"                     10
2001:db8::100/120                               Remote  ISIS      01d23h12m  18
       fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                       20
2001:db8::200/120                               Remote  ISIS      01d23h12m  18
       fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                       30
2001:db8::500/120                               Local   Local     01d23h13m  0
       int-1-bng-1-p-1                                                0
2001:db8::600/120                               Remote  ISIS      01d23h12m  18
       fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                       30
2001:db8::700/120                               Local   Local     01d23h13m  0
       int-1-bng-1-bng-2                                              0
2001:db8::800/120                               Local   Local     01d23h13m  0
       int-1-bng-1-pe-2                                               0
2001:db8::900/120 [L]                           Remote  ISIS(1)   01d23h12m  18
       fe80::b6a0:ffff:fe00:0-"int-1-bng-1-pe-2"                      20
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
2001:db8::a00/120                              Remote  ISIS      01d23h12m  18
       fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                    20
2001:db8::6400/120                             Remote  ISIS      01d23h12m  18
       fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                    30
2001:db8::6500/120 [L]                         Remote  ISIS(1)   01d23h12m  18
       fe80::b6a0:ffff:fe00:0-"int-1-bng-1-pe-2"                   20
2001:db8:aaaa:4::/64                           Remote  ISIS      01d23h12m  18
       2001:db8:aaaa:4::/64 (tunneled:SRV6-ISIS)                   20
2001:db8:aaaa:5::/64                           Remote  ISIS      01d23h12m  18
       2001:db8:aaaa:5::/64 (tunneled:SRV6-ISIS)                   30
2001:db8:aaaa:6::/64                           Remote  ISIS(1)   01d23h12m  18
       2001:db8:aaaa:6::/64 (tunneled:SRV6-ISIS)                   20
2001:db8:aaaa:a::/64                           Remote  ISIS      01d23h12m  18
       2001:db8:aaaa:a::/64 (tunneled:SRV6-ISIS)                   30
2001:db8:aaaa:14::/64                          Local   SRV6      01d23h13m  3
       fe80::201-"_tmnx_fpe_2.a"                                   0
2001:db8:aaaa:14:0:1000::/128                  Local   SRV6      01d23h13m  3
       Black Hole                                                  0
2001:db8:aaaa:14:0:2000::/128                  Local   SRV6-Pol* 01d23h13m  14
       2001:db8::5 (tunneled:SRV6-Policy:917506)                   1
2001:db8:aaaa:14:7fff:f000::/128               Local   ISIS(1)   01d23h12m  18
       2001:db8:aaaa:14:7fff:f000:: (tunneled:SRV6-ISIS)           10
2001:db8:aaaa:14:8000::/128                    Local   ISIS      01d23h12m  18
       2001:db8:aaaa:14:8000:: (tunneled:SRV6-ISIS)                10
2001:db8:aaaa:14:8000:1000::/128               Local   ISIS(1)   01d23h12m  18
       2001:db8:aaaa:14:8000:1000:: (tunneled:SRV6-ISIS)           10
2001:db8:aaaa:15::/64                          Remote  ISIS(1)   01d23h12m  18
       2001:db8:aaaa:15::/64 (tunneled:SRV6-ISIS)                  20
-------------------------------------------------------------------------------
No. of Routes: 27
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
* indicates that the corresponding row element may have been truncated.

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show router tunnel-table ipv6

===============================================================================
IPv6 Tunnel Table (Router: Base)
===============================================================================
Destination                                    Owner     Encap TunnelId  Pref
Nexthop                                        Color           Metric
-------------------------------------------------------------------------------
2001:db8::5/128                                srv6-pol  SRV6  917506    14
  fpe_1.a                                        10              0
2001:db8:aaaa:4::/64                           srv6-isis SRV6  524296    0
  fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                       20
2001:db8:aaaa:5::/64                           srv6-isis SRV6  524299    0
  fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                       30
2001:db8:aaaa:6::/64 [L]                       srv6-isis SRV6  524295    0
  fe80::b6a0:ffff:fe00:0-"int-1-bng-1-pe-2"                      20
2001:db8:aaaa:a::/64                           srv6-isis SRV6  524297    0
  fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                       30
2001:db8:aaaa:14:7fff:f000::/128 [L]           srv6-isis SRV6  524291    0
  fe80::b697:ffff:fe00:0-"int-1-bng-1-bng-2"                     10
2001:db8:aaaa:14:8000::/128                    srv6-isis SRV6  524292    0
  fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                       10
2001:db8:aaaa:14:8000:1000::/128 [L]           srv6-isis SRV6  524293    0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
  fe80::b6a0:ffff:fe00:0-"int-1-bng-1-pe-2"                        10
2001:db8:aaaa:15::/64 [L]                       srv6-isis SRV6  524294    0
  fe80::b697:ffff:fe00:0-"int-1-bng-1-bng-2"                       20
-------------------------------------------------------------------------------
Flags: B = BGP or MPLS backup hop available
       L = Loop-Free Alternate (LFA) hop available
       E = Inactive best-external BGP route
       k = RIB-API or Forwarding Policy backup hop
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show router fp-tunnel-table 1 ipv6

===============================================================================
IPv6 Tunnel Table Display

Legend:
label stack is ordered from bottom-most to top-most
B - FRR Backup
===============================================================================
Destination                                 Protocol         Tunnel-ID
  Lbl/SID
    NextHop                                                  Intf/Tunnel
  Lbl/SID (backup)
    NextHop   (backup)
-------------------------------------------------------------------------------
2001:db8:aaaa:4::/64                        SRV6              524296
  -
    fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                1/1/c1/1:1
2001:db8:aaaa:5::/64                        SRV6              524299
  -
    fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                1/1/c1/1:1
2001:db8:aaaa:6::/64                        SRV6              524295
  -
    fe80::b6a0:ffff:fe00:0-"int-1-bng-1-pe-2"               1/1/c1/7:1
  -
    fe80::b697:ffff:fe00:0-"int-1-bng-1-bng-2"(B)           1/1/c1/6:1
2001:db8:aaaa:a::/64                        SRV6              524297
  -
    fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                1/1/c1/1:1
2001:db8:aaaa:15::/64                       SRV6              524294
  -
    fe80::b697:ffff:fe00:0-"int-1-bng-1-bng-2"              1/1/c1/6:1
  -
    fe80::b6a0:ffff:fe00:0-"int-1-bng-1-pe-2"(B)            1/1/c1/7:1
2001:db8::5/128                             SRV6-Policy       -
  2001:db8:aaaa:a:0:1000::/2001:db8:aaaa:4:0:1000::
    0.140.1.1                                               pxc-1.b:1
2001:db8:aaaa:14:7fff:f000::/128            SRV6              524291
  -
    fe80::b697:ffff:fe00:0-"int-1-bng-1-bng-2"              1/1/c1/6:1
  2001:db8:aaaa:15:0:1000::
    fe80::b6a0:ffff:fe00:0-"int-1-bng-1-pe-2"(B)            1/1/c1/7:1
2001:db8:aaaa:14:8000::/128                 SRV6              524292
  -
    fe80::b69e:ffff:fe00:0-"int-1-bng-1-p-1"                1/1/c1/1:1
2001:db8:aaaa:14:8000:1000::/128            SRV6              524293
  -
    fe80::b6a0:ffff:fe00:0-"int-1-bng-1-pe-2"               1/1/c1/7:1
  2001:db8:aaaa:6:0:1000::
    fe80::b697:ffff:fe00:0-"int-1-bng-1-bng-2"(B)           1/1/c1/6:1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
Total Entries : 9
-------------------------------------------------------------------------------
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show router  bgp  summary
===============================================================================
 BGP Router ID:192.0.2.20       AS:64500       Local AS:64500
===============================================================================
BGP Admin State         : Up          BGP Oper State              : Up
Total Peer Groups       : 2           Total Peers                 : 3
Total VPN Peer Groups   : 0           Total VPN Peers             : 0
Current Internal Groups : 3           Max Internal Groups         : 3
Total BGP Paths         : 59          Total Path Memory           : 21280

Total IPv4 Remote Rts   : 0           Total IPv4 Rem. Active Rts  : 0
Total IPv6 Remote Rts   : 0           Total IPv6 Rem. Active Rts  : 0
Total IPv4 Backup Rts   : 0           Total IPv6 Backup Rts       : 0
Total LblIpv4 Rem Rts   : 0           Total LblIpv4 Rem. Act Rts  : 0
Total LblIpv6 Rem Rts   : 0           Total LblIpv6 Rem. Act Rts  : 0
Total LblIpv4 Bkp Rts   : 0           Total LblIpv6 Bkp Rts       : 0
Total Supressed Rts     : 0           Total Hist. Rts             : 0
Total Decay Rts         : 0

Total VPN-IPv4 Rem. Rts : 1           Total VPN-IPv4 Rem. Act. Rts: 1
Total VPN-IPv6 Rem. Rts : 1           Total VPN-IPv6 Rem. Act. Rts: 1
Total VPN-IPv4 Bkup Rts : 0           Total VPN-IPv6 Bkup Rts     : 0
Total VPN Local Rts     : 7           Total VPN Supp. Rts         : 0
Total VPN Hist. Rts     : 0           Total VPN Decay Rts         : 0

Total MVPN-IPv4 Rem Rts : 0           Total MVPN-IPv4 Rem Act Rts : 0
Total MVPN-IPv6 Rem Rts : 0           Total MVPN-IPv6 Rem Act Rts : 0
Total MDT-SAFI Rem Rts  : 0           Total MDT-SAFI Rem Act Rts  : 0
Total McIPv4 Remote Rts : 0           Total McIPv4 Rem. Active Rts: 0
Total McIPv6 Remote Rts : 0           Total McIPv6 Rem. Active Rts: 0
Total McVpnIPv4 Rem Rts : 0           Total McVpnIPv4 Rem Act Rts : 0
Total McVpnIPv6 Rem Rts : 0           Total McVpnIPv6 Rem Act Rts : 0

Total EVPN Rem Rts      : 4           Total EVPN Rem Act Rts      : 4
Total L2-VPN Rem. Rts   : 0           Total L2VPN Rem. Act. Rts   : 0
Total MSPW Rem Rts      : 0           Total MSPW Rem Act Rts      : 0
Total RouteTgt Rem Rts  : 0           Total RouteTgt Rem Act Rts  : 0
Total FlowIpv4 Rem Rts  : 0           Total FlowIpv4 Rem Act Rts  : 0
Total FlowIpv6 Rem Rts  : 0           Total FlowIpv6 Rem Act Rts  : 0
Total FlowVpnv4 Rem Rts : 0           Total FlowVpnv4 Rem Act Rts : 0
Total FlowVpnv6 Rem Rts : 0           Total FlowVpnv6 Rem Act Rts : 0
Total Link State Rem Rts: 0           Total Link State Rem Act Rts: 0
Total SrPlcyIpv4 Rem Rts: 0           Total SrPlcyIpv4 Rem Act Rts: 0
Total SrPlcyIpv6 Rem Rts: 0           Total SrPlcyIpv6 Rem Act Rts: 0

===============================================================================
BGP Summary
===============================================================================
Legend : D - Dynamic Neighbor
===============================================================================
Neighbor
Description
                   AS PktRcvd InQ  Up/Down   State|Rcv/Act/Sent (Addr Family)
                      PktSent OutQ
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
2001:db8::5
                64500    5669    0 01d23h12m 1/1/3 (Evpn)
                         5672    0
2001:db8::6
                64500    5670    0 01d23h12m 1/1/2 (VpnIPv4)
                         5672    0           1/1/2 (VpnIPv6)
2001:db8::15
                64500    5673    0 01d23h12m 3/3/3 (Evpn)
                         5672    0
-------------------------------------------------------------------------------

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show router isis  adjacency

===============================================================================
Rtr Base ISIS Instance 0 Adjacency
===============================================================================
System ID               Usage State Hold Interface                    MT-ID
-------------------------------------------------------------------------------
p-1                     L2    Up    7    int-1-bng-1-p-1              0
-------------------------------------------------------------------------------
Adjacencies : 1
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show router isis database

===============================================================================
Rtr Base ISIS Instance 0 Database
===============================================================================
LSP ID                                Sequence  Checksum Lifetime Attributes
-------------------------------------------------------------------------------

Displaying Level 1 database
-------------------------------------------------------------------------------
Level (1) LSP Count : 0

Displaying Level 2 database
-------------------------------------------------------------------------------
p-1.00-00                             0x13d     0x8ffb   762      L1L2
p-1.01-00                             0x134     0x4d03   902      L1L2
pe-1.00-00                            0x139     0x659e   863      L1L2
pe-1.01-00                            0x133     0xe585   1063     L1L2
p-2.00-00                             0x137     0xfd42   1085     L1L2
p-2.01-00                             0x135     0x6dc8   938      L1L2
p-2.02-00                             0x135     0xe56c   1042     L1L2
p-2.03-00                             0x132     0xfa58   764      L1L2
bng-1.00-00                           0x13d     0x160f   748      L1L2
bng-2.00-00                           0x137     0x777b   675      L1L2
Level (2) LSP Count : 10
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show router isis 1 adjacency
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Rtr Base ISIS Instance 1 Adjacency
===============================================================================
System ID               Usage State Hold Interface                    MT-ID
-------------------------------------------------------------------------------
bng-2                   L2    Up    27   int-1-bng-1-bng-2           0
pe-2                    L2    Up    7    int-1-bng-1-pe-2            0
-------------------------------------------------------------------------------
Adjacencies : 2
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show router isis 1 database

===============================================================================
Rtr Base ISIS Instance 1 Database
===============================================================================
LSP ID                                Sequence  Checksum Lifetime Attributes
-------------------------------------------------------------------------------

Displaying Level 1 database
-------------------------------------------------------------------------------
Level (1) LSP Count : 0

Displaying Level 2 database
-------------------------------------------------------------------------------
pe-2.00-00                            0x13a     0x105f   963      L1L2
pe-2.01-00                            0x133     0x51fb   708      L1L2
pe-2.02-00                            0x133     0x60ea   1055     L1L2
bng-1.00-00                           0x13c     0x29a7   747      L1L2
bng-1.01-00                           0x136     0x7b99   670      L1L2
bng-2.00-00                           0x138     0x8827   1129     L1L2
Level (2) LSP Count : 6
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show redundancy multi-chassis sync peer 192.0.2.21 detail

===============================================================================
Multi-chassis Peer Table
===============================================================================
Peer
-------------------------------------------------------------------------------
Peer IP Address         : 192.0.2.21
Description             : (Not Specified)
Authentication          : Disabled
Source IP Address       : 192.0.2.20
Admin State             : Enabled
Warm standby            : No
Remote warm standby     : No
Sub-mgmt options        :
  DHCP lease threshold  : Inactive
    Local / Remote      : -- / --
-------------------------------------------------------------------------------
Sync-status
-------------------------------------------------------------------------------
Client Applications     : SUBMGMT-IPOE SUBMGMT-PPPOE SRRP DHCPServer
Sync Admin State        : Up
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Sync Oper State         : Up
Sync Oper Flags         :
DB Sync State           : inSync
Num Entries             : 100
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
Rem Num Entries         : 100
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0

===============================================================================
MCS Application Stats
===============================================================================
Application             : igmp
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : igmpSnooping
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : subMgmtIpoe
Num Entries             : 10
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 10
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : srrp
Num Entries             : 26
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 26
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : mcRing
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : mldSnooping
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : dhcpServer
Num Entries             : 54
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 54
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : subHostTrk
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : subMgmtPppoe
Num Entries             : 10
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Rem Num Entries         : 10
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : ipsec
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : mld
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : python
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : l2tp
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : diameterProxy
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : pimSnpgSap
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : pimSnpgSdp
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : diameterNode
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : nat
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
Ports synced on peer 192.0.2.21
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Port/Encap                    Tag
-------------------------------------------------------------------------------
pw-1                          pw-port-1
===============================================================================

===============================================================================
SDPs synced on peer 192.0.2.21
===============================================================================
SDP/Vc-Id                     Tag
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
DHCP Server instances synced on peer 192.0.2.21
===============================================================================
Router-Name                      Server-Name
  Tag
-------------------------------------------------------------------------------
No instances found
===============================================================================

===============================================================================
Python cache instances synced on peer 192.0.2.21
===============================================================================
Python-Policy                 Tag
-------------------------------------------------------------------------------
No instances found
===============================================================================
No L2TP instances found.

===============================================================================
Track SRRP instances
===============================================================================
SRRP                    : 1
-------------------------------------------------------------------------------
L2TP tunnel ID start    : 0
L2TP tunnel ID end      : 0

===============================================================================

===============================================================================
Diameter proxy instances synced on peer 192.0.2.21
===============================================================================
Diameter-Peer-Policy          Tag
-------------------------------------------------------------------------------
No instances found
===============================================================================

===============================================================================
Diameter node instances synced on peer 192.0.2.21
===============================================================================
Diameter Node                                 Tag
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
No. of Diameter Nodes: 0
===============================================================================

===============================================================================
Nat groups synced on peer 192.0.2.21
===============================================================================
Nat group                                     Tag
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
No. of Nat groups: 0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /tools dump redundancy multi-chassis sync-database

The following totals are for:
 peer ip ALL, port/lag/sdp ALL, sync-tag ALL, application ALL
Valid Entries:                  100
Locally Deleted Entries:        0
Locally Deleted Alarmed Entries: 0
Pending Global Delete Entries:  0
Omcr Alarmed Entries:           0
Omcr Standby Entries:           0
Associated Shared Records (ALL): 0
Associated Shared Records (LD):  0

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  service  active-subscribers summary

===============================================================================
Active Subscriber table summary
===============================================================================
Total Count       : 20
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  service  active-subscribers hierarchy

===============================================================================
Active Subscribers Hierarchy
===============================================================================
-- ipoe-ds-1
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.11] - sla:sla-profile-1
       |
       +-- IPOE-session - mac:00:13:01:00:00:01 - svc:10
           |
           |-- 10.10.0.12 - DHCP
           |
           +-- 2001:db8:bbbb::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:1::/64 - DHCP6-PD-MR

-- ipoe-ds-10
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.20] - sla:sla-profile-1
       |
       +-- IPOE-session - mac:00:13:01:00:00:0a - svc:10
           |
           |-- 10.10.0.13 - DHCP
           |
           +-- 2001:db8:bbbb:10::1/128 - DHCP6
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
               |
               +-- 2001:db8:bbbb:11::/64 - DHCP6-PD-MR

-- ipoe-ds-2
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.12] - sla:sla-profile-1
       |
       +-- IPOE-session - mac:00:13:01:00:00:02 - svc:10
           |
           |-- 10.10.0.10 - DHCP
           |
           +-- 2001:db8:bbbb:2::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:3::/64 - DHCP6-PD-MR

-- ipoe-ds-3
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.13] - sla:sla-profile-1
       |
       +-- IPOE-session - mac:00:13:01:00:00:03 - svc:10
           |
           |-- 10.10.0.14 - DHCP
           |
           +-- 2001:db8:bbbb:c::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:d::/64 - DHCP6-PD-MR

-- ipoe-ds-4
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.14] - sla:sla-profile-1
       |
       +-- IPOE-session - mac:00:13:01:00:00:04 - svc:10
           |
           |-- 10.10.0.17 - DHCP
           |
           +-- 2001:db8:bbbb:8::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:9::/64 - DHCP6-PD-MR

-- ipoe-ds-5
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.15] - sla:sla-profile-1
       |
       +-- IPOE-session - mac:00:13:01:00:00:05 - svc:10
           |
           |-- 10.10.0.11 - DHCP
           |
           +-- 2001:db8:bbbb:4::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:5::/64 - DHCP6-PD-MR

-- ipoe-ds-6
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.16] - sla:sla-profile-1
       |
       +-- IPOE-session - mac:00:13:01:00:00:06 - svc:10
           |
           |-- 10.10.0.19 - DHCP
           |
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            +-- 2001:db8:bbbb:e::1/128 - DHCP6
                |
                +-- 2001:db8:bbbb:f::/64 - DHCP6-PD-MR

-- ipoe-ds-7
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.17] - sla:sla-profile-1
       |
       +-- IPOE-session - mac:00:13:01:00:00:07 - svc:10
           |
           |-- 10.10.0.18 - DHCP
           |
           +-- 2001:db8:bbbb:a::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:b::/64 - DHCP6-PD-MR

-- ipoe-ds-8
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.18] - sla:sla-profile-1
       |
       +-- IPOE-session - mac:00:13:01:00:00:08 - svc:10
           |
           |-- 10.10.0.16 - DHCP
           |
           +-- 2001:db8:bbbb:6::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:7::/64 - DHCP6-PD-MR

-- ipoe-ds-9
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.19] - sla:sla-profile-1
       |
       +-- IPOE-session - mac:00:13:01:00:00:09 - svc:10
           |
           |-- 10.10.0.15 - DHCP
           |
           +-- 2001:db8:bbbb:12::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:13::/64 - DHCP6-PD-MR

-- pppoe-1
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.1] - sla:sla-profile-1
       |
       +-- PPP-session - mac:00:11:01:00:00:01 - sid:1 - svc:10
           |
           |-- 10.10.1.10 - IPCP
           |
           +-- 2001:db8:bbbb:14::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:15::/64 - DHCP6-PD-MR

-- pppoe-10
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.10] - sla:sla-profile-1
       |
       +-- PPP-session - mac:00:11:01:00:00:0a - sid:1 - svc:10
           |
           |-- 10.10.1.19 - IPCP
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
           |
           +-- 2001:db8:bbbb:26::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:27::/64 - DHCP6-PD-MR

-- pppoe-2
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.2] - sla:sla-profile-1
       |
       +-- PPP-session - mac:00:11:01:00:00:02 - sid:1 - svc:10
           |
           |-- 10.10.1.11 - IPCP
           |
           +-- 2001:db8:bbbb:18::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:19::/64 - DHCP6-PD-MR

-- pppoe-3
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.3] - sla:sla-profile-1
       |
       +-- PPP-session - mac:00:11:01:00:00:03 - sid:1 - svc:10
           |
           |-- 10.10.1.12 - IPCP
           |
           +-- 2001:db8:bbbb:16::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:17::/64 - DHCP6-PD-MR

-- pppoe-4
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.4] - sla:sla-profile-1
       |
       +-- PPP-session - mac:00:11:01:00:00:04 - sid:1 - svc:10
           |
           |-- 10.10.1.13 - IPCP
           |
           +-- 2001:db8:bbbb:1a::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:1b::/64 - DHCP6-PD-MR

-- pppoe-5
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.5] - sla:sla-profile-1
       |
       +-- PPP-session - mac:00:11:01:00:00:05 - sid:1 - svc:10
           |
           |-- 10.10.1.14 - IPCP
           |
           +-- 2001:db8:bbbb:1c::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:1d::/64 - DHCP6-PD-MR

-- pppoe-6
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.6] - sla:sla-profile-1
       |
       +-- PPP-session - mac:00:11:01:00:00:06 - sid:1 - svc:10
           |
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
           |-- 10.10.1.15 - IPCP
           |
           +-- 2001:db8:bbbb:1e::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:1f::/64 - DHCP6-PD-MR

-- pppoe-7
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.7] - sla:sla-profile-1
       |
       +-- PPP-session - mac:00:11:01:00:00:07 - sid:1 - svc:10
           |
           |-- 10.10.1.16 - IPCP
           |
           +-- 2001:db8:bbbb:20::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:21::/64 - DHCP6-PD-MR

-- pppoe-8
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.8] - sla:sla-profile-1
       |
       +-- PPP-session - mac:00:11:01:00:00:08 - sid:1 - svc:10
           |
           |-- 10.10.1.17 - IPCP
           |
           +-- 2001:db8:bbbb:22::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:23::/64 - DHCP6-PD-MR

-- pppoe-9
   (sub-profile-1)
   |
   +-- sap:[pw-1:2.9] - sla:sla-profile-1
       |
       +-- PPP-session - mac:00:11:01:00:00:09 - sid:1 - svc:10
           |
           |-- 10.10.1.18 - IPCP
           |
           +-- 2001:db8:bbbb:24::1/128 - DHCP6
               |
               +-- 2001:db8:bbbb:25::/64 - DHCP6-PD-MR

-------------------------------------------------------------------------------
Number of active subscribers : 20
Flags: (N) = the host or the managed route is in non-forwarding state
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  service id 10 subscriber-hosts detail

===============================================================================
Subscriber Host table
===============================================================================
Sap
  IP Address
    MAC Address               PPPoE-SID       Origin          Fwding State
      Subscriber
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
[pw-1:2.1]
  10.10.1.10
    00:11:01:00:00:01          1               IPCP            Fwding
      pppoe-1
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000029636FE951
Acct-Q-Inst-Session-Id: B496FF0000002A636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.1]
  2001:db8:bbbb:14::1/128
    00:11:01:00:00:01          1               PPP-DHCP6       Fwding
      pppoe-1
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000047636FE951
Acct-Q-Inst-Session-Id: B496FF0000002A636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.2]
  10.10.1.11
    00:11:01:00:00:02          1               IPCP            Fwding
      pppoe-2
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF0000002C636FE951
Acct-Q-Inst-Session-Id: B496FF0000002D636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.2]
  2001:db8:bbbb:18::1/128
    00:11:01:00:00:02          1               PPP-DHCP6       Fwding
      pppoe-2
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000049636FE951
Acct-Q-Inst-Session-Id: B496FF0000002D636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.3]
  10.10.1.12
    00:11:01:00:00:03          1               IPCP            Fwding
      pppoe-3
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF0000002F636FE951
Acct-Q-Inst-Session-Id: B496FF00000030636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.3]
  2001:db8:bbbb:16::1/128
    00:11:01:00:00:03          1               PPP-DHCP6       Fwding
      pppoe-3
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000046636FE951
Acct-Q-Inst-Session-Id: B496FF00000030636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.4]
  10.10.1.13
    00:11:01:00:00:04          1               IPCP            Fwding
      pppoe-4
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000032636FE951
Acct-Q-Inst-Session-Id: B496FF00000033636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.4]
  2001:db8:bbbb:1a::1/128
    00:11:01:00:00:04          1               PPP-DHCP6       Fwding
      pppoe-4
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000048636FE951
Acct-Q-Inst-Session-Id: B496FF00000033636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.5]
  10.10.1.14
    00:11:01:00:00:05          1               IPCP            Fwding
      pppoe-5
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000035636FE951
Acct-Q-Inst-Session-Id: B496FF00000036636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.5]
  2001:db8:bbbb:1c::1/128
    00:11:01:00:00:05          1               PPP-DHCP6       Fwding
      pppoe-5
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF0000004A636FE951
Acct-Q-Inst-Session-Id: B496FF00000036636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.6]
  10.10.1.15
    00:11:01:00:00:06          1               IPCP            Fwding
      pppoe-6
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000038636FE951
Acct-Q-Inst-Session-Id: B496FF00000039636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.6]
  2001:db8:bbbb:1e::1/128
    00:11:01:00:00:06          1               PPP-DHCP6       Fwding
      pppoe-6
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF0000004B636FE951
Acct-Q-Inst-Session-Id: B496FF00000039636FE951
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.7]
  10.10.1.16
    00:11:01:00:00:07          1               IPCP            Fwding
      pppoe-7
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF0000003B636FE951
Acct-Q-Inst-Session-Id: B496FF0000003C636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.7]
  2001:db8:bbbb:20::1/128
    00:11:01:00:00:07          1               PPP-DHCP6       Fwding
      pppoe-7
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF0000004C636FE951
Acct-Q-Inst-Session-Id: B496FF0000003C636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.8]
  10.10.1.17
    00:11:01:00:00:08          1               IPCP            Fwding
      pppoe-8
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Acct-Session-Id       : B496FF0000003E636FE951
Acct-Q-Inst-Session-Id: B496FF0000003F636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.8]
  2001:db8:bbbb:22::1/128
    00:11:01:00:00:08          1               PPP-DHCP6       Fwding
      pppoe-8
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF0000004D636FE951
Acct-Q-Inst-Session-Id: B496FF0000003F636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.9]
  10.10.1.18
    00:11:01:00:00:09          1               IPCP            Fwding
      pppoe-9
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000041636FE951
Acct-Q-Inst-Session-Id: B496FF00000042636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.9]
  2001:db8:bbbb:24::1/128
    00:11:01:00:00:09          1               PPP-DHCP6       Fwding
      pppoe-9
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF0000004E636FE951
Acct-Q-Inst-Session-Id: B496FF00000042636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.10]
  10.10.1.19
    00:11:01:00:00:0a          1               IPCP            Fwding
      pppoe-10
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000044636FE951
Acct-Q-Inst-Session-Id: B496FF00000045636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.10]
  2001:db8:bbbb:26::1/128
    00:11:01:00:00:0a          1               PPP-DHCP6       Fwding
      pppoe-10
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF0000004F636FE951
Acct-Q-Inst-Session-Id: B496FF00000045636FE951
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.11]
  10.10.0.12
    00:13:01:00:00:01          N/A             DHCP            Fwding
      ipoe-ds-1
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000013636FE94C
Acct-Q-Inst-Session-Id: B496FF00000010636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.11]
  2001:db8:bbbb::1/128
    00:13:01:00:00:01          N/A             IPoE-DHCP6      Fwding
      ipoe-ds-1
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF0000000F636FE94C
Acct-Q-Inst-Session-Id: B496FF00000010636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.12]
  10.10.0.10
    00:13:01:00:00:02          N/A             DHCP            Fwding
      ipoe-ds-2
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF0000000A636FE94C
Acct-Q-Inst-Session-Id: B496FF00000005636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.12]
  2001:db8:bbbb:2::1/128
    00:13:01:00:00:02          N/A             IPoE-DHCP6      Fwding
      ipoe-ds-2
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000004636FE94C
Acct-Q-Inst-Session-Id: B496FF00000005636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.13]
  10.10.0.14
    00:13:01:00:00:03          N/A             DHCP            Fwding
      ipoe-ds-3
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF0000001B636FE94C
Acct-Q-Inst-Session-Id: B496FF0000001C636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.13]
  2001:db8:bbbb:c::1/128
    00:13:01:00:00:03          N/A             IPoE-DHCP6      Fwding
      ipoe-ds-3
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF0000001F636FE94C
Acct-Q-Inst-Session-Id: B496FF0000001C636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.14]
  10.10.0.17
    00:13:01:00:00:04          N/A             DHCP            Fwding
      ipoe-ds-4
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000025636FE951
Acct-Q-Inst-Session-Id: B496FF00000012636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.14]
  2001:db8:bbbb:8::1/128
    00:13:01:00:00:04          N/A             IPoE-DHCP6      Fwding
      ipoe-ds-4
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000011636FE94C
Acct-Q-Inst-Session-Id: B496FF00000012636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.15]
  10.10.0.11
    00:13:01:00:00:05          N/A             DHCP            Fwding
      ipoe-ds-5
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF0000000B636FE94C
Acct-Q-Inst-Session-Id: B496FF00000007636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.15]
  2001:db8:bbbb:4::1/128
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    00:13:01:00:00:05          N/A             IPoE-DHCP6      Fwding
      ipoe-ds-5
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000006636FE94C
Acct-Q-Inst-Session-Id: B496FF00000007636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.16]
  10.10.0.19
    00:13:01:00:00:06          N/A             DHCP            Fwding
      ipoe-ds-6
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000027636FE951
Acct-Q-Inst-Session-Id: B496FF0000001E636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.16]
  2001:db8:bbbb:e::1/128
    00:13:01:00:00:06          N/A             IPoE-DHCP6      Fwding
      ipoe-ds-6
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF0000001D636FE94C
Acct-Q-Inst-Session-Id: B496FF0000001E636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
[pw-1:2.17]
  10.10.0.18
    00:13:01:00:00:07          N/A             DHCP            Fwding
      ipoe-ds-7
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000026636FE951
Acct-Q-Inst-Session-Id: B496FF00000015636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.17]
  2001:db8:bbbb:a::1/128
    00:13:01:00:00:07          N/A             IPoE-DHCP6      Fwding
      ipoe-ds-7
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000014636FE94C
Acct-Q-Inst-Session-Id: B496FF00000015636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.18]
  10.10.0.16
    00:13:01:00:00:08          N/A             DHCP            Fwding
      ipoe-ds-8
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000024636FE950
Acct-Q-Inst-Session-Id: B496FF00000009636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.18]
  2001:db8:bbbb:6::1/128
    00:13:01:00:00:08          N/A             IPoE-DHCP6      Fwding
      ipoe-ds-8
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000008636FE94C
Acct-Q-Inst-Session-Id: B496FF00000009636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.19]
  10.10.0.15
    00:13:01:00:00:09          N/A             DHCP            Fwding
      ipoe-ds-9
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000020636FE94C
Acct-Q-Inst-Session-Id: B496FF00000021636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.19]
  2001:db8:bbbb:12::1/128
    00:13:01:00:00:09          N/A             IPoE-DHCP6      Fwding
      ipoe-ds-9
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000023636FE94E
Acct-Q-Inst-Session-Id: B496FF00000021636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.20]
  10.10.0.13
    00:13:01:00:00:0a          N/A             DHCP            Fwding
      ipoe-ds-10
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000016636FE94C
Acct-Q-Inst-Session-Id: B496FF00000017636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:2.20]
  2001:db8:bbbb:10::1/128
    00:13:01:00:00:0a          N/A             IPoE-DHCP6      Fwding
      ipoe-ds-10
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : B496FF00000022636FE94E
Acct-Q-Inst-Session-Id: B496FF00000017636FE94C
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
Number of subscriber hosts : 40
===============================================================================

[/]
A:admin@bng-1# /show  service id 10 ipoe session

===============================================================================
IPoE sessions for svc-id 10
===============================================================================
Sap Id                          Mac Address         Up Time         MC-Stdby
    Subscriber-Id
        [CircuitID] | [RemoteID]
-------------------------------------------------------------------------------
[pw-1:2.11]                     00:13:01:00:00:01   2d 02:51:27
    ipoe-ds-1
[pw-1:2.12]                     00:13:01:00:00:02   2d 02:51:27
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    ipoe-ds-2
[pw-1:2.13]                     00:13:01:00:00:03   2d 02:51:26
    ipoe-ds-3
[pw-1:2.14]                     00:13:01:00:00:04   2d 02:51:26
    ipoe-ds-4
[pw-1:2.15]                     00:13:01:00:00:05   2d 02:51:27
    ipoe-ds-5
[pw-1:2.16]                     00:13:01:00:00:06   2d 02:51:26
    ipoe-ds-6
[pw-1:2.17]                     00:13:01:00:00:07   2d 02:51:26
    ipoe-ds-7
[pw-1:2.18]                     00:13:01:00:00:08   2d 02:51:26
    ipoe-ds-8
[pw-1:2.19]                     00:13:01:00:00:09   2d 02:51:26
    ipoe-ds-9
[pw-1:2.20]                     00:13:01:00:00:0a   2d 02:51:26
    ipoe-ds-10
-------------------------------------------------------------------------------
CID | RID displayed when included in session-key
Number of sessions : 10
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  service id 10 pppoe session

===============================================================================
PPPoE sessions for svc-id 10
===============================================================================
Sap Id              Mac Address       Sid    Up Time         Type
    IP/L2TP-Id/Interface-Id                                   MC-Stdby
-------------------------------------------------------------------------------
[pw-1:2.1]          00:11:01:00:00:01 1      2d 02:51:21     local
    10.10.1.10
    00:01:00:02:00:01:00:01
[pw-1:2.2]          00:11:01:00:00:02 1      2d 02:51:21     local
    10.10.1.11
    00:02:00:02:00:02:00:01
[pw-1:2.3]          00:11:01:00:00:03 1      2d 02:51:21     local
    10.10.1.12
    00:03:00:02:00:03:00:01
[pw-1:2.4]          00:11:01:00:00:04 1      2d 02:51:21     local
    10.10.1.13
    00:04:00:02:00:04:00:01
[pw-1:2.5]          00:11:01:00:00:05 1      2d 02:51:21     local
    10.10.1.14
    00:05:00:02:00:05:00:01
[pw-1:2.6]          00:11:01:00:00:06 1      2d 02:51:21     local
    10.10.1.15
    00:06:00:02:00:06:00:01
[pw-1:2.7]          00:11:01:00:00:07 1      2d 02:51:21     local
    10.10.1.16
    00:07:00:02:00:07:00:01
[pw-1:2.8]          00:11:01:00:00:08 1      2d 02:51:21     local
    10.10.1.17
    00:08:00:02:00:08:00:01
[pw-1:2.9]          00:11:01:00:00:09 1      2d 02:51:21     local
    10.10.1.18
    00:09:00:02:00:09:00:01
[pw-1:2.10]         00:11:01:00:00:0a 1      2d 02:51:21     local
    10.10.1.19
    00:0A:00:02:00:0A:00:01
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
Number of sessions   : 10
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  service id 10 dhcp lease-state

===============================================================================
DHCP lease state table, service 10
===============================================================================
IP Address      Mac Address       Sap/Sdp Id          Remaining  Lease    MC
                                                      LeaseTime  Origin   Stdby
-------------------------------------------------------------------------------
10.10.0.10      00:13:01:00:00:02 [pw-1:2.12]         00h13m42s  DHCP
10.10.0.11      00:13:01:00:00:05 [pw-1:2.15]         00h13m42s  DHCP
10.10.0.12      00:13:01:00:00:01 [pw-1:2.11]         00h13m42s  DHCP
10.10.0.13      00:13:01:00:00:0a [pw-1:2.20]         00h13m42s  DHCP
10.10.0.14      00:13:01:00:00:03 [pw-1:2.13]         00h13m42s  DHCP
10.10.0.15      00:13:01:00:00:09 [pw-1:2.19]         00h13m42s  DHCP
10.10.0.16      00:13:01:00:00:08 [pw-1:2.18]         00h13m43s  DHCP
10.10.0.17      00:13:01:00:00:04 [pw-1:2.14]         00h13m44s  DHCP
10.10.0.18      00:13:01:00:00:07 [pw-1:2.17]         00h13m44s  DHCP
10.10.0.19      00:13:01:00:00:06 [pw-1:2.16]         00h13m44s  DHCP
-------------------------------------------------------------------------------
Number of lease states : 10
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  service id 10 dhcp6 lease-state

===============================================================================
DHCP lease state table, service 10
===============================================================================
IP Address      Mac Address       Sap/Sdp Id          Remaining  Lease    MC
                                                      LeaseTime  Origin   Stdby
-------------------------------------------------------------------------------
2001:db8:bbbb::1/128
                00:13:01:00:00:01 [pw-1:2.11]         23h43m37s  DHCP
2001:db8:bbbb:1::/64
                00:13:01:00:00:01 [pw-1:2.11]         23h43m37s  DHCP
2001:db8:bbbb:2::1/128
                00:13:01:00:00:02 [pw-1:2.12]         23h43m36s  DHCP
2001:db8:bbbb:3::/64
                00:13:01:00:00:02 [pw-1:2.12]         23h43m36s  DHCP
2001:db8:bbbb:4::1/128
                00:13:01:00:00:05 [pw-1:2.15]         23h43m36s  DHCP
2001:db8:bbbb:5::/64
                00:13:01:00:00:05 [pw-1:2.15]         23h43m36s  DHCP
2001:db8:bbbb:6::1/128
                00:13:01:00:00:08 [pw-1:2.18]         23h43m36s  DHCP
2001:db8:bbbb:7::/64
                00:13:01:00:00:08 [pw-1:2.18]         23h43m36s  DHCP
2001:db8:bbbb:8::1/128
                00:13:01:00:00:04 [pw-1:2.14]         23h43m37s  DHCP
2001:db8:bbbb:9::/64
                00:13:01:00:00:04 [pw-1:2.14]         23h43m37s  DHCP
2001:db8:bbbb:a::1/128
                00:13:01:00:00:07 [pw-1:2.17]         23h43m37s  DHCP
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
2001:db8:bbbb:b::/64
                00:13:01:00:00:07 [pw-1:2.17]         23h43m37s  DHCP
2001:db8:bbbb:c::1/128
                00:13:01:00:00:03 [pw-1:2.13]         23h43m37s  DHCP
2001:db8:bbbb:d::/64
                00:13:01:00:00:03 [pw-1:2.13]         23h43m37s  DHCP
2001:db8:bbbb:e::1/128
                00:13:01:00:00:06 [pw-1:2.16]         23h43m37s  DHCP
2001:db8:bbbb:f::/64
                00:13:01:00:00:06 [pw-1:2.16]         23h43m37s  DHCP
2001:db8:bbbb:10::1/128
                00:13:01:00:00:0a [pw-1:2.20]         23h38m40s  DHCP
2001:db8:bbbb:11::/64
                00:13:01:00:00:0a [pw-1:2.20]         23h38m40s  DHCP
2001:db8:bbbb:12::1/128
                00:13:01:00:00:09 [pw-1:2.19]         23h38m40s  DHCP
2001:db8:bbbb:13::/64
                00:13:01:00:00:09 [pw-1:2.19]         23h38m40s  DHCP
2001:db8:bbbb:14::1/128
                00:11:01:00:00:01 [pw-1:2.1]          23h43m41s  DHCP
2001:db8:bbbb:15::/64
                00:11:01:00:00:01 [pw-1:2.1]          23h43m41s  DHCP
2001:db8:bbbb:16::1/128
                00:11:01:00:00:03 [pw-1:2.3]          23h43m41s  DHCP
2001:db8:bbbb:17::/64
                00:11:01:00:00:03 [pw-1:2.3]          23h43m41s  DHCP
2001:db8:bbbb:18::1/128
                00:11:01:00:00:02 [pw-1:2.2]          23h43m41s  DHCP
2001:db8:bbbb:19::/64
                00:11:01:00:00:02 [pw-1:2.2]          23h43m41s  DHCP
2001:db8:bbbb:1a::1/128
                00:11:01:00:00:04 [pw-1:2.4]          23h43m41s  DHCP
2001:db8:bbbb:1b::/64
                00:11:01:00:00:04 [pw-1:2.4]          23h43m41s  DHCP
2001:db8:bbbb:1c::1/128
                00:11:01:00:00:05 [pw-1:2.5]          23h43m42s  DHCP
2001:db8:bbbb:1d::/64
                00:11:01:00:00:05 [pw-1:2.5]          23h43m42s  DHCP
2001:db8:bbbb:1e::1/128
                00:11:01:00:00:06 [pw-1:2.6]          23h43m41s  DHCP
2001:db8:bbbb:1f::/64
                00:11:01:00:00:06 [pw-1:2.6]          23h43m41s  DHCP
2001:db8:bbbb:20::1/128
                00:11:01:00:00:07 [pw-1:2.7]          23h43m41s  DHCP
2001:db8:bbbb:21::/64
                00:11:01:00:00:07 [pw-1:2.7]          23h43m41s  DHCP
2001:db8:bbbb:22::1/128
                00:11:01:00:00:08 [pw-1:2.8]          23h43m42s  DHCP
2001:db8:bbbb:23::/64
                00:11:01:00:00:08 [pw-1:2.8]          23h43m42s  DHCP
2001:db8:bbbb:24::1/128
                00:11:01:00:00:09 [pw-1:2.9]          23h43m41s  DHCP
2001:db8:bbbb:25::/64
                00:11:01:00:00:09 [pw-1:2.9]          23h43m41s  DHCP
2001:db8:bbbb:26::1/128
                00:11:01:00:00:0a [pw-1:2.10]         23h43m41s  DHCP
2001:db8:bbbb:27::/64
                00:11:01:00:00:0a [pw-1:2.10]         23h43m41s  DHCP
-------------------------------------------------------------------------------
Number of lease states : 40
===============================================================================

[/]
A:admin@bng-1#
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
[/]
A:admin@bng-1# /show  srrp 1 detail

===============================================================================
SRRP Instance 1
===============================================================================
Description         : (Not Specified)
Admin State         : Up                Oper State       : master
Preempt             : yes               One GARP per SAP : no
Monitor Oper Group  : None
System IP           : 192.0.2.20
Service ID          : VPRN 10
Group If            : group-int-1       MAC Address      : b4:96:ff:00:00:00
Grp If Description  : N/A
Grp If Admin State  : Up                Grp If Oper State: Up
Subscriber If       : sub-int-1
Sub If Admin State  : Up                Sub If Oper State: Up
Address             : 10.10.0.1/24      Gateway IP       : 10.10.0.254
Address             : 10.10.1.1/24      Gateway IP       : 10.10.1.254
Redundant If        : red-int-bng-1-bng*
Red If Admin State  : Up                Red If Oper State: Up
Address             : 192.168.11.1/24
Red Spoke-sdp       : 1:1
Msg Path SAP        : pw-1:2.4094       Passive          : no
Admin Gateway MAC   :                   Oper Gateway MAC : 00:00:5e:00:01:01
Standby-Forwarding  : Disabled
Config Priority     : 100               In-use Priority  : 100
Master Priority     : 100
Keep-alive Interval : 2 deci-seconds    Master Since     : 11/12/2022 16:23:34
Fib Population Mode : all
VRRP Policy 1       : None              VRRP Policy 2    : None

-------------------------------------------------------------------------------
Statistics
-------------------------------------------------------------------------------
Become Master       : 1                 Master Changes   : 1
Become Bkup Routing : 0                 Become Bkup Shunt: 1
Become Non-Master   : 0
Adv Sent            : 849384            Adv Received     : 0
Pri 0 Pkts Sent     : 0                 Pri 0 Pkts Rcvd  : 0
Preempt Events      : 0                 Preempted Events : 0
Mesg Intvl Discards : 0                 Mesg Intvl Errors: 0

===============================================================================
* indicates that the corresponding row element may have been truncated.

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  router 10 dhcp local-dhcp-server "dhcpv4" summary
===============================================================================
DHCP server dhcpv4  router 10
===============================================================================
Admin State            : inService
Operational State      : inService
Persistency State      : shutdown
User Data Base         : N/A
Use gateway IP address : enabled (scope pool)
Use pool from client   : enabled
Send force-renewals    : disabled
Creation Origin        : manual
Lease Hold Time        : 0h0m0s
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Lease Hold Time For    : N/A
User-ident             : mac-circuit-id

Failover Admin State   : outOfService
Failover Oper State    : shutdown
Failover Persist Key   : N/A
Administrative MCLT    : 0h10m0s
Operational MCLT       : 0h10m0s
Startup wait time      : 0h2m0s
Partner down delay     : 23h59m59s
  Ignore MCLT          : disabled

-------------------------------------------------------------------------------
Pool name : dhcpv4-1
-------------------------------------------------------------------------------
Failover Admin State   : inService
Failover Oper State    : normal
Failover Persist Key   : N/A
Administrative MCLT    : 0h10m0s
Operational MCLT       : 0h10m0s
Startup wait time      : 0h2m0s
Partner down delay     : 23h59m59s
  Ignore MCLT          : disabled
-------------------------------------------------------------------------------
Subnet                 Free     %    Stable   Declined Offered  Rem-pend Drain
-------------------------------------------------------------------------------
10.10.0.0/24       (A) 81       89%  10       0        0        0        N

Totals for pool        81       89%  10       0        0        0
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
Pool name : pppoev4-1
-------------------------------------------------------------------------------
Failover Admin State   : outOfService
Failover Oper State    : shutdown
Failover Persist Key   : N/A
Administrative MCLT    : 0h10m0s
Operational MCLT       : 0h10m0s
Startup wait time      : 0h2m0s
Partner down delay     : 23h59m59s
  Ignore MCLT          : disabled
-------------------------------------------------------------------------------
Subnet                 Free     %    Stable   Declined Offered  Rem-pend Drain
-------------------------------------------------------------------------------
10.10.1.0/24           81       89%  10       0        0        0        N

Totals for pool        81       89%  10       0        0        0
-------------------------------------------------------------------------------

Totals for server      162      89%  20       0        0        0

-------------------------------------------------------------------------------
Interface associations
Interface                       Admin
-------------------------------------------------------------------------------
loopback-1                      Up

-------------------------------------------------------------------------------
Local Address Assignment associations
Group interface                 Admin
-------------------------------------------------------------------------------
group-int-1                     Up
No associated firewall domains found.
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  router 10 dhcp6 local-dhcp-server "dhcpv6" summary
===============================================================================
DHCP server dhcpv6  router 10
===============================================================================
Admin State            : inService
Operational State      : inService
Persistency State      : shutdown
Use Link Address       : disabled
Use pool from client   : enabled
Creation Origin        : manual
Lease Hold Time        : 0h0m0s
Lease Hold Time For    : N/A
User-ident             : duid
Interface-id-mapping   : disabled
Ignore-rapid-commit    : disabled
Allow-lease-query      : disabled
Auto-provisioned       : false

Failover Admin State   : outOfService
Failover Oper State    : shutdown
Failover Persist Key   : N/A
Administrative MCLT    : 0h10m0s
Operational MCLT       : 0h10m0s
Startup wait time      : 0h2m0s
Partner down delay     : 23h59m59s
  Ignore MCLT          : disabled

-------------------------------------------------------------------------------
Pool name : dhcpv6-1
-------------------------------------------------------------------------------
Failover Admin State   : inService
Failover Oper State    : normal
Failover Persist Key   : N/A
Administrative MCLT    : 0h10m0s
Operational MCLT       : 0h10m0s
Startup wait time      : 0h2m0s
Partner down delay     : 23h59m59s
  Ignore MCLT          : disabled
-------------------------------------------------------------------------------
Prefix
                                    Stable   Declined Advert   Rem-pend Drain
-------------------------------------------------------------------------------
2001:db8:bbbb::/56
                               (A)  40       0        0        0        N
2001:db8:bbbb:100::/56
                               (A)  0        0        0        0        N

Totals for pool                     40       0        0        0
-------------------------------------------------------------------------------

Totals for server                   40       0        0        0

-------------------------------------------------------------------------------
Interface associations
Interface                       Admin
-------------------------------------------------------------------------------
loopback-1                      Up
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
Local Address Assignment associations
Group interface                 Admin
-------------------------------------------------------------------------------
No associated firewall domains found.
===============================================================================

[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  router 10 dhcp local-dhcp-server "dhcpv4" leases

===============================================================================
Leases for DHCP server dhcpv4 router 10
===============================================================================
IP Address      Lease State       Mac Address       Remaining   Clnt  Fail
  PPP user name/Opt82 Circuit Id                    LifeTime    Type  Ctrl
  User-db/Sticky-lease Hostname
-------------------------------------------------------------------------------
10.10.0.10      stable            00:13:01:00:00:02 0h23m42s    dhcp  local
  cid-2
10.10.0.11      stable            00:13:01:00:00:05 0h23m42s    dhcp  local
  cid-5
10.10.0.12      stable            00:13:01:00:00:01 0h23m42s    dhcp  local
  cid-1
10.10.0.13      stable            00:13:01:00:00:0a 0h23m42s    dhcp  local
  cid-10
10.10.0.14      stable            00:13:01:00:00:03 0h23m42s    dhcp  local
  cid-3
10.10.0.15      stable            00:13:01:00:00:09 0h23m42s    dhcp  local
  cid-9
10.10.0.16      stable            00:13:01:00:00:08 0h23m43s    dhcp  local
  cid-8
10.10.0.17      stable            00:13:01:00:00:04 0h23m44s    dhcp  local
  cid-4
10.10.0.18      stable            00:13:01:00:00:07 0h23m44s    dhcp  local
  cid-7
10.10.0.19      stable            00:13:01:00:00:06 0h23m44s    dhcp  local
  cid-6
10.10.1.10      internal          N/A               N/A         ppp   N/A

10.10.1.11      internal          N/A               N/A         ppp   N/A

10.10.1.12      internal          N/A               N/A         ppp   N/A

10.10.1.13      internal          N/A               N/A         ppp   N/A

10.10.1.14      internal          N/A               N/A         ppp   N/A

10.10.1.15      internal          N/A               N/A         ppp   N/A

10.10.1.16      internal          N/A               N/A         ppp   N/A

10.10.1.17      internal          N/A               N/A         ppp   N/A

10.10.1.18      internal          N/A               N/A         ppp   N/A

10.10.1.19      internal          N/A               N/A         ppp   N/A

-------------------------------------------------------------------------------
20 leases found
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
[/]
A:admin@bng-1#

[/]
A:admin@bng-1# /show  router 10 dhcp6 local-dhcp-server "dhcpv6" leases

===============================================================================
Leases for DHCPv6 server dhcpv6
===============================================================================
IP Address/Prefix                          Lease State       Remaining   Fail
  Link-local Address                                         LifeTime    Ctrl
-------------------------------------------------------------------------------
2001:db8:bbbb::1/64
  fe80::213:1ff:fe00:1                     stable            1d0h13m     local
2001:db8:bbbb:1::/64
  fe80::213:1ff:fe00:1                     stable            1d0h13m     local
2001:db8:bbbb:2::1/64
  fe80::213:1ff:fe00:2                     stable            1d0h13m     local
2001:db8:bbbb:3::/64
  fe80::213:1ff:fe00:2                     stable            1d0h13m     local
2001:db8:bbbb:4::1/64
  fe80::213:1ff:fe00:5                     stable            1d0h13m     local
2001:db8:bbbb:5::/64
  fe80::213:1ff:fe00:5                     stable            1d0h13m     local
2001:db8:bbbb:6::1/64
  fe80::213:1ff:fe00:8                     stable            1d0h13m     local
2001:db8:bbbb:7::/64
  fe80::213:1ff:fe00:8                     stable            1d0h13m     local
2001:db8:bbbb:8::1/64
  fe80::213:1ff:fe00:4                     stable            1d0h13m     local
2001:db8:bbbb:9::/64
  fe80::213:1ff:fe00:4                     stable            1d0h13m     local
2001:db8:bbbb:a::1/64
  fe80::213:1ff:fe00:7                     stable            1d0h13m     local
2001:db8:bbbb:b::/64
  fe80::213:1ff:fe00:7                     stable            1d0h13m     local
2001:db8:bbbb:c::1/64
  fe80::213:1ff:fe00:3                     stable            1d0h13m     local
2001:db8:bbbb:d::/64
  fe80::213:1ff:fe00:3                     stable            1d0h13m     local
2001:db8:bbbb:e::1/64
  fe80::213:1ff:fe00:6                     stable            1d0h13m     local
2001:db8:bbbb:f::/64
  fe80::213:1ff:fe00:6                     stable            1d0h13m     local
2001:db8:bbbb:10::1/64
  fe80::213:1ff:fe00:a                     stable            1d0h8m      local
2001:db8:bbbb:11::/64
  fe80::213:1ff:fe00:a                     stable            1d0h8m      local
2001:db8:bbbb:12::1/64
  fe80::213:1ff:fe00:9                     stable            1d0h8m      local
2001:db8:bbbb:13::/64
  fe80::213:1ff:fe00:9                     stable            1d0h8m      local
2001:db8:bbbb:14::1/64
  fe80::1:2:1:1                            stable            1d0h13m     local
2001:db8:bbbb:15::/64
  fe80::1:2:1:1                            stable            1d0h13m     local
2001:db8:bbbb:16::1/64
  fe80::3:2:3:1                            stable            1d0h13m     local
2001:db8:bbbb:17::/64
  fe80::3:2:3:1                            stable            1d0h13m     local
2001:db8:bbbb:18::1/64
  fe80::2:2:2:1                            stable            1d0h13m     local
2001:db8:bbbb:19::/64
  fe80::2:2:2:1                            stable            1d0h13m     local
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
2001:db8:bbbb:1a::1/64
  fe80::4:2:4:1                           stable           1d0h13m     local
2001:db8:bbbb:1b::/64
  fe80::4:2:4:1                           stable           1d0h13m     local
2001:db8:bbbb:1c::1/64
  fe80::5:2:5:1                           stable           1d0h13m     local
2001:db8:bbbb:1d::/64
  fe80::5:2:5:1                           stable           1d0h13m     local
2001:db8:bbbb:1e::1/64
  fe80::6:2:6:1                           stable           1d0h13m     local
2001:db8:bbbb:1f::/64
  fe80::6:2:6:1                           stable           1d0h13m     local
2001:db8:bbbb:20::1/64
  fe80::7:2:7:1                           stable           1d0h13m     local
2001:db8:bbbb:21::/64
  fe80::7:2:7:1                           stable           1d0h13m     local
2001:db8:bbbb:22::1/64
  fe80::8:2:8:1                           stable           1d0h13m     local
2001:db8:bbbb:23::/64
  fe80::8:2:8:1                           stable           1d0h13m     local
2001:db8:bbbb:24::1/64
  fe80::9:2:9:1                           stable           1d0h13m     local
2001:db8:bbbb:25::/64
  fe80::9:2:9:1                           stable           1d0h13m     local
2001:db8:bbbb:26::1/64
  fe80::a:2:a:1                           stable           1d0h13m     local
2001:db8:bbbb:27::/64
  fe80::a:2:a:1                           stable           1d0h13m     local
-------------------------------------------------------------------------------
40 leases found
===============================================================================
```

Show commands on P-1

```
A:admin@p-1# /show port

===============================================================================
Ports on Slot 1
===============================================================================
Port          Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id            State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
1/1/c1        Up         Link Up                        conn   100GBASE-LR4*
1/1/c1/1      Up    Yes  Up      8936 8936    - hybr dotq xgige
1/1/c1/2      Up    Yes  Up      8936 8936    - hybr dotq xgige
1/1/c1/3      Up    Yes  Up      8936 8936    - hybr dotq xgige
1/1/c1/4      Up    Yes  Up      8936 8936    - hybr dotq xgige
1/1/c1/5      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c1/6      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c1/7      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c1/8      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c1/9      Up    Yes  Up      8936 8936    - hybr dotq xgige
1/1/c1/10     Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c2        Up         Link Up                        conn   100G CWDM4 M*
1/1/c2/1      Up    Yes  Up      8936 8936    - hybr dotq xgige
1/1/c2/2      Down  No   Down    8936 8936    - hybr dotq xgige
1/1/c2/3      Down  No   Down    8936 8936    - hybr dotq xgige
1/1/c2/4      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c2/5      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c2/6      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c2/7      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c2/8      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c2/9      Up    No   Down    8936 8936    - hybr dotq xgige
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
1/1/c2/10     Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c3        Up         Link Up                         conn   100G CLR4   *
1/1/c3/1      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c3/2      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c3/3      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c3/4      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c3/5      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c3/6      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c3/7      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c3/8      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c3/9      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c3/10     Up    Yes  Up      8936 8936    - hybr dotq xgige
1/1/c4        Up         Link Up                         conn   100GBASE-LR4*
1/1/c4/1      Up    Yes  Up      8936 8936    - hybr dotq xgige
1/1/c4/2      Down  No   Down    8936 8936    - netw null xgige
1/1/c4/3      Down  No   Down    8936 8936    - netw null xgige
1/1/c4/4      Down  No   Down    8936 8936    - netw null xgige
1/1/c5        Down       Down                            conn   100G CWDM4 M*
1/1/c6        Down       Down                            conn   100G CLR4   *
1/2/c1        Down       Down                            conn   100GBASE-LR4*
1/2/c2        Down       Down                            conn   100G CWDM4 M*
1/2/c3        Down       Down                            conn   100G CLR4   *
1/2/c4        Down       Down                            conn   100GBASE-LR4*
1/2/c5        Down       Down                            conn   100G CWDM4 M*
1/2/c6        Down       Down                            conn   100G CLR4   *

===============================================================================
Ports on Slot A
===============================================================================
Port          Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id            State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
A/1           Up    Yes  Up      1514 1514    - netw null faste  MDI
A/3           Down  No   Down    1514 1514    - netw null faste
A/4           Down  No   Down    1514 1514    - netw null faste

===============================================================================
Ports on Port Cross Connect 1
===============================================================================
Port          Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id            State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
pxc-1.a       Up    Yes  Up      8932 8932    - hybr dotq xgige
pxc-1.b       Up    Yes  Up      8932 8932    - hybr dotq xgige

===============================================================================
Ports on Port Cross Connect 2
===============================================================================
Port          Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id            State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
pxc-2.a       Up    Yes  Up      8932 8932    - hybr dotq xgige
pxc-2.b       Up    Yes  Up      8932 8932    - hybr dotq xgige
===============================================================================

[/]
A:admin@p-1#

[/]
A:admin@p-1# /show router interface

===============================================================================
Interface Table (Router: Base)
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Interface-Name                   Adm       Opr(v4/v6)  Mode    Port/SapId
   IP-Address                                                   PfxState
-------------------------------------------------------------------------------
_tmnx_fpe_1.a                    Up        Up/Up       Network pxc-1.a:1
   fe80::100/64                                                 PREFERRED
_tmnx_fpe_1.b                    Up        Up/Up       Network pxc-1.b:1
   fe80::101/64                                                 PREFERRED
_tmnx_fpe_2.a                    Up        Up/Up       Network pxc-2.a:1
   fe80::200/64                                                 PREFERRED
_tmnx_fpe_2.b                    Up        Up/Up       Network pxc-2.b:1
   fe80::201/64                                                 PREFERRED
int-1-p-1-bng-1                  Up        Down/Up     Network 1/1/c2/1:1
   2001:db8::502/120                                            PREFERRED
   fe80::b69e:ffff:fe00:0/64                                    PREFERRED
int-1-p-1-p-2                    Up        Down/Up     Network 1/1/c1/9:1
   2001:db8::a01/120                                            PREFERRED
   fe80::b69e:ffff:fe00:0/64                                    PREFERRED
int-1-p-1-pe-1                   Up        Down/Up     Network 1/1/c1/1:1
   2001:db8::102/120                                            PREFERRED
   fe80::b69e:ffff:fe00:0/64                                    PREFERRED
system                           Up        Up/Up       Network system
   192.0.2.4/32                                                 n/a
   2001:db8::4/128                                              PREFERRED
to-radius                        Up        Up/Down     Network 1/1/c3/10:114
   192.168.114.4/24                                             n/a
-------------------------------------------------------------------------------
Interfaces : 9
===============================================================================

[/]
A:admin@p-1#

[/]
A:admin@p-1# /show  fwd-path-ext fpe 1

===============================================================================
FPE Id: 1
===============================================================================
Description        : (Not Specified)
Multi-Path         : Disabled
Path               : pxc 1
Pw Port Extension  : Disabled                          Oper     : down
Sub Mgmt Extension : Disabled                          Oper     : N/A
Vxlan              : Disabled                          Oper     : down
Segment-Routing V6 : Enabled                           Oper     : up
SRv6 Type          : origination
If-A Qos Policy    : default
If-B MTU           : 9786 bytes                        Oper MTU : 8914 bytes
If-B Qos Policy    : default
===============================================================================

[/]
A:admin@p-1#

[/]
A:admin@p-1# /show  fwd-path-ext fpe 2

===============================================================================
FPE Id: 2
===============================================================================
Description        : (Not Specified)
Multi-Path         : Disabled
Path               : pxc 2
Pw Port Extension  : Disabled                          Oper     : down
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Sub Mgmt Extension : Disabled                          Oper     : N/A
Vxlan              : Disabled                          Oper     : down
Segment-Routing V6 : Enabled                           Oper     : up
SRv6 Type          : termination
If-A Qos Policy    : default
If-B MTU           : 0 bytes                           Oper MTU : 8914 bytes
If-B Qos Policy    : default
===============================================================================

[/]
A:admin@p-1#

[/]
A:admin@p-1# /show router segment-routing-v6 base-routing-instance all

===============================================================================
Segment Routing v6 Base Routing Instance
===============================================================================
Locator
  Type       Function    SID                                     Status/InstId
    SRH-mode Protection  Interface
-------------------------------------------------------------------------------
p-1-loc
  End                 1 2001:db8:aaaa:4:0:1000::                ok
    USP
-------------------------------------------------------------------------------
  Auto-allocated End.X: USP Protected,
-------------------------------------------------------------------------------
  End.X              *2 2001:db8:aaaa:4:0:2000::                0
    USP        Protected int-1-p-1-pe-1
    ISIS Level: L2 Mac Address: b4:9f:01:01:00:01 Nbr Sys Id: 1920.0000.2005
  End.X              *4 2001:db8:aaaa:4:0:4000::                0
    USP        Protected int-1-p-1-p-2
    ISIS Level: L2 Mac Address: b4:a1:01:01:00:09 Nbr Sys Id: 1920.0000.2010
  End.X              *5 2001:db8:aaaa:4:0:5000::                0
    USP        Protected int-1-p-1-bng-1
    ISIS Level: L2 Mac Address: b4:99:01:01:00:01 Nbr Sys Id: 1920.0000.2020
-------------------------------------------------------------------------------
===============================================================================
Legend: * - System allocated

===============================================================================
Micro Segment Routing v6 Base Routing Instance
===============================================================================
Micro Segment Locator
  Type       Function    SID                                     Status/InstId
    SRH-mode Oper Func   Interface                               Protection
-------------------------------------------------------------------------------
===============================================================================
Legend: * - System allocated

[/]
A:admin@p-1#

[/]
A:admin@p-1# /show router segment-routing-v6 local-sid context  "Base"

===============================================================================
Segment Routing v6 Local SIDs
===============================================================================
SID                                              Type           Function
  Locator
  Context
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
2001:db8:aaaa:4:0:1000::                          End            1
  p-1-loc
  Base
-------------------------------------------------------------------------------
SIDs : 1
-------------------------------------------------------------------------------
===============================================================================

[/]
A:admin@p-1#

[/]
A:admin@p-1# /show router route-table ipv6

===============================================================================
IPv6 Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
2001:db8::4/128                               Local   Local     02d03h09m  0
       system                                                         0
2001:db8::5/128 [L]                           Remote  ISIS      02d03h08m  18
       fe80::b69f:ffff:fe00:0-"int-1-p-1-pe-1"                        10
2001:db8::a/128 [L]                           Remote  ISIS      02d03h08m  18
       fe80::b6a1:ffff:fe00:0-"int-1-p-1-p-2"                         10
2001:db8::14/128                              Remote  ISIS      01d23h20m  18
       fe80::b696:ffff:fe00:0-"int-1-p-1-bng-1"                       10
2001:db8::15/128 [L]                          Remote  ISIS      02d03h08m  18
       fe80::b6a1:ffff:fe00:0-"int-1-p-1-p-2"                         20
2001:db8::100/120                             Local   Local     02d03h08m  0
       int-1-p-1-pe-1                                                 0
2001:db8::200/120 [L]                         Remote  ISIS      02d03h08m  18
       fe80::b69f:ffff:fe00:0-"int-1-p-1-pe-1"                        20
2001:db8::500/120                             Local   Local     02d03h08m  0
       int-1-p-1-bng-1                                                0
2001:db8::600/120 [L]                         Remote  ISIS      02d03h08m  18
       fe80::b6a1:ffff:fe00:0-"int-1-p-1-p-2"                         20
2001:db8::a00/120                             Local   Local     02d03h08m  0
       int-1-p-1-p-2                                                  0
2001:db8::6400/120 [L]                        Remote  ISIS      02d03h08m  18
       fe80::b69f:ffff:fe00:0-"int-1-p-1-pe-1"                        20
2001:db8:aaaa:4::/64                          Local   SRV6      02d03h08m  3
       fe80::201-"_tmnx_fpe_2.a"                                      0
2001:db8:aaaa:4:0:1000::/128                  Local   SRV6      02d03h09m  3
       Black Hole                                                     0
2001:db8:aaaa:4:0:2000::/128                  Local   ISIS      02d03h08m  18
       2001:db8:aaaa:4:0:2000:: (tunneled:SRV6-ISIS)                  10
2001:db8:aaaa:4:0:4000::/128                  Local   ISIS      02d03h08m  18
       2001:db8:aaaa:4:0:4000:: (tunneled:SRV6-ISIS)                  10
2001:db8:aaaa:4:0:5000::/128                  Local   ISIS      01d23h20m  18
       2001:db8:aaaa:4:0:5000:: (tunneled:SRV6-ISIS)                  10
2001:db8:aaaa:5::/64                          Remote  ISIS      02d03h08m  18
       2001:db8:aaaa:5::/64 (tunneled:SRV6-ISIS)                      20
2001:db8:aaaa:a::/64                          Remote  ISIS      02d03h08m  18
       2001:db8:aaaa:a::/64 (tunneled:SRV6-ISIS)                      20
2001:db8:aaaa:14::/64                         Remote  ISIS      01d23h20m  18
       2001:db8:aaaa:14::/64 (tunneled:SRV6-ISIS)                     20
2001:db8:aaaa:15::/64                         Remote  ISIS      02d03h08m  18
       2001:db8:aaaa:15::/64 (tunneled:SRV6-ISIS)                     30
-------------------------------------------------------------------------------
No. of Routes: 20
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================

[/]
A:admin@p-1#

[/]
A:admin@p-1# /show router tunnel-table ipv6

===============================================================================
IPv6 Tunnel Table (Router: Base)
===============================================================================
Destination                                 Owner     Encap TunnelId  Pref
Nexthop                                     Color           Metric
-------------------------------------------------------------------------------
2001:db8:aaaa:4:0:2000::/128 [L]            srv6-isis SRV6  524289    0
  fe80::b69f:ffff:fe00:0-"int-1-p-1-pe-1"                     10
2001:db8:aaaa:4:0:4000::/128 [L]            srv6-isis SRV6  524293    0
  fe80::b6a1:ffff:fe00:0-"int-1-p-1-p-2"                      10
2001:db8:aaaa:4:0:5000::/128                srv6-isis SRV6  524300    0
  fe80::b696:ffff:fe00:0-"int-1-p-1-bng-1"                    10
2001:db8:aaaa:5::/64 [L]                    srv6-isis SRV6  524290    0
  fe80::b69f:ffff:fe00:0-"int-1-p-1-pe-1"                     20
2001:db8:aaaa:a::/64 [L]                    srv6-isis SRV6  524294    0
  fe80::b6a1:ffff:fe00:0-"int-1-p-1-p-2"                      20
2001:db8:aaaa:14::/64                       srv6-isis SRV6  524301    0
  fe80::b696:ffff:fe00:0-"int-1-p-1-bng-1"                    20
2001:db8:aaaa:15::/64 [L]                   srv6-isis SRV6  524295    0
  fe80::b6a1:ffff:fe00:0-"int-1-p-1-p-2"                      30
-------------------------------------------------------------------------------
Flags: B = BGP or MPLS backup hop available
       L = Loop-Free Alternate (LFA) hop available
       E = Inactive best-external BGP route
       k = RIB-API or Forwarding Policy backup hop
===============================================================================

[/]
A:admin@p-1#

[/]
A:admin@p-1# /show router fp-tunnel-table 1 ipv6

===============================================================================
IPv6 Tunnel Table Display

Legend:
label stack is ordered from bottom-most to top-most
B - FRR Backup
===============================================================================
Destination                                 Protocol         Tunnel-ID
  Lbl/SID
    NextHop                                                  Intf/Tunnel
  Lbl/SID (backup)
    NextHop   (backup)
-------------------------------------------------------------------------------
2001:db8:aaaa:5::/64                        SRV6              524290
  -
    fe80::b69f:ffff:fe00:0-"int-1-p-1-pe-1"                 1/1/c1/1:1
  -
    fe80::b6a1:ffff:fe00:0-"int-1-p-1-p-2"(B)               1/1/c1/9:1
2001:db8:aaaa:a::/64                        SRV6              524294
  -
    fe80::b6a1:ffff:fe00:0-"int-1-p-1-p-2"                  1/1/c1/9:1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
  -
    fe80::b69f:ffff:fe00:0-"int-1-p-1-pe-1"(B)              1/1/c1/1:1
2001:db8:aaaa:14::/64                       SRV6              524301
  -
    fe80::b696:ffff:fe00:0-"int-1-p-1-bng-1"                1/1/c2/1:1
2001:db8:aaaa:15::/64                       SRV6              524295
  -
    fe80::b6a1:ffff:fe00:0-"int-1-p-1-p-2"                  1/1/c1/9:1
  -
    fe80::b69f:ffff:fe00:0-"int-1-p-1-pe-1"(B)              1/1/c1/1:1
2001:db8:aaaa:4:0:2000::/128                SRV6              524289
  -
    fe80::b69f:ffff:fe00:0-"int-1-p-1-pe-1"                 1/1/c1/1:1
  2001:db8:aaaa:5:0:1000::
    fe80::b6a1:ffff:fe00:0-"int-1-p-1-p-2"(B)               1/1/c1/9:1
2001:db8:aaaa:4:0:4000::/128                SRV6              524293
  -
    fe80::b6a1:ffff:fe00:0-"int-1-p-1-p-2"                  1/1/c1/9:1
  2001:db8:aaaa:a:0:1000::
    fe80::b69f:ffff:fe00:0-"int-1-p-1-pe-1"(B)              1/1/c1/1:1
2001:db8:aaaa:4:0:5000::/128                SRV6              524300
  -
    fe80::b696:ffff:fe00:0-"int-1-p-1-bng-1"                1/1/c2/1:1
-------------------------------------------------------------------------------
Total Entries : 7
-------------------------------------------------------------------------------
===============================================================================

[/]
A:admin@p-1#

[/]
A:admin@p-1# /show router isis adjacency

===============================================================================
Rtr Base ISIS Instance 0 Adjacency
===============================================================================
System ID               Usage State Hold Interface                    MT-ID
-------------------------------------------------------------------------------
bng-1                   L2    Up    22   int-1-p-1-bng-1              0
p-2                     L2    Up    7    int-1-p-1-p-2                0
pe-1                    L2    Up    7    int-1-p-1-pe-1               0
-------------------------------------------------------------------------------
Adjacencies : 3
===============================================================================

[/]
A:admin@p-1#

[/]
A:admin@p-1# /show router isis database

===============================================================================
Rtr Base ISIS Instance 0 Database
===============================================================================
LSP ID                                Sequence  Checksum Lifetime Attributes
-------------------------------------------------------------------------------

Displaying Level 1 database
-------------------------------------------------------------------------------
Level (1) LSP Count : 0

Displaying Level 2 database
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
p-1.00-00                                0x13e     0x8dfc   952      L1L2
p-1.01-00                                0x135     0x4b04   1027     L1L2
pe-1.00-00                               0x13a     0x639f   1012     L1L2
pe-1.01-00                               0x133     0xe585   626      L1L2
p-2.00-00                                0x137     0xfd42   647      L1L2
p-2.01-00                                0x136     0x6bc9   1146     L1L2
p-2.02-00                                0x136     0xe36d   1149     L1L2
p-2.03-00                                0x133     0xf859   886      L1L2
bng-1.00-00                              0x13e     0x1410   947      L1L2
bng-2.00-00                              0x138     0x757c   837      L1L2
Level (2) LSP Count : 10
===============================================================================

[/]
```

Show commands on PE-1

```
A:admin@pe-1# /show port

===============================================================================
Ports on Slot 1
===============================================================================
Port          Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id            State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
1/1/c1        Up         Link Up                          conn   100GBASE-LR4*
1/1/c1/1      Up    Yes  Up      8936 8936    - hybr dotq xgige
1/1/c1/2      Up    Yes  Up      8936 8936    - hybr dotq xgige
1/1/c1/3      Up    Yes  Up      8936 8936    - hybr dotq xgige
1/1/c1/4      Up    Yes  Up      8936 8936    - hybr dotq xgige
1/1/c1/5      Up    Yes  Up      8936 8936    - hybr dotq xgige
1/1/c1/6      Up    Yes  Up      8936 8936    - hybr dotq xgige
1/1/c1/7      Up    Yes  Up      8936 8936    - hybr dotq xgige
1/1/c1/8      Up    Yes  Up      8936 8936    - hybr dotq xgige
1/1/c1/9      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c1/10     Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c2        Up         Link Up                          conn   100G CWDM4 M*
1/1/c2/1      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c2/2      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c2/3      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c2/4      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c2/5      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c2/6      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c2/7      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c2/8      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c2/9      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c2/10     Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c3        Up         Link Up                          conn   100G CLR4   *
1/1/c3/1      Up    Yes  Up      8936 8936    - hybr qinq xgige
1/1/c3/2      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c3/3      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c3/4      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c3/5      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c3/6      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c3/7      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c3/8      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c3/9      Up    No   Down    8936 8936    - hybr dotq xgige
1/1/c3/10     Up    Yes  Up      8936 8936    - hybr dotq xgige
1/1/c4        Up         Link Up                          conn   100GBASE-LR4*
1/1/c4/1      Up    Yes  Up      8936 8936    - hybr dotq xgige
1/1/c4/2      Down  No   Down    8936 8936    - netw null xgige
1/1/c4/3      Down  No   Down    8936 8936    - netw null xgige
1/1/c4/4      Down  No   Down    8936 8936    - netw null xgige
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
1/1/c5        Down       Down                              conn   100G CWDM4 M*
1/1/c6        Down       Down                              conn   100G CLR4   *
1/2/c1        Down       Down                              conn   100GBASE-LR4*
1/2/c2        Down       Down                              conn   100G CWDM4 M*
1/2/c3        Down       Down                              conn   100G CLR4   *
1/2/c4        Down       Down                              conn   100GBASE-LR4*
1/2/c5        Down       Down                              conn   100G CWDM4 M*
1/2/c6        Down       Down                              conn   100G CLR4   *
1/2/m1/1      Down       Link Up                           anchor

===============================================================================
Ports on Slot A
===============================================================================
Port          Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id            State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
A/1           Up    Yes  Up      1514 1514    - netw null faste  MDI
A/3           Down  No   Down    1514 1514    - netw null faste
A/4           Down  No   Down    1514 1514    - netw null faste

===============================================================================
Ports on Port Cross Connect 1
===============================================================================
Port          Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id            State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
pxc-1.a       Up    Yes  Up      8932 8932    - hybr dotq xgige
pxc-1.b       Up    Yes  Up      8932 8932    - hybr dotq xgige

===============================================================================
Ports on Port Cross Connect 2
===============================================================================
Port          Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id            State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
pxc-2.a       Up    Yes  Up      8932 8932    - hybr dotq xgige
pxc-2.b       Up    Yes  Up      8932 8932    - hybr dotq xgige
===============================================================================

[/]
A:admin@pe-1#

[/]
A:admin@pe-1# /show router interface

===============================================================================
Interface Table (Router: Base)
===============================================================================
Interface-Name                   Adm       Opr(v4/v6)  Mode    Port/SapId
   IP-Address                                                   PfxState
-------------------------------------------------------------------------------
_tmnx_fpe_1.a                    Up        Up/Up       Network pxc-1.a:1
   fe80::100/64                                                 PREFERRED
_tmnx_fpe_1.b                    Up        Up/Up       Network pxc-1.b:1
   fe80::101/64                                                 PREFERRED
_tmnx_fpe_2.a                    Up        Up/Up       Network pxc-2.a:1
   fe80::200/64                                                 PREFERRED
_tmnx_fpe_2.b                    Up        Up/Up       Network pxc-2.b:1
   fe80::201/64                                                 PREFERRED
int-1-pe-1-p-1                   Up        Down/Up     Network 1/1/c1/1:1
   2001:db8::101/120                                            PREFERRED
   fe80::b69f:ffff:fe00:0/64                                    PREFERRED
int-1-pe-1-p-2                   Up        Down/Up     Network 1/1/c1/5:1
   2001:db8::201/120                                            PREFERRED
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
   fe80::b69f:ffff:fe00:0/64                                    PREFERRED
system                          Up        Up/Up       Network system
   192.0.2.5/32                                                 n/a
   2001:db8::5/128                                              PREFERRED
to-ixia                         Up        Up/Up       Network 1/1/c3/1:1.1
   172.16.100.1/24                                              n/a
   2001:db8::6401/120                                           PREFERRED
   fe80::b69f:ffff:fe00:0/64                                    PREFERRED
to-radius                       Up        Up/Down     Network 1/1/c3/10:114
   192.168.114.5/24                                             n/a
-------------------------------------------------------------------------------
Interfaces : 9
===============================================================================

[/]
A:admin@pe-1#

[/]
A:admin@pe-1# /show  fwd-path-ext fpe 1

===============================================================================
FPE Id: 1
===============================================================================
Description        : (Not Specified)
Multi-Path         : Disabled
Path               : pxc 1
Pw Port Extension  : Disabled                          Oper     : down
Sub Mgmt Extension : Disabled                          Oper     : N/A
Vxlan              : Disabled                          Oper     : down
Segment-Routing V6 : Enabled                           Oper     : up
SRv6 Type          : origination
If-A Qos Policy    : default
If-B MTU           : 9786 bytes                        Oper MTU : 8914 bytes
If-B Qos Policy    : default
===============================================================================

[/]
A:admin@pe-1#

[/]
A:admin@pe-1# /show  fwd-path-ext fpe 2

===============================================================================
FPE Id: 2
===============================================================================
Description        : (Not Specified)
Multi-Path         : Disabled
Path               : pxc 2
Pw Port Extension  : Disabled                          Oper     : down
Sub Mgmt Extension : Disabled                          Oper     : N/A
Vxlan              : Disabled                          Oper     : down
Segment-Routing V6 : Enabled                           Oper     : up
SRv6 Type          : termination
If-A Qos Policy    : default
If-B MTU           : 0 bytes                           Oper MTU : 8914 bytes
If-B Qos Policy    : default
===============================================================================

[/]
A:admin@pe-1#

[/]
A:admin@pe-1# /show service id "dual-homing" base
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Service Basic Information
===============================================================================
Service Id        : 11                  Vpn Id            : 0
Service Type      : Epipe
MACSec enabled    : no
Name              : dual-homing
Description       : (Not Specified)
Customer Id       : 1                   Creation Origin   : manual
Last Status Change: 11/12/2022 15:58:53
Last Mgmt Change  : 11/12/2022 12:08:24
Test Service      : No
Admin State       : Up                  Oper State        : Up
MTU               : 1514
Vc Switching      : False
SAP Count         : 1                   SDP Bind Count    : 0
Per Svc Hashing   : Disabled            Lbl Eth/IP L4 TEID: Disabled
Ignore MTU Mismat*: Disabled
Vxlan Src Tep Ip  : N/A
Force QTag Fwd    : Disabled
Lcl Switch Svc St : sap
Oper Group        : <none>

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                               Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/c3/1:*.*                         qinq         8936    8936    Up   Up
===============================================================================
* indicates that the corresponding row element may have been truncated.

[/]
A:admin@pe-1#

[/]
A:admin@pe-1# /show service id "dual-homing" bgp-evpn

===============================================================================
BGP EVPN Table
===============================================================================
EVI               : 11                  Creation Origin    : manual

-------------------------------------------------------------------------------
Local AC Name             Eth Tag  Endpoint                    Ingress Label
-------------------------------------------------------------------------------
access                    1                                    0
-------------------------------------------------------------------------------
Number of local ACs : 1

-------------------------------------------------------------------------------
Remote AC Name                    Eth Tag  Endpoint
-------------------------------------------------------------------------------
bng                               2
-------------------------------------------------------------------------------
Number of Remote ACs : 1
===============================================================================

===============================================================================
Segment Routing v6 Instance 1 Service 11
===============================================================================
Admin State            : Enabled
Srv6 Instance          : 1
Default Locator        : pe-1-loc
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Oper Group              : (Not Specified)
Default Route Tag       : 0xb
Source Address          : 2001:db8::5
ECMP                    : 1
Force Vlan VC Fwd       : disabled
Next Hop Type           : system-ipv6
Evi 3-byte Auto-RT      : disabled
Route Resolution        : fallback-tunnel-to-route-table
Force QinQ VC Fwd       : none
MH Mode                 : network
===============================================================================

[/]
A:admin@pe-1#

[/]
A:admin@pe-1# /show router bgp routes evpn auto-disc tag 2 detail
===============================================================================
 BGP Router ID:192.0.2.5        AS:64500       Local AS:64500
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP EVPN Auto-Disc Routes
===============================================================================
Original Attributes

Network        : n/a
Nexthop        : 2001:db8::14
Path Id        : None
From           : 2001:db8::14
Res. Nexthop   : fe80::b69e:ffff:fe00:0
Local Pref.    : 100                    Interface Name : int-1-pe-1-p-1
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : 30
Connector      : None
Community      : color:00:20 target:64500:11
                 l2-attribute:MTU: 1514 C: 0 P: 1 B: 0
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.20
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
EVPN type      : AUTO-DISC
ESI            : 01:01:01:01:01:01:01:01:01:01
Tag            : 2
Route Dist.    : 192.0.2.20:11
MPLS Label     : 524286
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
Add Paths Send : Default
Last Modified  : 01d23h16m
SRv6 TLV Type  : SRv6 L2 Service TLV (6)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:aaaa:14::
Full Sid       : 2001:db8:aaaa:14:7fff:e000::
Behavior       : End.DX2 (21)
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
SRv6 SubSubTLV : SRv6 SID Structure (1)
Loc-Block-Len  : 48                    Loc-Node-Len   : 16
Func-Len       : 20                    Arg-Len        : 0
Tpose-Len      : 20                    Tpose-offset   : 64

Modified Attributes

Network        : n/a
Nexthop        : 2001:db8::14
Path Id        : None
From           : 2001:db8::14
Res. Nexthop   : fe80::b69e:ffff:fe00:0
Local Pref.    : 100                   Interface Name : int-1-pe-1-p-1
Aggregator AS  : None                  Aggregator     : None
Atomic Aggr.   : Not Atomic            MED            : None
AIGP Metric    : None                  IGP Cost       : 30
Connector      : None
Community      : color:00:20 target:64500:11
                 l2-attribute:MTU: 1514 C: 0 P: 1 B: 0
Cluster        : No Cluster Members
Originator Id  : None                  Peer Router Id : 192.0.2.20
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
EVPN type      : AUTO-DISC
ESI            : 01:01:01:01:01:01:01:01:01:01
Tag            : 2
Route Dist.    : 192.0.2.20:11
MPLS Label     : 524286
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                     Dest Class     : 0
Add Paths Send : Default
Last Modified  : 01d23h16m
SRv6 TLV Type  : SRv6 L2 Service TLV (6)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:aaaa:14::
Full Sid       : 2001:db8:aaaa:14:7fff:e000::
Behavior       : End.DX2 (21)
SRv6 SubSubTLV : SRv6 SID Structure (1)
Loc-Block-Len  : 48                    Loc-Node-Len   : 16
Func-Len       : 20                    Arg-Len        : 0
Tpose-Len      : 20                    Tpose-offset   : 64

-------------------------------------------------------------------------------
Original Attributes

Network        : n/a
Nexthop        : 2001:db8::15
Path Id        : None
From           : 2001:db8::15
Res. Nexthop   : fe80::b6a1:ffff:fe00:0
Local Pref.    : 100                   Interface Name : int-1-pe-1-p-2
Aggregator AS  : None                  Aggregator     : None
Atomic Aggr.   : Not Atomic            MED            : None
AIGP Metric    : None                  IGP Cost       : 30
Connector      : None
Community      : color:00:30 target:64500:11
                 l2-attribute:MTU: 1514 C: 0 P: 0 B: 1
Cluster        : No Cluster Members
Originator Id  : None                  Peer Router Id : 192.0.2.21
Flags          : Used Valid Best IGP
Route Source   : Internal
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
AS-Path        : No As-Path
EVPN type      : AUTO-DISC
ESI            : 01:01:01:01:01:01:01:01:01:01
Tag            : 2
Route Dist.    : 192.0.2.21:11
MPLS Label     : 524286
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
Add Paths Send : Default
Last Modified  : 01d23h16m
SRv6 TLV Type  : SRv6 L2 Service TLV (6)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:aaaa:15::
Full Sid       : 2001:db8:aaaa:15:7fff:e000::
Behavior       : End.DX2 (21)
SRv6 SubSubTLV : SRv6 SID Structure (1)
Loc-Block-Len  : 48                     Loc-Node-Len   : 16
Func-Len       : 20                     Arg-Len        : 0
Tpose-Len      : 20                     Tpose-offset   : 64

Modified Attributes

Network        : n/a
Nexthop        : 2001:db8::15
Path Id        : None
From           : 2001:db8::15
Res. Nexthop   : fe80::b6a1:ffff:fe00:0
Local Pref.    : 100                    Interface Name : int-1-pe-1-p-2
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None                   IGP Cost       : 30
Connector      : None
Community      : color:00:30 target:64500:11
                 l2-attribute:MTU: 1514 C: 0 P: 0 B: 1
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.21
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
EVPN type      : AUTO-DISC
ESI            : 01:01:01:01:01:01:01:01:01:01
Tag            : 2
Route Dist.    : 192.0.2.21:11
MPLS Label     : 524286
Route Tag      : 0
Neighbor-AS    : n/a
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0
Add Paths Send : Default
Last Modified  : 01d23h16m
SRv6 TLV Type  : SRv6 L2 Service TLV (6)
SRv6 SubTLV    : SRv6 SID Information (1)
Sid            : 2001:db8:aaaa:15::
Full Sid       : 2001:db8:aaaa:15:7fff:e000::
Behavior       : End.DX2 (21)
SRv6 SubSubTLV : SRv6 SID Structure (1)
Loc-Block-Len  : 48                     Loc-Node-Len   : 16
Func-Len       : 20                     Arg-Len        : 0
Tpose-Len      : 20                     Tpose-offset   : 64

-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Routes : 2
===============================================================================

[/]
A:admin@pe-1#

[/]
A:admin@pe-1# /show  service id "dual-homing"  segment-routing-v6 detail

===============================================================================
Segment Routing v6 Instance 1 Service 11
===============================================================================
Locator
  Type         Function  SID                                        Status
-------------------------------------------------------------------------------
pe-1-loc
  End.DX2        *524288 2001:db8:aaaa:5:8000::                   ok
===============================================================================
Legend: * - System allocated

[/]
A:admin@pe-1#

[/]
A:admin@pe-1# /show service id "dual-homing" segment-routing-v6 instance 1 destinations

===============================================================================
TEP, SID
===============================================================================
Instance  TEP Address                        Segment Id
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================

===============================================================================
Segment Routing v6 Ethernet Segment Dest
===============================================================================
Instance  Eth SegId                       Num. Macs     Last Change
-------------------------------------------------------------------------------
1         01:01:01:01:01:01:01:01:01:01   0             11/12/2022 15:58:53
-------------------------------------------------------------------------------
Number of entries: 1
-------------------------------------------------------------------------------
===============================================================================

[/]
A:admin@pe-1#

[/]
A:admin@pe-1# /show service id "dual-homing" segment-routing-v6 instance 1 end-dx2

===============================================================================
Segment Routing v6 Instance 1 Service 11
===============================================================================
Locator
  Type         Function  SID                                        Status
-------------------------------------------------------------------------------
pe-1-loc
  End.DX2        *524288 2001:db8:aaaa:5:8000::                   ok
===============================================================================
Legend: * - System allocated

[/]
A:admin@pe-1#
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
[/]
A:admin@pe-1# /show  service id "dual-homing"  segment-routing-v6 instance 1 locator "pe-1-loc"

===============================================================================
Segment Routing v6 Instance 1 Service 11
===============================================================================
Locator
  Type         Function  SID                                       Status
-------------------------------------------------------------------------------
pe-1-loc
  End.DX2        *524288 2001:db8:aaaa:5:8000::                   ok
===============================================================================
Legend: * - System allocated

[/]
A:admin@pe-1#

[/]
A:admin@pe-1# /show router segment-routing-v6 base-routing-instance all

===============================================================================
Segment Routing v6 Base Routing Instance
===============================================================================
Locator
  Type       Function    SID                                       Status/InstId
    SRH-mode Protection  Interface
-------------------------------------------------------------------------------
pe-1-loc
  End                  1 2001:db8:aaaa:5:0:1000::                 ok
    USP
-------------------------------------------------------------------------------
  Auto-allocated End.X: USP Protected,
-------------------------------------------------------------------------------
  End.X          *524289 2001:db8:aaaa:5:8000:1000::             0
    USP        Protected int-1-pe-1-p-1
    ISIS Level: L2 Mac Address: b4:9e:01:01:00:01 Nbr Sys Id: 1920.0000.2004
  End.X          *524290 2001:db8:aaaa:5:8000:2000::             0
    USP        Protected int-1-pe-1-p-2
    ISIS Level: L2 Mac Address: b4:a1:01:01:00:05 Nbr Sys Id: 1920.0000.2010
-------------------------------------------------------------------------------
===============================================================================
Legend: * - System allocated

===============================================================================
Micro Segment Routing v6 Base Routing Instance
===============================================================================
Micro Segment Locator
  Type       Function    SID                                       Status/InstId
    SRH-mode Oper Func   Interface                                 Protection
-------------------------------------------------------------------------------
===============================================================================
Legend: * - System allocated

[/]
A:admin@pe-1#

[/]
A:admin@pe-1# /show router segment-routing-v6 local-sid context  "Base"

===============================================================================
Segment Routing v6 Local SIDs
===============================================================================
SID                                             Type           Function
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
  Locator
  Context
-------------------------------------------------------------------------------
2001:db8:aaaa:5:0:1000::                          End            1
  pe-1-loc
  Base
-------------------------------------------------------------------------------
SIDs : 1
-------------------------------------------------------------------------------
===============================================================================

[/]
A:admin@pe-1#

[/]
A:admin@pe-1# /show router route-table ipv6

===============================================================================
IPv6 Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
2001:db8::4/128 [L]                           Remote  ISIS      02d03h05m  18
       fe80::b69e:ffff:fe00:0-"int-1-pe-1-p-1"                        10
2001:db8::5/128                               Local   Local     02d03h06m  0
       system                                                         0
2001:db8::a/128 [L]                           Remote  ISIS      02d03h05m  18
       fe80::b6a1:ffff:fe00:0-"int-1-pe-1-p-2"                        10
2001:db8::14/128 [L]                          Remote  ISIS      01d23h17m  18
       fe80::b69e:ffff:fe00:0-"int-1-pe-1-p-1"                        20
2001:db8::15/128 [L]                          Remote  ISIS      02d03h05m  18
       fe80::b6a1:ffff:fe00:0-"int-1-pe-1-p-2"                        20
2001:db8::100/120                             Local   Local     02d03h06m  0
       int-1-pe-1-p-1                                                 0
2001:db8::200/120                             Local   Local     02d03h06m  0
       int-1-pe-1-p-2                                                 0
2001:db8::500/120 [L]                         Remote  ISIS      02d03h05m  18
       fe80::b69e:ffff:fe00:0-"int-1-pe-1-p-1"                        20
2001:db8::600/120 [L]                         Remote  ISIS      02d03h05m  18
       fe80::b6a1:ffff:fe00:0-"int-1-pe-1-p-2"                        20
2001:db8::a00/120 [L]                         Remote  ISIS      02d03h05m  18
       fe80::b69e:ffff:fe00:0-"int-1-pe-1-p-1"                        20
2001:db8::6400/120                            Local   Local     02d03h06m  0
       to-ixia                                                        0
2001:db8:aaaa:4::/64                          Remote  ISIS      02d03h05m  18
       2001:db8:aaaa:4::/64 (tunneled:SRV6-ISIS)                      20
2001:db8:aaaa:5::/64                          Local   SRV6      02d03h06m  3
       fe80::201-"_tmnx_fpe_2.a"                                      0
2001:db8:aaaa:5:0:1000::/128                  Local   SRV6      02d03h06m  3
       Black Hole                                                     0
2001:db8:aaaa:5:0:6000::/128                  Local   SRV6-Pol* 02d03h06m  14
       2001:db8::15 (tunneled:SRV6-Policy:917507)                     1
2001:db8:aaaa:5:0:7000::/128                  Local   SRV6-Pol* 02d03h05m  14
       2001:db8::14 (tunneled:SRV6-Policy:917506)                     1
2001:db8:aaaa:5:8000:1000::/128               Local   ISIS      02d03h05m  18
       2001:db8:aaaa:5:8000:1000:: (tunneled:SRV6-ISIS)               10
2001:db8:aaaa:5:8000:2000::/128               Local   ISIS      02d03h05m  18
       2001:db8:aaaa:5:8000:2000:: (tunneled:SRV6-ISIS)               10
2001:db8:aaaa:a::/64                          Remote  ISIS      02d03h05m  18
       2001:db8:aaaa:a::/64 (tunneled:SRV6-ISIS)                      20
2001:db8:aaaa:14::/64                         Remote  ISIS      01d23h17m  18
       2001:db8:aaaa:14::/64 (tunneled:SRV6-ISIS)                     30
2001:db8:aaaa:15::/64                         Remote  ISIS      02d03h05m  18
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
       2001:db8:aaaa:15::/64 (tunneled:SRV6-ISIS)                  30
-------------------------------------------------------------------------------
No. of Routes: 21
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
* indicates that the corresponding row element may have been truncated.

[/]
A:admin@pe-1#

[/]
A:admin@pe-1# /show router tunnel-table ipv6

===============================================================================
IPv6 Tunnel Table (Router: Base)
===============================================================================
Destination                                 Owner     Encap TunnelId  Pref
Nexthop                                     Color           Metric
-------------------------------------------------------------------------------
2001:db8::14/128                            srv6-pol  SRV6  917506    14
  fpe_1.a                                     20              0
2001:db8::15/128                            srv6-pol  SRV6  917507    14
  fpe_1.a                                     30              0
2001:db8:aaaa:4::/64 [L]                    srv6-isis SRV6  524290    0
  fe80::b69e:ffff:fe00:0-"int-1-pe-1-p-1"                     20
2001:db8:aaaa:5:8000:1000::/128 [L]         srv6-isis SRV6  524289    0
  fe80::b69e:ffff:fe00:0-"int-1-pe-1-p-1"                     10
2001:db8:aaaa:5:8000:2000::/128 [L]         srv6-isis SRV6  524292    0
  fe80::b6a1:ffff:fe00:0-"int-1-pe-1-p-2"                     10
2001:db8:aaaa:a::/64 [L]                    srv6-isis SRV6  524293    0
  fe80::b6a1:ffff:fe00:0-"int-1-pe-1-p-2"                     20
2001:db8:aaaa:14::/64 [L]                   srv6-isis SRV6  524297    0
  fe80::b69e:ffff:fe00:0-"int-1-pe-1-p-1"                     30
2001:db8:aaaa:15::/64 [L]                   srv6-isis SRV6  524294    0
  fe80::b6a1:ffff:fe00:0-"int-1-pe-1-p-2"                     30
-------------------------------------------------------------------------------
Flags: B = BGP or MPLS backup hop available
       L = Loop-Free Alternate (LFA) hop available
       E = Inactive best-external BGP route
       k = RIB-API or Forwarding Policy backup hop
===============================================================================

[/]
A:admin@pe-1#

[/]
A:admin@pe-1# /show router fp-tunnel-table 1 ipv6

===============================================================================
IPv6 Tunnel Table Display

Legend:
label stack is ordered from bottom-most to top-most
B - FRR Backup
===============================================================================
Destination                                 Protocol         Tunnel-ID
  Lbl/SID
    NextHop                                                  Intf/Tunnel
  Lbl/SID (backup)
    NextHop   (backup)
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
2001:db8:aaaa:4::/64                          SRV6              524290
  -
    fe80::b69e:ffff:fe00:0-"int-1-pe-1-p-1"                   1/1/c1/1:1
  -
    fe80::b6a1:ffff:fe00:0-"int-1-pe-1-p-2"(B)                1/1/c1/5:1
2001:db8:aaaa:a::/64                          SRV6              524293
  -
    fe80::b6a1:ffff:fe00:0-"int-1-pe-1-p-2"                   1/1/c1/5:1
  -
    fe80::b69e:ffff:fe00:0-"int-1-pe-1-p-1"(B)                1/1/c1/1:1
2001:db8:aaaa:14::/64                         SRV6              524297
  -
    fe80::b69e:ffff:fe00:0-"int-1-pe-1-p-1"                   1/1/c1/1:1
  -
    fe80::b6a1:ffff:fe00:0-"int-1-pe-1-p-2"(B)                1/1/c1/5:1
2001:db8:aaaa:15::/64                         SRV6              524294
  -
    fe80::b6a1:ffff:fe00:0-"int-1-pe-1-p-2"                   1/1/c1/5:1
  -
    fe80::b69e:ffff:fe00:0-"int-1-pe-1-p-1"(B)                1/1/c1/1:1
2001:db8::14/128                              SRV6-Policy       -
  2001:db8:aaaa:4:0:1000::/2001:db8:aaaa:a:0:1000::
    0.140.1.1                                                 pxc-1.b:1
2001:db8::15/128                              SRV6-Policy       -
  2001:db8:aaaa:a:0:1000::/2001:db8:aaaa:4:0:1000::
    0.140.1.1                                                 pxc-1.b:1
2001:db8:aaaa:5:8000:1000::/128               SRV6              524289
  -
    fe80::b69e:ffff:fe00:0-"int-1-pe-1-p-1"                   1/1/c1/1:1
  2001:db8:aaaa:4:0:1000::
    fe80::b6a1:ffff:fe00:0-"int-1-pe-1-p-2"(B)                1/1/c1/5:1
2001:db8:aaaa:5:8000:2000::/128               SRV6              524292
  -
    fe80::b6a1:ffff:fe00:0-"int-1-pe-1-p-2"                   1/1/c1/5:1
  2001:db8:aaaa:a:0:1000::
    fe80::b69e:ffff:fe00:0-"int-1-pe-1-p-1"(B)                1/1/c1/1:1
-------------------------------------------------------------------------------
Total Entries : 8
-------------------------------------------------------------------------------
===============================================================================

[/]
A:admin@pe-1#

[/]
A:admin@pe-1# /show router  bgp  summary
===============================================================================
 BGP Router ID:192.0.2.5        AS:64500       Local AS:64500
===============================================================================
BGP Admin State         : Up          BGP Oper State              : Up
Total Peer Groups       : 1           Total Peers                 : 2
Total VPN Peer Groups   : 0           Total VPN Peers             : 0
Current Internal Groups : 1           Max Internal Groups         : 1
Total BGP Paths         : 45          Total Path Memory           : 16136

Total IPv4 Remote Rts   : 0           Total IPv4 Rem. Active Rts  : 0
Total IPv6 Remote Rts   : 0           Total IPv6 Rem. Active Rts  : 0
Total IPv4 Backup Rts   : 0           Total IPv6 Backup Rts       : 0
Total LblIpv4 Rem Rts   : 0           Total LblIpv4 Rem. Act Rts  : 0
Total LblIpv6 Rem Rts   : 0           Total LblIpv6 Rem. Act Rts  : 0
Total LblIpv4 Bkp Rts   : 0           Total LblIpv6 Bkp Rts       : 0
Total Supressed Rts     : 0           Total Hist. Rts             : 0
Total Decay Rts         : 0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Total VPN-IPv4 Rem. Rts : 0          Total VPN-IPv4 Rem. Act. Rts: 0
Total VPN-IPv6 Rem. Rts : 0          Total VPN-IPv6 Rem. Act. Rts: 0
Total VPN-IPv4 Bkup Rts : 0          Total VPN-IPv6 Bkup Rts     : 0
Total VPN Local Rts     : 1          Total VPN Supp. Rts         : 0
Total VPN Hist. Rts     : 0          Total VPN Decay Rts         : 0

Total MVPN-IPv4 Rem Rts : 0          Total MVPN-IPv4 Rem Act Rts : 0
Total MVPN-IPv6 Rem Rts : 0          Total MVPN-IPv6 Rem Act Rts : 0
Total MDT-SAFI Rem Rts  : 0          Total MDT-SAFI Rem Act Rts  : 0
Total McIPv4 Remote Rts : 0          Total McIPv4 Rem. Active Rts: 0
Total McIPv6 Remote Rts : 0          Total McIPv6 Rem. Active Rts: 0
Total McVpnIPv4 Rem Rts : 0          Total McVpnIPv4 Rem Act Rts : 0
Total McVpnIPv6 Rem Rts : 0          Total McVpnIPv6 Rem Act Rts : 0

Total EVPN Rem Rts      : 4          Total EVPN Rem Act Rts      : 4
Total L2-VPN Rem. Rts   : 0          Total L2VPN Rem. Act. Rts   : 0
Total MSPW Rem Rts      : 0          Total MSPW Rem Act Rts      : 0
Total RouteTgt Rem Rts  : 0          Total RouteTgt Rem Act Rts  : 0
Total FlowIpv4 Rem Rts  : 0          Total FlowIpv4 Rem Act Rts  : 0
Total FlowIpv6 Rem Rts  : 0          Total FlowIpv6 Rem Act Rts  : 0
Total FlowVpnv4 Rem Rts : 0          Total FlowVpnv4 Rem Act Rts : 0
Total FlowVpnv6 Rem Rts : 0          Total FlowVpnv6 Rem Act Rts : 0
Total Link State Rem Rts: 0          Total Link State Rem Act Rts: 0
Total SrPlcyIpv4 Rem Rts: 0          Total SrPlcyIpv4 Rem Act Rts: 0
Total SrPlcyIpv6 Rem Rts: 0          Total SrPlcyIpv6 Rem Act Rts: 0

===============================================================================
BGP Summary
===============================================================================
Legend : D - Dynamic Neighbor
===============================================================================
Neighbor
Description
                   AS PktRcvd InQ  Up/Down   State|Rcv/Act/Sent (Addr Family)
                      PktSent OutQ
-------------------------------------------------------------------------------
2001:db8::14
                64500    5682    0 01d23h17m 2/2/1 (Evpn)
                         5680    0
2001:db8::15
                64500    6146    0 02d03h05m 2/2/1 (Evpn)
                         6135    0
-------------------------------------------------------------------------------

[/]
A:admin@pe-1#

[/]
A:admin@pe-1# /show router isis adjacency

===============================================================================
Rtr Base ISIS Instance 0 Adjacency
===============================================================================
System ID               Usage State Hold Interface                    MT-ID
-------------------------------------------------------------------------------
p-1                     L2    Up    26   int-1-pe-1-p-1               0
p-2                     L2    Up    9    int-1-pe-1-p-2               0
-------------------------------------------------------------------------------
Adjacencies : 2
===============================================================================

[/]
A:admin@pe-1#
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
[/]
A:admin@pe-1# /show router isis database

===============================================================================
Rtr Base ISIS Instance 0 Database
===============================================================================
LSP ID                                Sequence  Checksum Lifetime Attributes
-------------------------------------------------------------------------------

Displaying Level 1 database
-------------------------------------------------------------------------------
Level (1) LSP Count : 0

Displaying Level 2 database
-------------------------------------------------------------------------------
p-1.00-00                             0x13e     0x8dfc   1100     L1L2
p-1.01-00                             0x135     0x4b04   1175     L1L2
pe-1.00-00                            0x13a     0x639f   1161     L1L2
pe-1.01-00                            0x133     0xe585   775      L1L2
p-2.00-00                             0x137     0xfd42   796      L1L2
p-2.01-00                             0x135     0x6dc8   648      L1L2
p-2.02-00                             0x135     0xe56c   752      L1L2
p-2.03-00                             0x133     0xf859   1034     L1L2
bng-1.00-00                           0x13e     0x1410   1095     L1L2
bng-2.00-00                           0x138     0x757c   985      L1L2
Level (2) LSP Count : 10
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Diameter Base Protocol: Establishing a Diameter Peer Connection

This chapter provides information about configuring and troubleshooting the Diameter Base protocol to establish a Diameter peer connection.

Topics in this chapter include:

- Applicability
- Overview
- Configuration
- Conclusion

Applicability

This information and configuration in this chapter are based on SR OS Release 19.10.R1.

Note:
This chapter covers the Diameter base protocol implementation that is available from SR OS Release 16.0.R4 onward (configured in **aaa** CLI context as **diameter node**). The legacy Diameter base implementation (configured in the **aaa** CLI context as **diameter-peer-policy**) is supported in maintenance mode only, without any further feature enhancement planned. Nokia recommends using or transitioning to the new Diameter base protocol implementation.

Overview

Diameter is an Authentication, Authorization and Accounting (AAA) protocol defined by the IETF in RFC 6733, *Diameter Base Protocol*. While historically wireline access networks were largely based on RADIUS for subscriber authentication, authorization, and accounting, it was decided by 3rd Generation Partnership Project (3GPP) that wireless access networks will be largely based on Diameter. Over time, operators are looking to converge both types of networks, and one of the aspects of this is to replace RADIUS in wireline access networks by Diameter.

Diameter is based on three layers: the transport layer, the Diameter base protocol layer and the Diameter applications as shown in Figure 342: Diameter protocol stack.

Figure 342: Diameter protocol stack

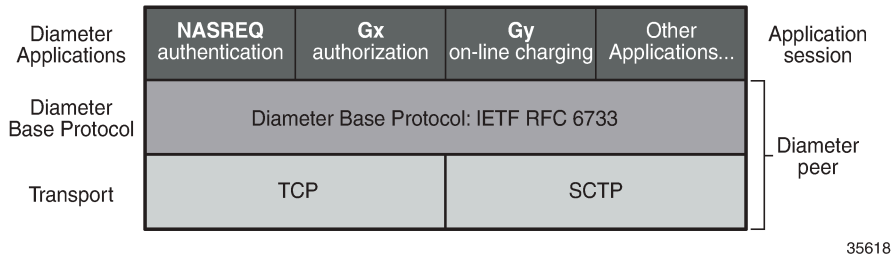

The bottom layer is the transport layer and can be either TCP or SCTP. SR OS supports TCP. The Diameter base protocol implementation in SR OS is based on RFC 6733. The top layer contains

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

the Diameter applications. SR OS supports NASREQ for authentication, Gx for authorization, policy management and usage monitoring and Gy or Diameter Credit Control Application (DCCA) for online charging.

Figure 343: Diameter network topology shows a Diameter network topology that will be used in the configuration examples in this chapter.

A Diameter Client (BNG) is connected via peer-1a and peer-1b to two Diameter Agents (DRA-1 and DRA-2) that provide connectivity to the Diameter Application Servers (PCRF). Via these peers, the BNG can authenticate and perform policy control of subscriber sessions using the NASREQ and Gx applications. The same Diameter Client (BNG) is also directly connected to another Diameter Application Server (OCS) via peer-2. Via this peer, on-line charging can be done for the subscriber sessions using the Gy application.

Figure 343: Diameter network topology

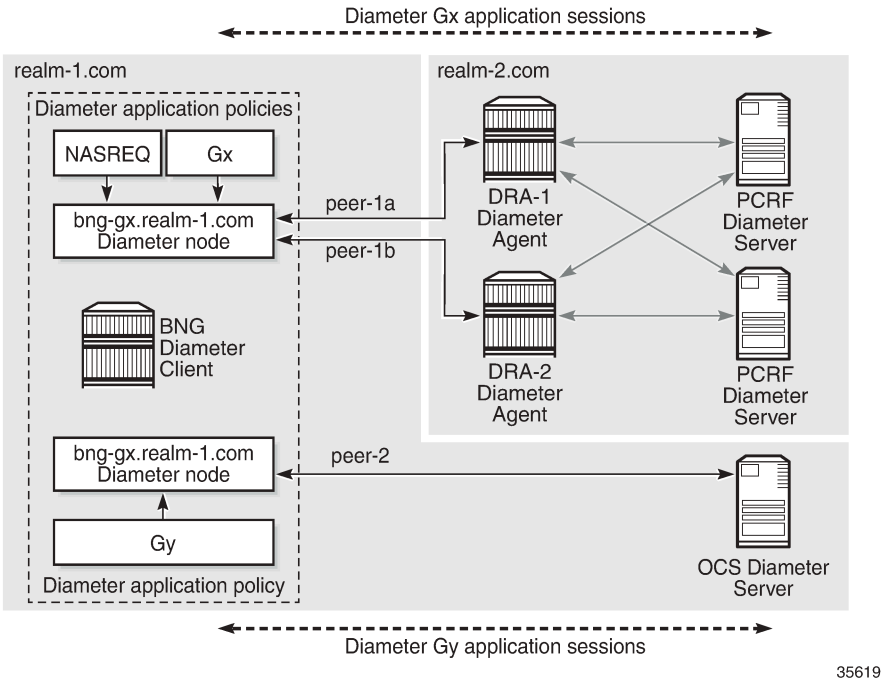

Configuration

The Diameter base protocol and the Diameter applications are configured separately, where the Diameter base protocol must be configured first, and the Diameter applications next. The transport layer configuration is part of the Diameter base protocol layer. This example describes the Diameter base protocol configuration.

The Diameter base protocol and the corresponding transport layer configuration is based on Diameter Nodes. Each Diameter Node represents a Diameter routing instance and contains a list of peers in the routing domain that provide direct or indirect connectivity to application servers.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

An example Diameter node configuration that corresponds with the topology in Figure 343: Diameter network topology is shown below.

```
configure {
    aaa {
        diameter {
            node "bng-gx.realm-1.com" {
                description "Authentication and Policy Management"
                connection {
                    ipv4 {
                        local-address 192.0.2.2
                    }
                    ipv6 {
                        local-address 2001:db8::2
                    }
                }
                peer index 1 {
                    admin-state enable
                    address 2001:db8:2:6::1
                    destination-host "dra-1.realm-2.com"
                    preference 10
                }
                peer index 2 {
                    admin-state enable
                    address 172.16.7.2
                    destination-host "dra-2.realm-2.com"
                    preference 20
                }
            }
            node "bng-gy.realm-1.com" {
                description "Credit Control"
                origin-realm "realm-1.com"
                router-instance "management"
                peer index 1 {
                    admin-state enable
                    address 192.99.3.0
                    destination-host "ocs.realm-1.com"
                }
            }
        }
```

A Diameter node configuration requires a unique origin host as key. The origin host is used in Diameter application policies to associate the application with the node. All Diameter base and application messages forwarded via the peers of that node use the configured origin host in the Origin-Host AVP. The value for the Origin-Realm AVP is by default derived from the configured origin host: the realm is the part of the origin host after the first dot (".") as delimiter or equal to the origin host when it does not contain a delimiter. For example, for **node** "*bng-gx.realm-1.com*": Origin-Host = "bng-gx.realm-1.com", Origin-Realm = "realm-1.com". It is also possible to explicitly configure an origin realm as shown in the example for the **node** "*bng-gy.realm-1.com*".

A node configuration can include a routing context, and an IPv4 and/or IPv6 source address. These parameters are used to establish the TCP transport connection for all peers in the node. The specified **local-address** must be a reachable local interface address in the specified or in the default router instance. For **node** "*bng-gx.realm-1.com*" in the example, no router instance is specified. By default, the TCP connections are established in the Base router using the specified local addresses. For **node** "*bng-gy.realm-1.com*" in the example, the out of band router instance "*management*" is used to establish the TCP connection of its peer. As no local address is specified, the system will automatically select an interface address, in this case an out of band IP address configured in the Boot Options File (BOF).

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Within a Diameter node, up to 5 peers can be configured with an index value between 1 and 5 as key, an IPv4 or IPv6 destination address for the TCP connection, and a mandatory destination host that is used as Destination-Host AVP value for all Diameter base messages on the peer. In a Diameter node, one peer is selected to forward application messages for a specific application session. The other peers provide redundancy when supported by the Diameter application, such as Gy session failover. A Diameter peer for application messages is selected based on following criteria:

1. Forwarding:

   If the application message contains a Destination-Host AVP, select the peer in the peer table with a matching configured destination host. This is the forwarding phase.

2. Routing:

   When the lookup in the peer table fails, perform a lookup in the realm routing table and select the peer with realm name equal to the Destination-Realm AVP in the application message, and with matching application ID. When multiple peers are matched, select in order or priority until a single peer is found:

   a. the peer with the lowest configured preference (default preference is 50)

   b. the peer with the lowest index

3. Default peer:

   When both forwarding lookup in the peer table and routing lookup in the realm routing table were unsuccessful, use the peer configured as **default-peer**. Only a single peer in a node can be configured as a default-peer. Multiple peers in a node configured as default-peer results in a validation error:

```
MINOR: MGMT_CORE #5001: configure aaa diameter node "bng-gx.realm-1.com" peer index 2 -
 Multiple default peer is not allowed
```

For **node** "*bng-gx.realm-1.com*" in the example, peer-1a with index 1 has a configured preference of 10 and peer-2 with index 2 has a configured preference of 20. Diameter Gx application messages will fail the peer table lookup as the destination host of the PCRF will not be present (no direct connection between Diameter client and Diameter application server):

```
# /show aaa diameter-node "bng-gx.realm-1.com" peers

===============================================================================
Peers
===============================================================================
Host identity                       Status        Default Preference Active
-------------------------------------------------------------------------------
dra-1.realm-2.com                   I-Open        No      10         Yes
dra-2.realm-2.com                   I-Open        No      20         Yes
-------------------------------------------------------------------------------
No. of peers: 2
===============================================================================
```

Instead a realm routing table lookup is performed to find the peer for forwarding the application messages. In this case peer-1a (dra-1.realm-2.com) is selected based on the matching destination realm (realm-2.com), application ID (Gx) and the lower preference value:

```
# /show aaa diameter-node "bng-gx.realm-1.com" routing-table

===============================================================================
Routes
===============================================================================
Realm-Name
           Application   Pref. Id  Server-Identifier
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
realm-2.com
          nasreq gx     10    1   dra-1.realm-2.com
realm-2.com
          nasreq gx     20    2   dra-2.realm-2.com
-------------------------------------------------------------------------------
No. of routes: 2
===============================================================================
```

The realm routing table is populated based on the Origin-Realm AVP and Application-Id AVP received in the Capability Exchange Answer message together with the configured index and preference values.

Note that Diameter answer messages do not rely on peer or realm routing table lookups. Answers are forwarded over the same route in the reverse direction of the matching requests. This is achieved with a transactional cache in each traversed Diameter node, using the Hop-by-Hop AVP to match requests with answers.

When enabling the peer (**admin-state enable**), the system tries to establish the transport TCP connection. Once the TCP session is up, the system starts a Diameter Capability Exchange using the configured Diameter identity (Origin-Host and Origin-Realm AVPs) and advertising support for all SR OS Diameter applications in Application-Id AVP's (NASREQ, Gx, and Gy). When the Origin-Host AVP in the received CEA message corresponds with the destination host configured for the peer (case insensitive) and at least one application in the CEA is common with the SR OS advertised applications, then the peer moves to the I-Open state (I from Initiator). An example of a Capability Exchange is illustrated in detail in the troubleshooting section.

Optionally, a connection and a watchdog timer can be configured in the Diameter node:

```
configure {
    aaa {
        diameter {
            node "bng-gx.realm-1.com" {
                connection {
                    timer 30
                    ---snip---
                }
                peer index 1 {
                    connection-timer 30
                    watchdog-timer 30
                    ---snip---
```

- **connection-timer**

  The connection timer or Tc timer controls the frequency at which a transport connection is attempted to be established. The default value is 30 seconds. This timer can be configured per node to be used by all peers or overridden per peer.

- **watchdog-timer**

  The watchdog timer controls the frequency at which Device-Watchdog-Request messages are transmitted to the peer, and is called the Tw timer in RFC 3539, *Authentication, Authorization and Accounting (AAA) Transport Profile*. A small timer results in a faster detection of a peer failure at the expense of generating more messages. The timer is configured per peer and its default value is 30 seconds.

A Python policy can be configured in the Diameter node to manipulate Diameter messages transmitted to and/or received on its peers.

```
configure {
    aaa {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        diameter {
            node "bng-gy.realm-1.com" {
                python-policy "py-diameter-1"
                ---snip---
```

Manipulating Diameter messages, such as changing the content or format of AVPs using Python is out of the scope of this chapter.

By default, Diameter messages are sent with a DSCP set to AF41. The DSCP value can be changed with the sgt-qos configuration:

```
# /configure router sgt-qos dscp application diameter dscp nc1
```

SR OS uses TCP as transport and the TCP destination port number is fixed to the standard Diameter base protocol port 3868. The source port is randomly chosen from the ephemeral port range.

Troubleshooting

The status and statistics of the Diameter peers can be verified with following show commands:

```
# /show aaa diameter-node "bng-gx.realm-1.com" peers

===============================================================================
Peers
===============================================================================
Host identity                       Status        Default Preference Active
-------------------------------------------------------------------------------
dra-1.realm-2.com                   I-Open        No      10         Yes
dra-2.realm-2.com                   I-Open        No      20         Yes
-------------------------------------------------------------------------------
No. of peers: 2
===============================================================================
```

```
# /show aaa diameter-node "bng-gx.realm-1.com" peer "dra-1.realm-2.com"

===============================================================================
Peer "dra-1.realm-2.com"
===============================================================================
Index                       : 1
Status                      : I-Open
Administrative state        : enabled
Active                      : Yes
Active applications         : nasreq gx
Last disconnect cause       : rebooting
Preference                  : 10
Default peer                : No
Connection timer (s)        : N/A
Watchdog timer (s)          : 13
Pending messages            : 0
Remote realm                : realm-2.com
Remote IP address           : 2001:db8:2:6::1
Remote TCP port             : 3868
Remote Origin-State-Id      : 1574235027
Local host identity         : bng-gx.realm-1.com
Local realm                 : realm-1.com
Local IP address            : 2001:db8::2
Local TCP port              : 53734
Last management change      : 11/19/2019 15:05:48
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
```

```
# /show aaa diameter-node "bng-gx.realm-1.com" peer "dra-1.realm-2.com" statistics

===============================================================================
Peer "dra-1.realm-2.com"
===============================================================================
Message                              Sent                 Received
-------------------------------------------------------------------------------
Capabilities-Exchange-Request        7                    0
Capabilities-Exchange-Answer         0                    7
Disconnect-Peer-Request              1                    4
Disconnect-Peer-Answer               4                    1
Device-Watchdog-Request              1217                 778
Device-Watchdog-Answer               778                  1217
Application message request          0                    0
Application message answer           0                    0

Last cleared time: N/A
===============================================================================
```

To clear the peer statistics, use following command:

```
# /clear aaa diameter-node "bng-gx.realm-1.com" peer "dra-1.realm-2.com" statistics
```

Diameter debugging is split between node and application level:

```
debug
    diameter
        application
            policy "diam-nasreq-1"
                session-messages
            exit
        exit
        node "bng-gx.realm-1.com"
            peer "dra-1.realm-2.com"
                peer-to-peer
            exit
        exit
    exit
exit
```

In this chapter, the Diameter base protocol debugging for peer messages is explained, configured at the node level debug. When a Python script is active for the node, the debug messages are logged after Python processing.

To debug the Diameter base protocol messages for **peer** "*dra-1.realm-2.com*", use the following debug commands:

```
debug
    diameter
        node "bng-gx.realm-1.com"
            peer "dra-1.realm-2.com"
                peer-to-peer
            exit
        exit
    exit
exit
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The **peer-to-peer** option enables debug output for all Diameter base messages of the specified peer: Capabilities Exchange, Device Watchdog and Disconnect Peer messages. By default, error conditions are also logged in the debug output. Debug for error conditions can be disabled per Diameter node or per peer with the debug option **no on-error**. Errors reported at the node level include Diameter base errors that are unrelated to a peer, such as a routing problem for a Diameter application message. Errors reported at the peer level include all errors that occur after peer selection and peer connection errors.

Let's start with the peer connection in Closed state (remote end rebooting):

```
# /show aaa diameter-node "bng-gx.realm-1.com" peer "dra-1.realm-2.com"

===============================================================================
Peer "dra-1.realm-2.com"
===============================================================================
Index                       : 1
Status                      : Closed
Administrative state        : enabled
Active                      : No
Active applications         :
Last disconnect cause       : rebooting
Preference                  : 10
Default peer                : No
Connection timer (s)        : 18
Watchdog timer (s)          : N/A
Pending messages            : 0
Remote realm                : (Not Specified)
Remote IP address           : (Not Specified)
Remote TCP port             : (Not Specified)
Remote Origin-State-Id      : (Not Specified)
Local host identity         : bng-gx.realm-1.com
Local realm                 : realm-1.com
Local IP address            : (Not Specified)
Local TCP port              : (Not Specified)
Last management change      : 11/19/2019 15:05:48
===============================================================================
```

The *Connection timer(s)* field in above peer details output show that in 18 seconds, a new connection attempt will be made, followed by a Capabilities Exchange when successful. The transmitted Capabilities-Exchange-Request (CER) and received Capabilities-Exchange-Answer (CEA) are shown in the debug output:

```
233997 2019/11/20 19:17:16.271 UTC minor: DEBUG #2001 Base DIAMETER
DIAMETER: Message Transmission
Transmit: "CER"
Application policy: N/A
Node: "bng-gx.realm-1.com"
Received peer: N/A
Transmit peer: "dra-1.realm-2.com"
Python policy: N/A
  Header
    ver 1  len 284  flags R-------  code 257
    app-id 0 hbh-id 19864  e2e-id 3486524428
  AVPs
    origin-host (264) -M------ [26]
      data [18] (DiameterIdentity) : bng-gx.realm-1.com
    origin-realm (296) -M------ [19]
      data [11] (DiameterIdentity) : realm-1.com
    host-ip-addr (257) -M------ [26]
      data [18] (Address) : ipv6 2001:db8::2
    vendor-id (266) -M------ [12]
      data [4] (Unsigned32) : 6527
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    product-name (269) -------- [13]
      data [5] (UTF8String) : SR-OS
    auth-appl-id (258) -M------ [12]
      data [4] (Unsigned32) : 1 : Nasreq
    auth-appl-id (258) -M------ [12]
      data [4] (Unsigned32) : 4 : Gy
    auth-appl-id (258) -M------ [12]
      data [4] (Unsigned32) : 16777238 : Gx
    vend-specific-appl-id (260) -M------ [32]
      data [24] (Grouped)
        vendor-id (266) -M------ [12]
          data [4] (Unsigned32) : 10415
        auth-appl-id (258) -M------ [12]
          data [4] (Unsigned32) : 4 : Gy
    vend-specific-appl-id (260) -M------ [32]
      data [24] (Grouped)
        vendor-id (266) -M------ [12]
          data [4] (Unsigned32) : 10415
        auth-appl-id (258) -M------ [12]
          data [4] (Unsigned32) : 16777238 : Gx
    supported-vendor-id (265) -M------ [12]
      data [4] (Unsigned32) : 3561
    supported-vendor-id (265) -M------ [12]
      data [4] (Unsigned32) : 6527
    supported-vendor-id (265) -M------ [12]
      data [4] (Unsigned32) : 10415
    supported-vendor-id (265) -M------ [12]
      data [4] (Unsigned32) : 13019
    firmware-revision (267) -------- [12]
      data [4] (Unsigned32) : 191001


233998 2019/11/20 19:17:16.275 UTC minor: DEBUG #2001 Base DIAMETER
DIAMETER: Message Reception
Receive: "CEA"
Application policy: N/A
Node: "bng-gx.realm-1.com"
Received peer: "dra-1.realm-2.com"
Transmit peer: N/A
Python policy: N/A
  Header
    ver 1  len 240  flags --------  code 257
    app-id 0 hbh-id 19864  e2e-id 3486524428
  AVPs
    result-code (268) -M------ [12]
      data [4] (Unsigned32) : 2001 : DIAM_RESCODE_SUCCESS
    origin-host (264) -M------ [25]
      data [17] (DiameterIdentity) : dra-1.realm-2.com
    origin-realm (296) -M------ [19]
      data [11] (DiameterIdentity) : realm-2.com
    host-ip-addr (257) -M------ [26]
      data [18] (Address) : ipv6 2001:db8:2:6::1
    vendor-id (266) -M------ [12]
      data [4] (Unsigned32) : 6527
    product-name (269) -------- [28]
      data [20] (UTF8String) : PythonDiameterAgent1
    origin-state-id (278) -M------ [12]
      data [4] (Unsigned32) : 1574277432
    supported-vendor-id (265) -M------ [12]
      data [4] (Unsigned32) : 10415
    auth-appl-id (258) -M------ [12]
      data [4] (Unsigned32) : 1 : Nasreq
    auth-appl-id (258) -M------ [12]
      data [4] (Unsigned32) : 16777238 : Gx
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    vend-specific-appl-id (260) -M------ [32]
      data [24] (Grouped)
        vendor-id (266) -M------ [12]
          data [4] (Unsigned32) : 10415
        auth-appl-id (258) -M------ [12]
          data [4] (Unsigned32) : 16777238 : Gx
    firmware-revision (267) -------- [12]
      data [4] (Unsigned32) : 1
```

The result of the successful Capabilities Exchange is that the peer moved to the I-Open state, ready to forward NASREQ and Gx application messages:

```
# /show aaa diameter-node "bng-gx.realm-1.com" peer "dra-1.realm-2.com"

===============================================================================
Peer "dra-1.realm-2.com"
===============================================================================
Index                       : 1
Status                      : I-Open
Administrative state        : enabled
Active                      : Yes
Active applications         : nasreq gx
Last disconnect cause       : rebooting
Preference                  : 10
Default peer                : No
Connection timer (s)        : N/A
Watchdog timer (s)          : 9
Pending messages            : 0
Remote realm                : realm-2.com
Remote IP address           : 2001:db8:2:6::1
Remote TCP port             : 3868
Remote Origin-State-Id      : 1574277432
Local host identity         : bng-gx.realm-1.com
Local realm                 : realm-1.com
Local IP address            : 2001:db8::2
Local TCP port              : 55199
Last management change      : 11/19/2019 15:05:48
===============================================================================
```

The *Watchdog timer(s)* field in preceding peer details output shows that in 9 seconds, a Device Watchdog exchange will be initiated. The transmitted Device-Watchdog-Request (DWR) and received Device-Watchdog-Answer (DWA) are shown in the debug output:

```
233999 2019/11/20 19:17:44.268 UTC minor: DEBUG #2001 Base DIAMETER
DIAMETER: Message Transmission
Transmit: "DWR"
Application policy: N/A
Node: "bng-gx.realm-1.com"
Received peer: N/A
Transmit peer: "dra-1.realm-2.com"
Python policy: N/A
  Header
    ver 1  len 68  flags R-------  code 280
    app-id 0 hbh-id 19865  e2e-id 3486524431
  AVPs
    origin-host (264) -M------ [26]
      data [18] (DiameterIdentity) : bng-gx.realm-1.com
    origin-realm (296) -M------ [19]
      data [11] (DiameterIdentity) : realm-1.com


234000 2019/11/20 19:17:44.271 UTC minor: DEBUG #2001 Base DIAMETER
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
DIAMETER: Message Reception
Receive: "DWA"
Application policy: N/A
Node: "bng-gx.realm-1.com"
Received peer: "dra-1.realm-2.com"
Transmit peer: N/A
Python policy: N/A
  Header
    ver 1  len 92  flags --------  code 280
    app-id 0 hbh-id 19865  e2e-id 3486524431
  AVPs
    result-code (268) -M------ [12]
      data [4] (Unsigned32) : 2001 : DIAM_RESCODE_SUCCESS
    origin-host (264) -M------ [25]
      data [17] (DiameterIdentity) : dra-1.realm-2.com
    origin-realm (296) -M------ [19]
      data [11] (DiameterIdentity) : realm-2.com
    origin-state-id (278) -M------ [12]
      data [4] (Unsigned32) : 1574277432
```

Now let's try to bring up the peer in the **node** *bng-gy.realm-1.com*:

```
# /show aaa diameter-node "bng-gy.realm-1.com" peers

===============================================================================
Peers
===============================================================================
Host identity                       Status        Default Preference Active
-------------------------------------------------------------------------------
ocs.realm-1.com                     Closed        No      50         No
-------------------------------------------------------------------------------
No. of peers: 1
===============================================================================
```

Debug is enabled at the peer level for error conditions without the **peer-to-peer** option. Failures are reported, but not all transmitted and received peer messages.

```
debug
    diameter
        node "bng-gy.realm-1.com"
            peer "ocs.realm-1.com"
            exit
        exit
    exit
exit
```

The Diameter server is provisioned with an origin host different from the configured destination host for the peer, resulting in a failure and peer reset:

```
234330 2019/11/22 14:57:32.272 UTC minor: DEBUG #2001 management DIAMETER
DIAMETER: Failure
Receive: "CEA"
Application policy: N/A
Node: "bng-gy.realm-1.com"
Received peer: "ocs.realm-1.com"
Transmit peer: N/A
Python policy: N/A
Result code: "DIAM_RESCODE_INVALID_AVP_VALUE"
Error message: "mismatch with locally stored information"
Failed AVP:
    origin-host (264) -M------ [27]
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
      data [19] (DiameterIdentity) : ocs.wrong-realm.com

Message:
  Header
    ver 1  len 176  flags --------  code 257
    app-id 0 hbh-id 6050  e2e-id 3486524894
  AVPs
    result-code (268) -M------ [12]
      data [4] (Unsigned32) : 2001 : DIAM_RESCODE_SUCCESS
    origin-host (264) -M------ [27]
      data [19] (DiameterIdentity) : ocs.wrong-realm.com
    origin-realm (296) -M------ [23]
      data [15] (DiameterIdentity) : wrong-realm.com
    host-ip-addr (257) -M------ [14]
      data [6] (Address) : ipv4 192.99.3.0
    vendor-id (266) -M------ [12]
      data [4] (Unsigned32) : 6527
    product-name (269) -------- [28]
      data [20] (UTF8String) : PythonDiameterServer
    origin-state-id (278) -M------ [12]
      data [4] (Unsigned32) : 1574434643
    auth-appl-id (258) -M------ [12]
      data [4] (Unsigned32) : 4 : Gy
    firmware-revision (267) -------- [12]
      data [4] (Unsigned32) : 1


234331 2019/11/22 14:57:32.272 UTC minor: DEBUG #2001 management DIAMETER
DIAMETER: Peer Reset
Node: "bng-gy.realm-1.com"
Peer: "ocs.realm-1.com"
Reason: "failed to parse received CEA"
```

Events

Following events are defined for the Diameter base protocol:

```
===============================================================================
Application
 ID#    Event Name                       P   g/s     Logged     Dropped
-------------------------------------------------------------------------------
   2007 tmnxDiamMessageDropped           MI  thr          0           0
   2008 tmnxDiamNdPeerStatActiveChanged  MI  thr         46           0
===============================================================================
```

The **tmnxDiamNdPeerStatActiveChanged** event is generated when the state of a Diameter peer toggles between active / not active:

```
38080 2019/11/22 14:52:02.269 UTC MINOR: DIAMETER #2008 management peer state change
"DIAMETER node bng-gy.realm-1.com, peer ocs.realm-1.com is active"
```

The **tmnxDiamMessageDropped** event is generated when Diameter base drops a malformed message.

Conclusion

As a result of fixed mobile network convergence, Diameter is used in fixed access networks as an alternative for Radius based AAA. Diameter peering provides reliable and secure transport with peer

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

redundancy. Its functionality is defined in a base Diameter protocol specified in RFC 6733. Various applications can be layered on top of base Diameter and they can utilize the robust transport capabilities that Diameter provides.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

L2TP for Subscriber Access — LAC

This chapter provides information about L2TP for subscriber access.

Topics in this chapter include:

Applicability

This chapter describes L2TP Access Concentrator (LAC) support for the L2TP Access Aggregation (LAA) architecture model and was initially written for SR OS Release 11.0.R4. The MD-CLI in the current edition is based on Release 19.5.R1. PPP hosts are supported in a Routed CO model (with IES or VPRN services) using ATM, Ethernet or Ethernet over Pseudowire SAPs. A description of the L2TP Tunnel Switch (LTS) and L2TP Network Server (LNS) functions are out of the scope of this chapter.

Overview

PPP access architectures

The Broadband Forum proposes two architectures for Point-to-Point Protocol (PPP) access.

- The PPP Terminated Aggregation architecture (PTA)
- The L2TP Access Aggregation architecture (LAA)

The PTA architecture (local-access model) uses the Broadband Network Gateway (BNG) to terminate user PPP sessions (see scenario PPPoE-5 in Figure 344: PPP access architectures).

The LAA architecture (which is a tunneled access model) uses a LAC and an LNS to transport PPP sessions from the LAC to the LNS which performs tunnel termination (see scenario PPPoE-1 and PPPoE -2 in Figure 344: PPP access architectures).

Optionally, an LTS can be used in the transport network to perform the grooming of traffic between tunnels (see scenarios PPPoE-3 and PPPoE-4 in Figure 344: PPP access architectures).

The LNS is the logical termination point of the PPP sessions originated by the remote clients and tunneled by the LAC/LTS.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 344: PPP access architectures

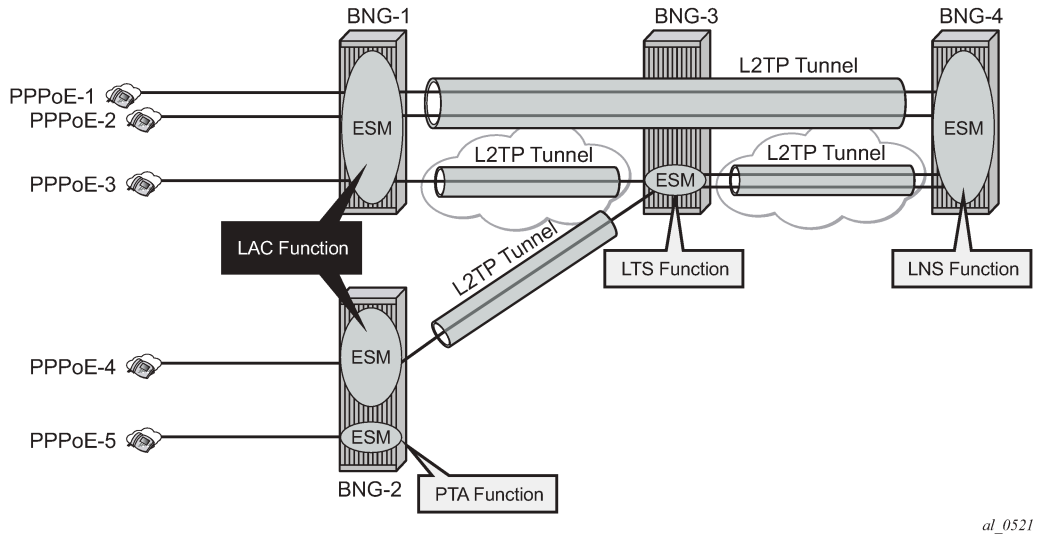

L2TP tunnels - LAC and LNS reachability options

The router instance where the L2TP tunnel starts and where ESM is handled can be one and the same, but does not need to be the same. The LNS peer address can be reachable via IP, BGP/IGP shortcuts, over a spoke SDP (GRE/MPLS), RFC 4364 VPRNs (*BGP/MPLS IP Virtual Private Networks*), but cannot be an address belonging to a directly connected interface. See Figure 345: Supported LT2P reachability options.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 345: Supported LT2P reachability options

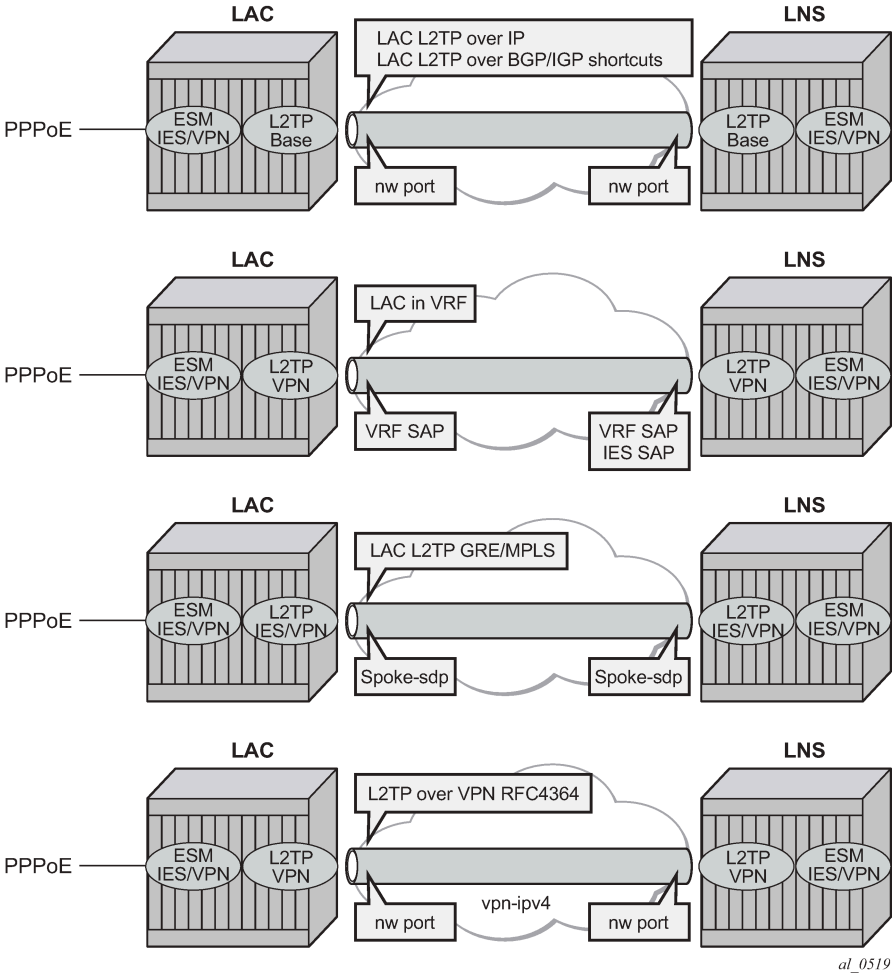

Recap of the L2TPv2 protocol

L2TPv2 is a client-server protocol relying on UDP and encapsulates Layer 2 packets such as PPP for transmission across a network. L2TPv2 passes control and data messages over separate control and data channels, thus defines following message types:

- Control messages—The in-band control channel passes sequenced control messages, supporting connection management, call management, error reporting, and session control. Optionally, a shared-secret challenge authentication method can be used between the tunnel endpoints.

  The following messages are used for L2TP tunnel management:

  – Tunnel setup (Control Connection Management)

    - Start-Control-Connection-Request (SCCRQ)
    - Start-Control-Connection-Reply (SCCRP)
    - Start-Control-Connection-Connected (SCCCN)
    - Stop-Control-Connection-Notification (StopCCN)

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

  - Tunnel keepalive
    - Hello (HELLO)

  The following messages are used for L2TP session (call) management:

  - Session setup over an existing tunnel
    - Incoming-Call-Request (ICRQ)
    - Incoming-Call-Reply (ICRP)
    - Incoming-Call-Connected (ICCN)
    - Call-Disconnect-Notify (CDN)

  Zero-Length Body (ZLB) messages are control packets with an L2TP header only and are used to explicitly acknowledge packets, making the control channel reliable.

  L2TP message encoding is done through Attribute Value Pairs (AVP).

- Data messages — Data messages encapsulate the PPP frames that are sent into the L2TP tunnel.

L2TPv2 sessions run over an L2TP tunnel and are referenced by an L2TP session-id. An L2TP tunnel can carry none, one, or multiple L2TP sessions. An L2TP session corresponds to a PPPoE session. L2TPv3 for LAC-LNS dynamic tunnel setup is not supported.

L2TP header and AVP layout

The L2TPv2 header consists of following fields (RFC 2611, *URN Namespace Definition Mechanisms*):

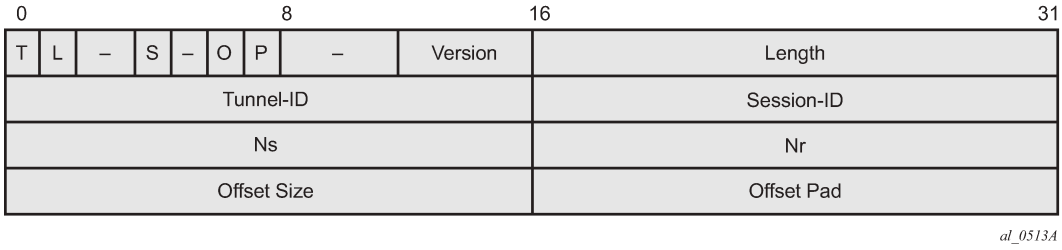

Table 20: L2TPv2 header fields and descriptions

Field	Description
T	Type of L2TP message (1 bit): 0—data message 1—control message
L	Indicates if the optional Length field is present in the message (1 bit): 0—the field is left out of the message entirely 1—the field is included (must be included in control messages)
-	Reserved for future use, must be set to zero.
S	Indicates if the Ns and Nr fields are present (1 bit): 0 — the fields are left out of the message; entirely 1 — the fields are included (must be included in control messages)
O	Indicates if the Offset field is present (1 bit): 0 — the field is left out of the message entirely (must be left out of control messages); 1 — the field is included

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Field	Description
P	Used with data messages only. Indicates priority of the data message (1 bit): 0 — no (this value is used for all control messages); 1 — yes
Version	The version of the message (4 bits): 2 — this is the latest version of the L2TP data message header; 1 — indicates an L2F packet as described in RFC 2341 Packets with an unknown version number are discarded.
Length	The total length (in bytes) of the L2TP message (16 bits).
Tunnel-ID	Identifies the L2TP tunnel (that is, the control connection).This number has local significance — each end gives the same tunnel different tunnel IDs. The ID refers to the receiver, not the sender, and is assigned during tunnel creation (16 bits).
Session-ID	Identifies the PPP session within a tunnel. This number has local significance — each end gives the same session different session IDs. The ID refers to the receiver, not the sender, and is assigned during session creation (16 bits).
Ns	The sequence number of the message. This is mandatory for control messages (to enable re-transmission of lost messages) but optional for data messages (to re-order data messages that were mis-sequenced during forwarding). The number, which starts at 0 and increments by 1, is assigned by an L2TP peer for each session in a tunnel (16 bits).
Nr	The sequence number of the next control message expected to be received. This is equal to the sequence number of last received control message plus 1. Used by the receiving peer to ensure that control messages are sent in order without duplication. In data messages, the field (if present as indicated by the S bit) is ignored (16 bits).
Offset Size	The location of the L2TP payload, expressed as the number of octets from the start of the message header (16 bits).
Offset pad	User-defined bytes used to pad the message header so that the payload starts at the location indicated by the Offset Size field (16 bits).

The AVP header consists of following fields (RFC 2611):

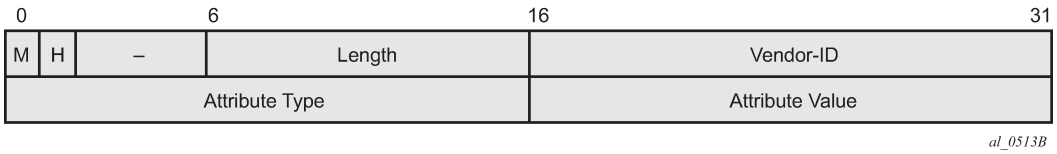

Table 21: AVP header fields and descriptions

Field	Description
M	Mandatory bit — If the M bit is set on an unrecognized AVP within a message associated with a particular session, the session associated with this message MUST be terminated (1 bit).
H	Hidden bit — Identifies the hiding of data in the Attribute-Value field of an AVP. This capability can be used to avoid the passing of sensitive data, such as user passwords, as clear text in an AVP. The H-bit MUST only be set if a shared secret exists between the LAC and LNS. The shared secret is the same secret that is used for tunnel authentication. If the H-bit is set in any

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Field	Description
	AVP(s) in a given control message, a Random Vector AVP must also be present in the message and MUST precede the first AVP having an H bit of 1 (1 bit).
-	Reserved for future use, must be set to zero (4 bits).
Length	Indicates the total number of bytes (including the overall length and bitmask fields) contained in this AVP (10 bits).
Vendor-id	Any vendor wishing to implement their own L2TP extensions can use their own Vendor ID along with private Attribute values. Vendor-ID=0 means that the standard AVPs are used (2 bytes).
Attribute Type	A value with a unique interpretation across all AVPs defined under a given Vendor (2 bytes).
Attribute Value	This is the actual value as indicated by the Vendor ID and Attribute Type (2 bytes).

RADIUS-triggered tunnel/session setup without LNS renegotiation

Figure 346: RADIUS triggered tunnel/session setup without LNS renegotiation depicts the complete PPP session setup, using RADIUS authentication on both LAC and LNS. After the discovery phase (PADI/PADO/PADR/PADS) and LCP negotiation phase (LCP config_request/Ack), the LAC initiates the L2TP tunnel setup based on RADIUS authentication information (RADIUS Request/Accept) and includes the negotiated PPP user-LAC information (called LCP proxy information). The LNS replies directly with a successful CHAP authentication if it agrees with the received proxy information. IP negotiation (IPCP config_request/Ack) is handled between the user and the LNS, and the LAC has no IP knowledge of this PPP session.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 346: RADIUS triggered tunnel/session setup without LNS renegotiation

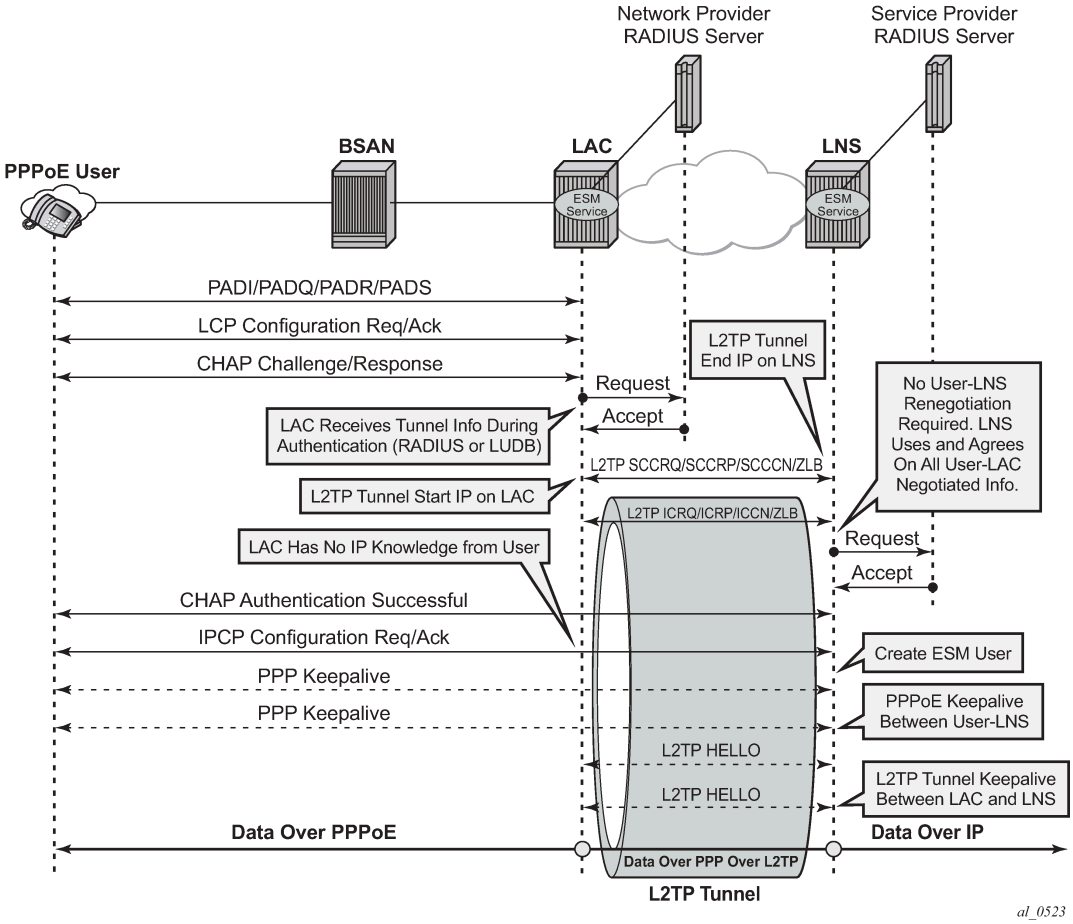

RADIUS-triggered tunnel/session setup with LNS renegotiation

Figure 347: RADIUS triggered tunnel/session setup with LNS renegotiation shows the scenario where the LNS does not agree with the received LCP proxy information and (re)starts the LCP phase (LCP config_request/Ack) directly with the PPP user. The rest of this scenario is the same as shown in Figure 346: RADIUS triggered tunnel/session setup without LNS renegotiation.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 347: RADIUS triggered tunnel/session setup with LNS renegotiation

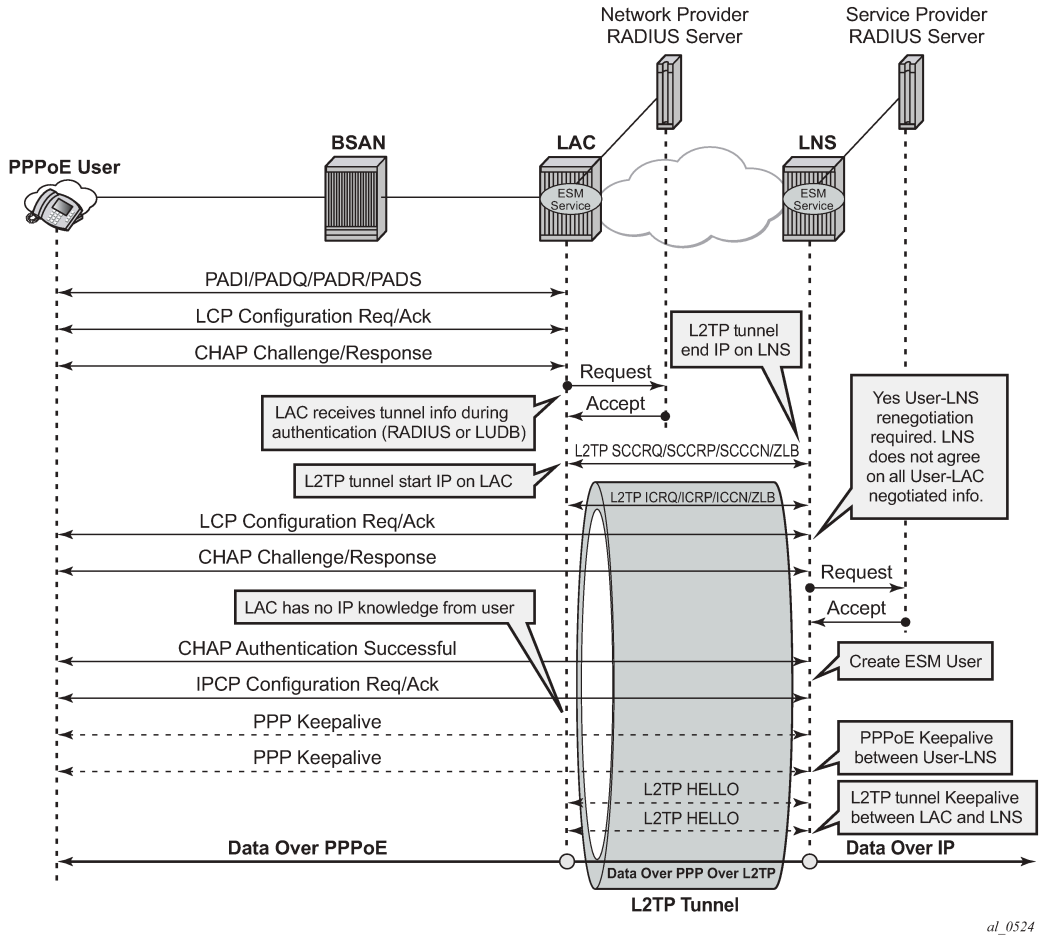

Running multiple PPP sessions over a single L2TP tunnel

Figure 348: Running multiple PPP sessions over a single L2TP tunnel shows multiple PPP sessions tunneled over a single L2TP Tunnel. The LAC encapsulates each PPP session with a different L2TP session-id (SID) but with the same L2TP tunnel-id (TID).

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 348: Running multiple PPP sessions over a single L2TP tunnel

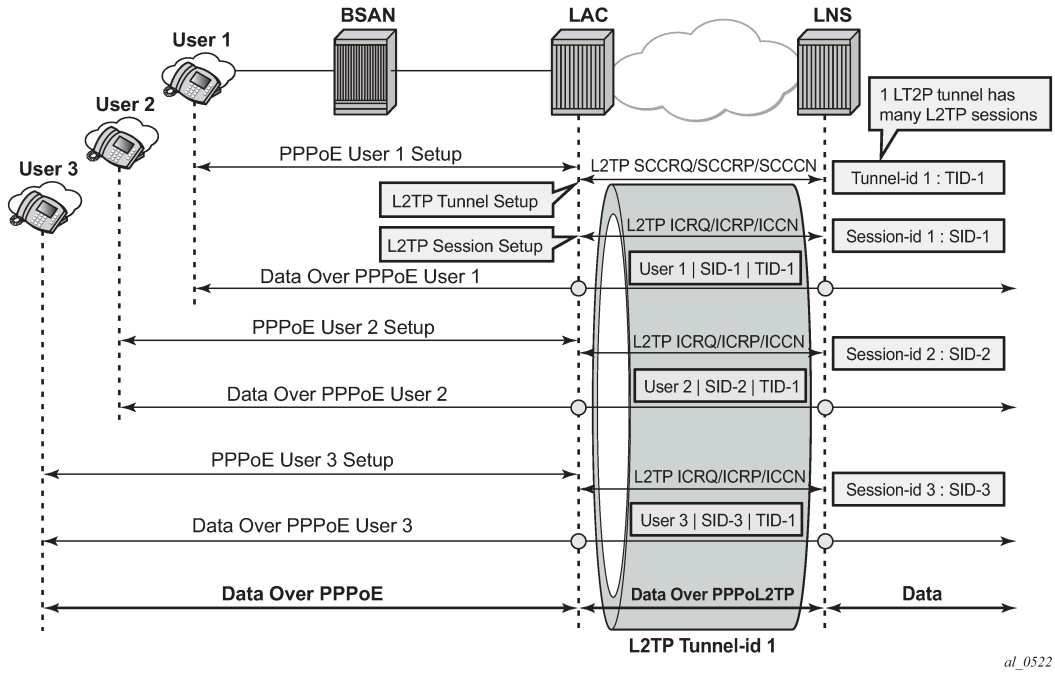

PPP user-initiated release/terminate

Figure 349: PPP user-initiated release/terminate shows the user initiated terminate_request tunneled by the LAC followed by the user initiated PADT terminated on the LAC. The LAC informs the LNS about the termination of the session via the L2TP CDN message. The L2TP tunnel can be optionally (idle-timeout) terminated via the L2TP StopCCN message.

Figure 349: PPP user-initiated release/terminate

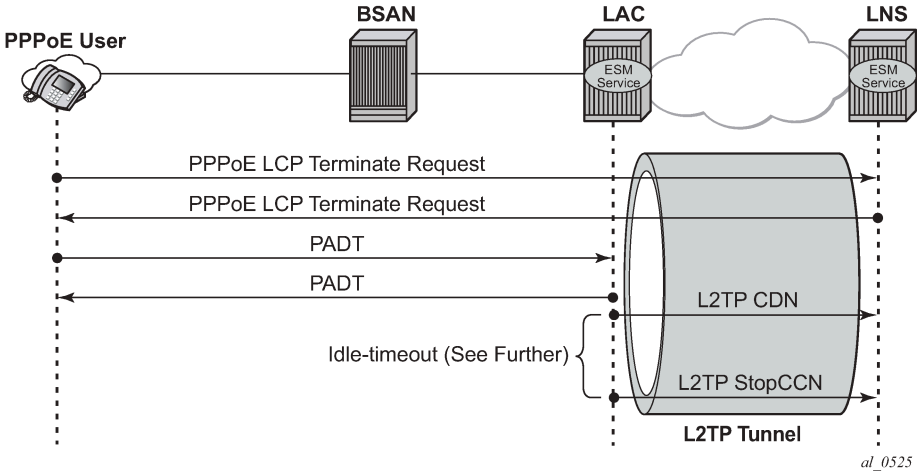

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

L2TP tunnel/session state diagram

Figure 350: L2TP tunnel and session state diagram gives an overview of the main L2TP tunnel and session states. An L2TP tunnel in the establishedIdle state is a tunnel without sessions. A **tools** command (see Advanced topics) can put an L2TP tunnel in a draining state (this prevents adding new sessions to the tunnel but leaves the current sessions intact) or in a drained state (moved from draining to drained when all sessions are terminated). The draining and drained state are not shown in the state diagram.

The L2TP tunnel setup occurs first with the triggers being: session activation, auto-establish, and a **tools start** command (see the Advanced topics section). An L2TP session setup trigger is always session based.

Figure 350: L2TP tunnel and session state diagram

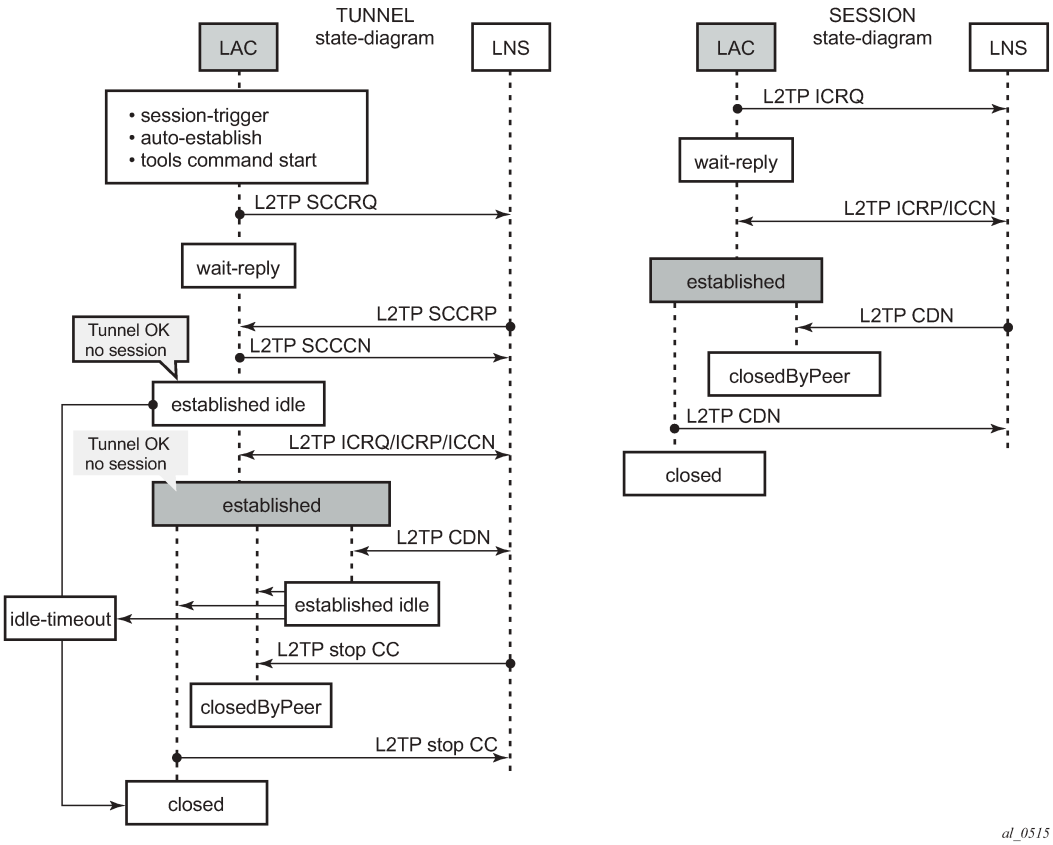

Configuration

Scenario 1: RADIUS-derived L2TP parameters

In the first scenario, the LAC receives an incoming connection and contacts the LAC RADIUS server. The RADIUS server retrieves the attributes for the user's domain (for example @wholesale.com) and passes the tunnel attributes to the LAC. Based on these RADIUS provided tunnel attributes, the LAC selects or initiates a new tunnel to the LTS or directly to the LNS. Once the tunnel is established, the LNS

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

authenticates the end user using its own RADIUS server. Configuring the LNS and the LTS are out of the scope of this example.

In a RADIUS driven L2TP setup, either all or some of the required L2TP attributes are returned via RADIUS. If the RADIUS server only returns the L2TP [67] Tunnel-Server-Endpoint attributes, then the L2TP tunnel/session is established using the 'l2tp node parameter values' for the other required L2TP parameters. The 'l2tp node parameters' are defined under the configure router/service l2tp hierarchy. If the RADIUS server does not return all of the L2TP attributes and the node values are not configured, then the system falls back to default settings for these L2TP parameters.

The standard and vendor specific [26-6572] L2TP RADIUS attributes are listed in the tables below, together with the corresponding l2tp node parameters and defaults.

Table 22: Generic L2TP RADIUS attributes

Attribute ID	Attribute name	Mandatory	CLI node parameter	Corresponding defaults	
64	Tunnel-Type	Y	-	-	-
65	Tunnel-Medium-Type	Y	-	-	-
66	Tunnel-Client-Endpoint: [0-31]	N	local-address	no local-address	system-ip
67	Tunnel-Server-Endpoint	N	-	-	-
69	Tunnel-Password	N	password	no password	-
82	Tunnel-Assignment-ID:0	N	-	-	default_radius_group
82	Tunnel-Assignment-ID: [1..31]	N	-	-	Unnamed
83	Tunnel-Preference	N	preference	no preference	50
90	Tunnel-Client-Auth-ID	N	local-name	no local-name	system-name
91	Tunnel-Server-Auth-ID	N	-	-	-

Table 23: Nokia specific L2TP RADIUS attributes

26-6527	Attribute name	Mandatory	CLI node parameter	Corresponding defaults	
-46	Alc-Tunnel-Group	N	-	-	-
-47	Alc-Tunnel-Algorithm	N	session-assign-method	no session-assign-method	existingFirst
-48	Alc-Tunnel-Max-Sessions:0	N	-	group-session-limit	131071

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

26-6527	Attribute name	Mandatory	CLI node parameter	Corresponding defaults	
-48	Alc-Tunnel-Max-Sessions:[1..31]	N	-	tunnel-session-limit	32767
-49	Alc-Tunnel-Idle-Timeout	N	idle-timeout	no idle-timeout	Infinite
-50	Alc-Tunnel-Hello-Interval	N	hello-interval	no hello-interval	300 sec
-51	Alc-Tunnel-Destruct-Timeout	N	destruct-timeout	no destruct-timeout	60 sec
-52	Alc-Tunnel-Max-Retries-Estab	N	max-retries-estab	no max-retries-estab	5
-53	Alc-Tunnel-Max-Retries-Not-Estab	N	max-retries-not-estab	no max-retries-not-estab	5
-54	Alc-Tunnel-AVP-Hiding	N	avp-hiding	no avp-hiding	Never
-97	Alc-Tunnel-Challenge	N	challenge	no challenge	Never
-104	Alc-Tunnel-Serv-Id	N	-	-	Base
-120	Alc-Tunnel-Rx-Window-Size	N	receive-window-size	no receive-window-size	64
-144	Alc-Tunnel-Acct-Policy	N	radius-accounting-policy	no radius-accounting-policy	-

Base router hosted LAC with single endpoint/single tunnel

Using the mandatory L2TP RADIUS attributes (see the following RADIUS user file) the LAC initiates an L2TP tunnel. The source address for the tunnel is the IPv4 address of a loopback interface in the Base router system (LAC tunnel endpoint).The destination for the tunnel is defined by the Tunnel-Server-Endpoint RADIUS attribute [67], and is also known as the peer tunnel LNS endpoint address.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 351: Base router hosted LAC with single endpoint/single tunnel

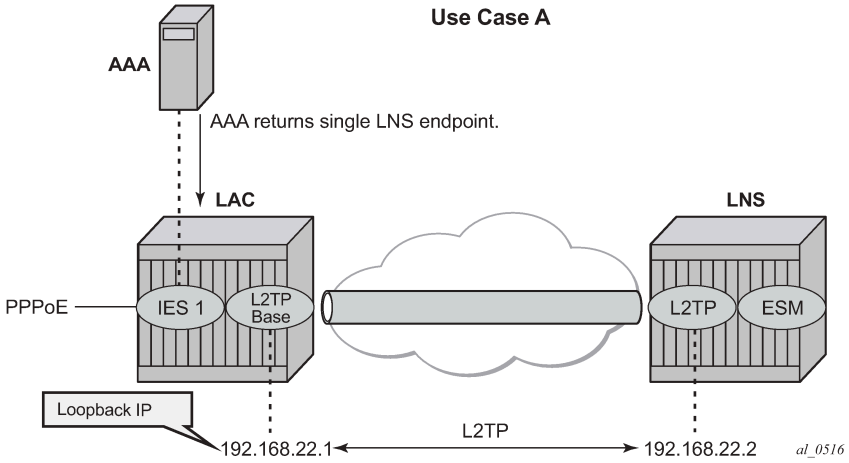

The PPPoE user terminates on IES service 1, sap 1/1/3:100, and is authenticated via RADIUS **authentication-policy** *radius-1* which provides wholesale/retail (L2TP) information.

```
configure {
    service {
        ies "ies-1" {
            admin-state enable
            service-id 1
            customer "1"
            subscriber-interface "sub-l2tp" {
                ipv4 {
                    unnumbered {
                        ip-int-name "system"
                    }
                }
                group-interface "grp-l2tp" {
                    radius-auth-policy "radius-1"
                    pppoe {
                        admin-state enable
                        session-limit 3
                        sap-session-limit 10
                    }
                    sap 1/1/3:100 {
                        sub-sla-mgmt {
                            admin-state enable
                            sub-ident-policy "all-subscribers"
                            subscriber-limit 1000
                        }
                    }
                }
            }
        }
    }
    commit
```

The excerpt from the FreeRADIUS users file below shows the attributes to be returned.

```
user1@wholesale.com     Cleartext-Password := "letmein", NAS-Identifier == "LAC"
                        Alc-Subsc-ID-Str = "%{User-name}",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Tunnel-Type:1 += L2TP,
                        Tunnel-Medium-Type:1 +=IP,
                        Tunnel-Server-Endpoint:1 += 192.168.22.2,
```

L2TP is enabled (no shutdown) in the related service instance.

The L2TP tunnel is set up in the base instance and not in a VRF because the attribute Alc-Tunnel-Serv-Id is not returned from RADIUS.

Missing L2TP parameters are taken from defaults defined in the router l2tp context.

```
configure router l2tp
       calling-number-format "%S %s"  # L2TP AVP 22 format
                                      # Default format 'system-name sap-id'
       ---snip---
       no local-name                  # default name equals system-name
       no max-retries-estab           # default value equals 5
       ---snip---
       no shutdown                    # enable L2TP
```

This scenario shows the PPPoE session termination (base IES service 1) and the L2TP tunnel setup in the base router instance.

Base router hosted LAC with multiple endpoints

Figure 352: Base router hosted LAC with multiple endpoints

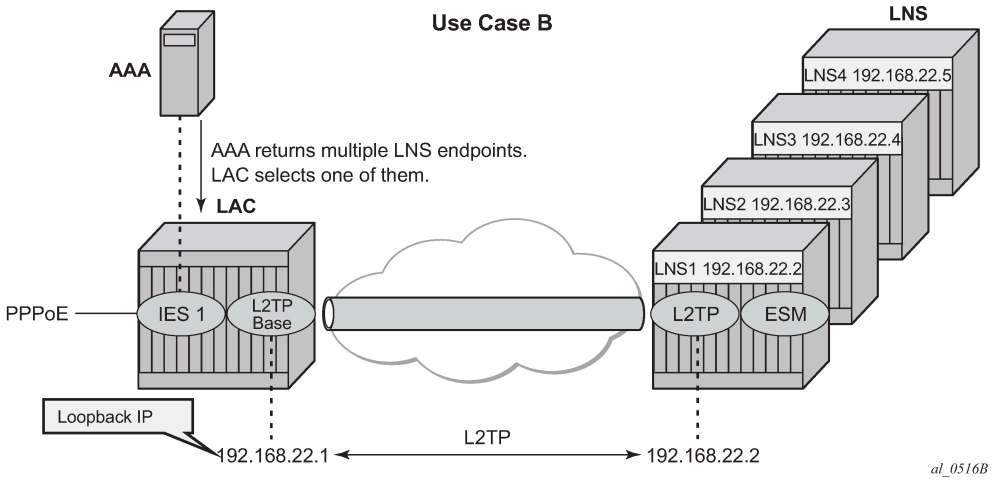

The following excerpt from the FreeRADIUS users file shows that user *user1@wholesale.com* has 4 possible endpoints (LNS), each with its own tunnel preference. The LAC selects one L2TP endpoint out of these 4 tunnel specifications according to the configured L2TP selection process. This use case uses weighted load balancing between RADIUS-tunnel-1 and RADIUS-tunnel-2. The L2TP tunnel selection process is out of the scope of this chapter.

```
user1@wholesale.com     Cleartext-Password := "letmein", NAS-Identifier == "LAC"
                        Alc-Subsc-ID-Str = "%{User-name}",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        Alc-SLA-Prof-Str = "sla-profile-1",
# group related info
                        Tunnel-Client-Endpoint:0 = 192.168.22.1,
                        Alc-Tunnel-Algorithm:0 = weighted-access,
                        Tunnel-Client-Auth-Id:0 = "lac-pe1",
                        Tunnel-Assignment-Id:0 = "RADIUS-group",
                        Alc-Tunnel-Max-Retries-Estab:0 = 2,
# tunnel-1 related info
                        Tunnel-Type:1 += L2TP,
                        Tunnel-Medium-Type:1 +=IP,
                        Tunnel-Server-Endpoint:1 += 192.168.22.2,
                        Tunnel-Assignment-Id:1 += "RADIUS-tunnel-1",
                        Tunnel-Preference:1 += 100,
# tunnel-2 related info
                        Tunnel-Type:2 += L2TP,
                        Tunnel-Medium-Type:2 +=IP,
                        Tunnel-Server-Endpoint:2 += 192.168.22.3,
                        Tunnel-Assignment-Id:2 += "RADIUS-tunnel-2",
                        Tunnel-Preference:2 += 100,
# tunnel-3 related info
                        Tunnel-Type:3 += L2TP,
                        Tunnel-Medium-Type:3 +=IP,
                        Tunnel-Server-Endpoint:3 += 192.168.22.4,
                        Tunnel-Assignment-Id:3 += "RADIUS-tunnel-3",

---snip---
```

This scenario shows the PPPoE session termination (base IES service 1) and the L2TP tunnel setup in the base router instance.

VRF hosted LAC

Figure 353: VRF hosted LAC shows the PPPoE session termination (base IES service 1) and the L2TP tunnel setup in a different router instance (VPRN 65536).

Figure 353: VRF hosted LAC

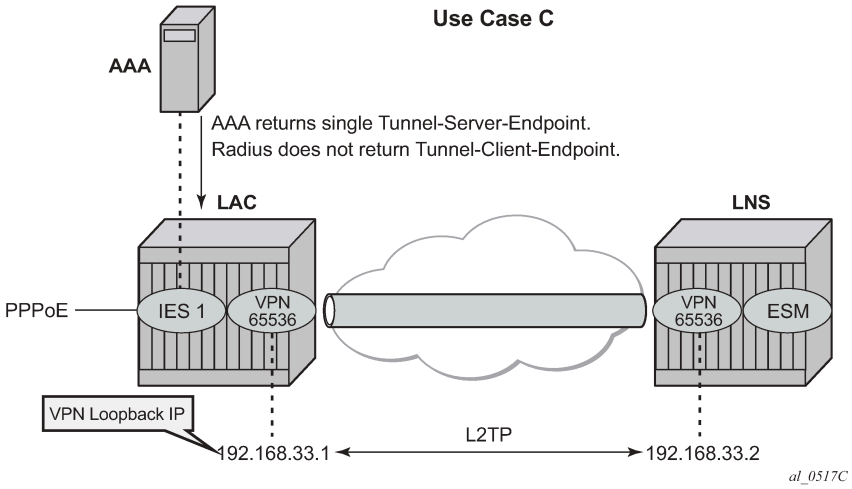

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Using the following L2TP RADIUS attributes, the LAC initiates an L2TP tunnel in VPRN 65536. The PPPoE session is still handled by IES service 1, which proves that both router instances can be different. (See use-case A for configuration details of IES service 1).

```
user1@wholesale.com     Cleartext-Password := "letmein", NAS-Identifier == "LAC"
                        Alc-Subsc-ID-Str = "%{User-name}",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-Tunnel-Serv-Id = 65536,
                        Tunnel-Client-Auth-Id:0 = "lac-pe1",
                        Tunnel-Assignment-Id:0 = "RADIUS-returned-TG",
                        Tunnel-Type:1 += L2TP,
                        Tunnel-Medium-Type:1 +=IP,
                        Tunnel-Server-Endpoint:1 += 192.168.33.2,
                        Tunnel-Assignment-Id:1 += "RADIUS-returned-TN",
```

If RADIUS does not return the L2TP source IP address (Tunnel-Client-Endpoint), then the IP address from the VPRN 65536 interface named 'system' is used as the L2TP source address. The tunnel setup fails if this system interface does not exist.

```
configure {
    service {
        vprn "vprn-65536" {
            admin-state enable
            customer "1"
            service-id 65536
            route-distinguisher "64496:65536"
            l2tp {
                admin-state enable
            }
            vrf-target {
                community "target:64496:65536"
            }
            interface "system" {
                loopback true
                ipv4 {
                    primary {
                        address 192.168.33.1
                        prefix-length 32
                    }
                }
            }
```

Scenario 2: Node-derived L2TP parameters

In the second scenario, the LAC receives the incoming connection and an 'L2TP tunnel-group-name' is assigned during LUDB or RADIUS authentication. This tunnel-group-name refers to the CLI preconfigured tunnel-group name context (**configure router** <*router-name*> **l2tp group** <*tunnel-group-name*>), which provides the context for all relevant tunnel attributes.

Based on these attributes, the LAC selects and initiates a tunnel to the LTS or directly to the LNS as in Scenario 1: RADIUS-derived L2TP parameters.

RADIUS returns L2TP tunnel group

In use case D, the L2TP tunnel-group-name is assigned during RADIUS authentication.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 354: RADIUS returns L2TP tunnel group

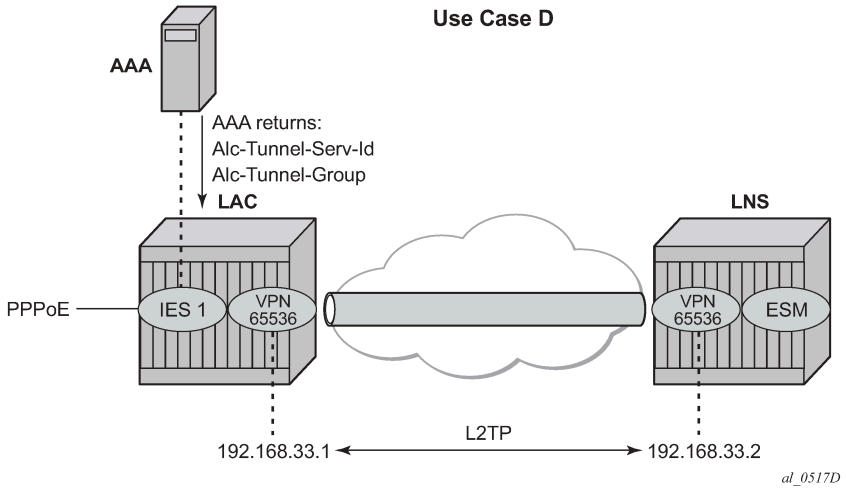


```
user1@wholesale.com     Cleartext-Password := "letmein", NAS-Identifier == "LAC"
                        Alc-Subsc-ID-Str = "%{User-name}",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-Tunnel-Serv-Id = 65536,
                        Alc-Tunnel-Group = "wholesale.com",
```

The L2TP tunnel is initiated from VPRN 65536 (Alc-Tunnel-Serv-Id) and all L2TP tunnel information is taken from the l2tp group wholesale.com hierarchy (Alc-Tunnel-Group) as defined on the node.

```
configure {
    service {
        vprn "vprn-65536" {
            admin-state enable
            customer "1"
            service-id 65536
            route-distinguisher "64496:65536"
            ---snip---
            l2tp {
                admin-state enable
                group "wholesale.com" {
                    admin-state enable
                    tunnel "wholesale.com" {
                        admin-state enable
                        auto-establish false
                        peer 192.168.33.2
                        local-address 192.168.33.1
                        local-name "lac-pe1"
                    }
                }
            }
            interface "system" {
                loopback true
                ipv4 {
                    primary {
                        address 192.168.33.1
                        prefix-length 32
                    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
            }
        }
    }
```

An L2TP tunnel is set up by either a PPP session-trigger, a **tools** command or by the l2tp group tunnel auto-establish parameter configuration. See the Advanced topics section for the non-session-triggered tunnel setup.

LUDB returns L2TP tunnel group

In use case E, the L2TP tunnel-group-name is assigned during LUDB authentication, so this essentially is a RADIUS-less scenario.

Figure 355: LUDB returns L2TP tunnel group

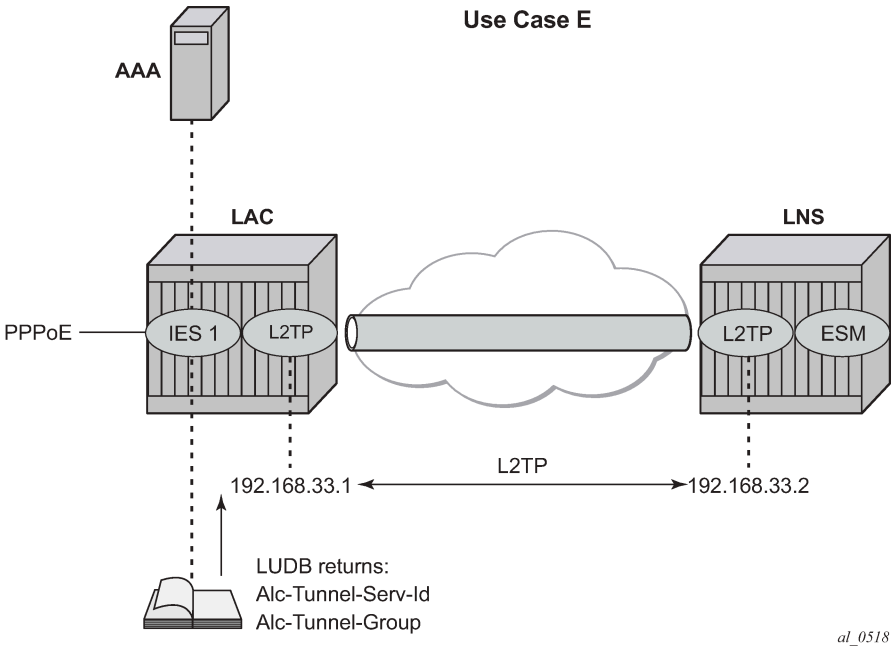

The PPPoE user enters on an IES service 1, sap 1/1/3:100, and is authenticated via the LUDB which provides L2TP wholesale/retail and ESM information. The PPPoE context refers to a local-user database *l2tp* to provide the subscriber authentication and the tunnel setup parameters, so no RADIUS is needed.

```
configure {
    service {
        ies "ies-1" {
            admin-state enable
            service-id 1
            customer "1"
            subscriber-interface "sub-l2tp" {
                ipv4 {
                    unnumbered {
                        ip-int-name "system"
                    }
                }
                group-interface "grp-l2tp" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    pppoe {
                        admin-state enable
                        session-limit 3
                        sap-session-limit 10
                        user-db "l2tp"
                    }
                    sap 1/1/3:100 {
                        sub-sla-mgmt {
                            admin-state enable
                            sub-ident-policy "all-subscribers"
                            subscriber-limit 1000
                        }
                    }
                }
            }
        }
    }
```

The referenced local user database *l2tp* configuration provides all of the required L2TP and ESM information.

```
configure {
    subscriber-mgmt {
        local-user-db "l2tp" {
            admin-state enable
            ppp {
                match-list [user-name]
                host "wholesale.com" {
                    admin-state enable
                    host-identification {
                        user-name {
                            name "wholesale.com"
                            format domain-only
                        }
                    }
                    identification {
                        sla-profile-string "sla-profile-1"
                        sub-profile-string "sub-profile-1"
                        subscriber-id "user@wholesale.com"
                    }
                    l2tp {
                        group {
                            name "wholesale.com"
                            service-id 65536
                        }
                    }
                    password {
                        ignore
                    }
                }
            }
        }
    }
```

Operation and troubleshooting

The subsequent sections explain how the use cases A to E described in the configuration section are verified using show, debug, and tools commands.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The standard router debugging tools can be used to monitor and troubleshoot the L2TP tunnel and session setup.

Useful show commands are:

```
show service id <service-id> ppp session [detail]
show router l2tp tunnel [detail]
show router l2tp session [detail]
show router l2tp peer [ip-address]
```

To debug and show PPPoE packets:

```
debug service id <service-id> ppp packet mode egr-ingr-and-dropped
debug service id <service-id> ppp packet detail-level medium
```

To debug and show RADIUS authentication:

```
debug router radius packet-type authentication
```

To debug and show LUDB authentication:

```
debug subscriber-mgmt local-user-db <local-user-db-name> detail all
```

To debug and show the LAC tunnel selection process and L2TP state machine:

```
debug router l2tp event lac-session-setup
debug router l2tp event finite-state-machine
```

To debug and show the L2TP tunnel and session setup:

```
debug router l2tp packet direction both
debug router l2tp packet detail-level high
```

Understanding the L2TP debug output

The following L2TP ICRQ message (**debug router l2tp packet**) is used to explain how the displayed debug output should be interpreted. See Recap of the L2TPv2 protocol -L2TP header and AVP layout for more details.

```
19 2019/05/28 11:15:51.766 CEST MINOR: DEBUG #2001 Base L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.0.2.1:1701 -> 192.168.22.2:1701
tunnel 4734 session 0, ns 2 nr 1, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        IncomingCallRequest(10)
    AVP AssignedSessionId(0,14), flags: mandatory, reserved=0
        8018
    AVP CallSerialNumber(0,15), flags: mandatory, reserved=0
        256975
    AVP CallingNumber(0,22), flags: mandatory, reserved=0
        "LAC 1/1/3:100"
    AVP AgentCircuitId(3561,1), flags:, reserved=0
        "circuit0"
    AVP AgentRemoteId(3561,2), flags:, reserved=0
        "remote0"
    AVP ActDataRateUp(3561,129), flags:, reserved=0
        2000000
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    AVP ActDataRateDown(3561,130), flags:, reserved=0
        4000000"
```

- L2TP(v2, ctrl, egress): UDP 192.0.2.1:1701 -> 192.168.22.2:1701
  - version: v2
  - type field (T-bit): control message (ctrl)
  - 192.0.2.1:1701 -> 192.168.22.2:1701
    - 192.0.2.1:1701 - source tunnel-end-point:source udp port
    - 192.168.22.2:1701 - destination tunnel-end-point:destination udp port
- tunnel 13008 session 0, ns 2 nr 1, flags:, reserved=0
  - tunnel-id: 13008
  - session-id: 0
  - ns:2
  - nr:1
  - flags: 0 (refers to T/L/S/O/P bits L2TP header)
  - reserved field:0
- AVP CallingNumber(0,22), flags: mandatory, reserved=0
  - AVP MessageType(0,22): "LAC 1/1/3:100"
    - Vendor-id: 0 - Standard Attribute
    - Attribute Type: 22 – Calling Number AVP
    - Attribute Value: "LAC 1/1/3:100"

Scenario 1: RADIUS-derived L2TP parameters

Base router hosted LAC with single endpoint/single tunnel

The **debug service id** <*service-id*> **ppp packet mode egr-ingr-and-dropped** command shows the PPPoE packet exchange. The following PADI packet shows the service, SAP, and received PPPoE tags. The received PPPoE DSL forum tags are by default copied during the LAC L2TP tunnel setup into the Incoming Call Request (ICRQ) DSL Forum AVP's (RFC 5515).

```
1 2019/05/28 11:15:51.724 CEST MINOR: DEBUG #2001 Base PPPoE
"PPPoE: RX Packet
   IES 1, SAP 1/1/3:100

   DMAC: ff:ff:ff:ff:ff:ff
   SMAC: 00:00:00:00:01:01
   Ether Type: 0x8863 (Discovery)

   PPPoE Header:
   Version: 1                Type      : 1
   Code   : 0x09 (PADI)      Session-Id: 0x0000 (0)
   Length : 48

   PPPoE Tags:
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
   [0x0101] Service-Name: ""
   [0x0103] Host-Uniq: len = 1, value = 31
   [0x0105] Vendor-Specific: vendor-id = 0x0de9 (ADSL Forum)
      [0x01] Agent-Circuit-Id: "circuit0"
      [0x02] Agent-Remote-Id: "remote0"
      [0x81] Actual-Upstream: 2000
      [0x82] Actual-Downstream: 4000
"
```

The **debug router radius packet-type authentication** command shows the actual authentication parameters returned by RADIUS. This example returns the minimum set of L2TP related RADIUS attributes.

```
12 2019/05/28 11:15:51.762 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Access-Accept(2) id 242 len 89 from 172.16.1.11:1812 vrid 1 pol rsp-radius-1
    VSA [26] 15 Nokia(6527)
      SUBSC PROF STR [12] 13 sub-profile-1
    VSA [26] 15 Nokia(6527)
      SLA PROF STR [13] 13 sla-profile-1
    TUNNEL TYPE [64] 4 1 L2TP(3)
    TUNNEL MEDIUM TYPE [65] 4 1 IPv4(1)
    TUNNEL SERVER ENDPOINT [67] 13 1 192.168.22.2
"
```

The **debug router l2tp event lac-session-setup** command shows the LAC tunnel selection for this example. An L2TP group-name '*default_radius_group*' with tunnel-name '*unnamed*' is created in this case, because RADIUS did not return an explicit group and tunnel name.

```
13 2019/05/28 11:15:51.763 CEST MINOR: DEBUG #2001 Base PPPoE 255217->L2TP
"PPPoE 255217->L2TP: UDP 192.0.2.1:1701 -> 192.168.22.2:1701
preference 50 tunnel default_radius_group:unnamed
    request to open new tunnel 12082"
```

```
14 2019/05/28 11:15:51.763 CEST MINOR: DEBUG #2001 Base PPPoE 255217->L2TP
"PPPoE 255217->L2TP: UDP 192.0.2.1:1701 -> 192.168.22.2:1701
preference 50 tunnel default_radius_group:unnamed
    create session 791813970"
```

The **debug router l2tp packet detail-level** command shows the L2TP tunnel and session setup for this example.

For the tunnel setup, the LAC sends a Start-Control-Connection-Request (SCCRQ) containing the assigned tunnel-id (no tunnel authentication in the example). The tunnel is now in a wait-reply state.

```
15 2019/05/28 11:15:51.763 CEST MINOR: DEBUG #2001 Base L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.0.2.1:1701 -> 192.168.22.2:1701
tunnel 0 session 0, ns 0 nr 0, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        StartControlConnectionRequest(1)
    AVP ProtocolVersion(0,2), flags: mandatory, reserved=0
        version=1, revision=0
    AVP HostName(0,7), flags: mandatory, reserved=0
        "lac-pe1"
    AVP WindowSize(0,10), flags: mandatory, reserved=0
        64
    AVP FramingCapabilities(0,3), flags: mandatory, reserved=0
        sync=no, async=no
    AVP BearerCapabilities(0,4), flags: mandatory, reserved=0
        digital=yes, analogue=no
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    AVP FirmwareRevision(0,6), flags:, reserved=0
        4869
    AVP VendorName(0,8), flags:, reserved=0
        "Nokia"
    AVP AssignedTunnelId(0,9), flags: mandatory, reserved=0
        12082"
```

The LNS can bring up the tunnel, so the LNS replies with a Start-Control-Connection-Reply (SCCRP) including the assigned tunnel-id.

```
16 2019/05/28 11:15:51.765 CEST MINOR: DEBUG #2001 Base L2TP(v2, ctrl, ingress)
"L2TP(v2, ctrl, ingress): UDP 192.168.22.2:1701 -> 192.0.2.1:1701
tunnel 12082 session 0, ns 0 nr 1, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        StartControlConnectionReply(2)
    AVP ProtocolVersion(0,2), flags: mandatory, reserved=0
        version=1, revision=0
    AVP HostName(0,7), flags: mandatory, reserved=0
        "lns-pe2"
    AVP WindowSize(0,10), flags: mandatory, reserved=0
        64
    AVP FramingCapabilities(0,3), flags: mandatory, reserved=0
        sync=no, async=no
    AVP BearerCapabilities(0,4), flags: mandatory, reserved=0
        digital=yes, analogue=no
    AVP FirmwareRevision(0,6), flags:, reserved=0
        4869
    AVP VendorName(0,8), flags:, reserved=0
        "Nokia"
    AVP AssignedTunnelId(0,9), flags: mandatory, reserved=0
        4734"
```

As the last step in the tunnel setup phase, the LAC responds with a Start-Control-Connection-Connected (SCCCN) message. After an LNS ZLB acknowledgment, the tunnel is in the establishedIdle state.

```
17 2019/05/28 11:15:51.765 CEST MINOR: DEBUG #2001 Base L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.0.2.1:1701 -> 192.168.22.2:1701
tunnel 4734 session 0, ns 1 nr 1, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        StartControlConnectionConnected(3)"
```

Once the tunnel exists the session setup starts, a three-way exchange for session establishment within the tunnel is performed. The LAC sends an Incoming-Call-Request (ICRQ) with the parameter information for the session. The session is now in the wait-reply state.

```
19 2019/05/28 11:15:51.766 CEST MINOR: DEBUG #2001 Base L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.0.2.1:1701 -> 192.168.22.2:1701
tunnel 4734 session 0, ns 2 nr 1, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        IncomingCallRequest(10)
    AVP AssignedSessionId(0,14), flags: mandatory, reserved=0
        8018
    AVP CallSerialNumber(0,15), flags: mandatory, reserved=0
        256975
    AVP CallingNumber(0,22), flags: mandatory, reserved=0
        "LAC 1/1/3:100"
    AVP AgentCircuitId(3561,1), flags:, reserved=0
        "circuit0"
    AVP AgentRemoteId(3561,2), flags:, reserved=0
        "remote0"
    AVP ActDataRateUp(3561,129), flags:, reserved=0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        2000000
    AVP ActDataRateDown(3561,130), flags:, reserved=0
        4000000"
```

The LNS then sends an Incoming-Call-Reply (ICRP) that contains the assigned session-id. The session is now in the connect state.

```
21 2019/05/28 11:15:51.768 CEST MINOR: DEBUG #2001 Base L2TP(v2, ctrl, ingress)
"L2TP(v2, ctrl, ingress): UDP 192.168.22.2:1701 -> 192.0.2.1:1701
tunnel 12082 session 8018, ns 1 nr 3, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        IncomingCallReply(11)
    AVP AssignedSessionId(0,14), flags: mandatory, reserved=0
        24477"
```

Finally the LAC sends an Incoming Call Connected (ICCN) and provides the LNS with additional information from the user initiated session. This information includes the LCP information from the negotiation that the LAC and remote user performed. This information is used by the LNS to decide whether to start LCP re-negotiation and/or Authentication re-negotiation with the PPP user or not. After an LNS ZLB acknowledgment, the session is in the established state.

```
24 2019/05/28 11:15:51.769 CEST MINOR: DEBUG #2001 Base L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.0.2.1:1701 -> 192.168.22.2:1701
tunnel 4734 session 24477, ns 3 nr 2, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        IncomingCallConnected(12)
    AVP FramingType(0,19), flags: mandatory, reserved=0
        sync=no, async=no
    AVP TxConnectSpeed(0,24), flags: mandatory, reserved=0
        4294967295
    AVP InitialRxLcpConfReq(0,26), flags:, reserved=0
         01 04 05 d4
         [1] MRU: 1492
    AVP LastTxLcpConfReq(0,27), flags:, reserved=0
         01 04 05 d4 03 05 c2 23 05 05 06 29 c0 85 ab
         [1] MRU: 1492
         [3] Authentication-Protocol: 0xc223 (CHAP), Algorithm = 5 (MD5)
         [5] Magic-Number: 0x29c085ab
    AVP LastRxLcpConfReq(0,28), flags:, reserved=0
         01 04 05 d4
         [1] MRU: 1492
    AVP ProxyAuthenType(0,29), flags:, reserved=0
        chap(2)
    AVP ProxyAuthenName(0,30), flags:, reserved=0
         "user1@wholesale.com"
    AVP ProxyAuthenChallenge(0,31), flags:, reserved=0
         4c 1c 3d e9 8f 11 7a 09 e0 2a 4e 9e d6 d4 c6 78
         eb d2 bc e0 72 27 41 a2 77 61 67 22 95 2b 1a 61
         c6 57 4b
    AVP ProxyAuthenId(0,32), flags:, reserved=0
        id=1, reserved=0
    AVP ProxyAuthenResponse(0,33), flags:, reserved=0
         65 2f 0f 3e 28 24 5c 8a ff 63 66 98 93 29 97 3d
    AVP RxConnectSpeed(0,38), flags:, reserved=0
        4294967295"
```

The operational PPPoE session information for the IES 1 (base router) instance is as follows.

```
[]
A:admin@LAC# show service id 1 ppp session
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
PPP sessions for service 1
===============================================================================
User-Name
  Descr.
          Up Time       Type  Termination     IP/L2TP-Id/Interface-Id MC-Stdby
-------------------------------------------------------------------------------
user1@wholesale.com
  svc:1 sap:1/1/3:100 mac:00:00:00:00:01:01 sid:1
          0d 00:01:21   oE    lac             791813970
-------------------------------------------------------------------------------
No. of PPP sessions: 1
===============================================================================
```

The operational tunnel information in the base instance shows that the tunnel is established.

```
[]
A:admin@LAC# show router l2tp tunnel
===============================================================================
Conn ID    Loc-Tu-ID Rem-Tu-ID State              Blacklist-state   Ses Active
  Group                                                             Ses Total
    Assignment
-------------------------------------------------------------------------------
791805952  12082     4734      established        not-blacklisted   1
  default_radius_group                                              1
    unnamed
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================
```

Detailed operational tunnel information is obtained using following command.

```
[]
A:admin@LAC# show router l2tp tunnel tunnel-id 12082 detail

===============================================================================
L2TP Tunnel Status
===============================================================================

Connection ID: 791805952
Protocol     : v2
State        : established
IP           : 192.0.2.1
UDP          : 1701
Peer IP      : 192.168.22.2
Peer UDP     : 1701
Tx dst-IP    : 192.168.22.2
Tx dst-UDP   : 1701
Rx src-IP    : 192.168.22.2
Rx src-UDP   : 1701
Name         : lac-pe1
Remote Name  : lns-pe2
Assignment ID: unnamed
Group Name   : default_radius_group
Acct. Policy : N/A
Error Message: N/A

                                       Remote Conn ID    : 310247424
Tunnel ID         : 12082              Remote Tunnel ID  : 4734
Preference        : 50                 Receive Window    : 64
Hello Interval (s): 300                AVP Hiding        : never
Idle TO (s)       : infinite           Destruct TO (s)   : 60
Max Retr Estab    : 5                  Max Retr Not Estab: 5
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Cfg'd Sess Limit  : unlimited           Oper Session Limit: 32767
Transport Type    : udpIp               Challenge         : never
Time Started      : 05/28/2019 11:15:52 Time Idle         : N/A
Time Established  : 05/28/2019 11:15:52 Time Closed       : N/A
Stop CCN Result   : noError             General Error     : noError
Blacklist-state   : not-blacklisted
Set Dont Fragment : true

Failover
State             : not-recoverable
Recovery Conn ID  : N/A
Recovery state    : not-applicable
Recovered Conn ID : N/A
Recovery method   : mcs
Track SRRP        : (Not specified)
Ctrl msg behavior : handle
Recovery time (ms)
Requested         : N/A
Peer              : N/A
-------------------------------------------------------------------------------
===============================================================================
```

The operational L2TP session information shows the L2TP session is established.

```
[]
A:admin@LAC# show router l2tp session

===============================================================================
L2TP Session Summary
===============================================================================
ID                  Control Conn ID     Tunnel-ID   Session-ID  State
-------------------------------------------------------------------------------
791813970           791805952           12082       8018        established
-------------------------------------------------------------------------------
No. of sessions: 1
===============================================================================
```

For detailed operational L2TP session information use the following command.

```
[]
A:admin@LAC# show router l2tp session session-id 8018 detail

===============================================================================
L2TP Session 791813970
===============================================================================
Connection ID: 791813970
State        : established
Tunnel Group : default_radius_group
Assignment ID: unnamed
Error Message: N/A

Control Conn ID   : 791805952           Rem Cntrl Conn ID : 310247424
Tunnel ID         : 12082               Remote Tunnel ID  : 4734
Session ID        : 8018                Remote Session ID : 24477
PW Type           : ppp                 Remote Conn ID    : 310271901
Time Started      : 05/28/2019 11:15:52
Time Established  : 05/28/2019 11:15:52 Time Closed       : N/A
CDN Result        : noError             General Error     : noError
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
No. of sessions: 1
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Base router hosted LAC with multiple endpoints

The **debug router radius packet-type authentication** command shows the actual RADIUS authentication parameters returned. This example returns multiple tunnel endpoints from which the LAC selects one. This example uses weighted load balancing. (The L2TP tunnel selection process is out of the scope of this example).

```
12 2019/05/28 15:04:49.128 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Access-Accept(2) id 235 len 225 from 172.16.1.11:1812 vrid 1 pol rsp-radius-1
    VSA [26] 15 Nokia(6527)
      SUBSC PROF STR [12] 13 sub-profile-1
    VSA [26] 15 Nokia(6527)
      SLA PROF STR [13] 13 sla-profile-1
    TUNNEL CLIENT ENDPOINT [66] 12 192.168.22.1
    VSA [26] 6 Nokia(6527)
      TUNNEL ALGORITHM [47] 4 weighted access(1)
    TUNNEL CLIENT AUTH ID [90] 7 lac-pe1
    TUNNEL ASSIGNMENT ID [82] 12 RADIUS-group
    VSA [26] 6 Nokia(6527)
      TUNNEL MAX RETRIES ESTAB [52] 4 0 2
    TUNNEL TYPE [64] 4 1 L2TP(3)
    TUNNEL MEDIUM TYPE [65] 4 1 IPv4(1)
    TUNNEL SERVER ENDPOINT [67] 13 1 192.168.22.2
    TUNNEL ASSIGNMENT ID [82] 16 1 RADIUS-tunnel-1
    TUNNEL PREFERENCE [83] 4 1 100
    TUNNEL TYPE [64] 4 2 L2TP(3)
    TUNNEL MEDIUM TYPE [65] 4 2 IPv4(1)
    TUNNEL SERVER ENDPOINT [67] 13 2 192.168.22.3
    TUNNEL ASSIGNMENT ID [82] 16 2 RADIUS-tunnel-2
    TUNNEL PREFERENCE [83] 4 2 100
"
```

The **debug router l2tp event lac-session-setup** command shows the LAC tunnel LNS2-T2 is selected for this example.

```
13 2019/05/28 15:04:49.129 CEST MINOR: DEBUG #2001 Base PPPoE 263408->L2TP
"PPPoE 263408->L2TP: UDP 192.168.22.1:1701 -> 192.168.22.2:1701
preference 100 tunnel RADIUS-group:RADIUS-tunnel-1
    request to open new tunnel 5288"
```

```
14 2019/05/28 15:04:49.129 CEST MINOR: DEBUG #2001 Base PPPoE 263408->L2TP
"PPPoE 263408->L2TP: UDP 192.168.22.1:1701 -> 192.168.22.2:1701
preference 100 tunnel RADIUS-group:RADIUS-tunnel-1
    create session 346586242"
```

The operational PPPoE session information in IES 1/base instance is shown as follows.

```
[]
A:admin@LAC# show service id 1 ppp session

===============================================================================
PPP sessions for service 1
===============================================================================
User-Name
  Descr.
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
           Up Time       Type  Termination     IP/L2TP-Id/Interface-Id MC-Stdby
-------------------------------------------------------------------------------
user1@wholesale.com
  svc:1 sap:1/1/3:100 mac:00:00:00:00:01:01 sid:1
           0d 00:00:56   oE    lac             346586242
-------------------------------------------------------------------------------
No. of PPP sessions: 1
===============================================================================
```

The operational L2TP tunnel information (base instance) is shown below.

```
[]
A:admin@LAC# show router l2tp tunnel
===============================================================================
Conn ID    Loc-Tu-ID Rem-Tu-ID State             Blacklist-state   Ses Active
  Group                                                            Ses Total
    Assignment
-------------------------------------------------------------------------------
346554368  5288      8368      established       not-blacklisted   1
  RADIUS-group                                                     1
    RADIUS-tunnel-1
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================
```

Operational session information (base instance) shows the session is in the established state.

```
[]
A:admin@LAC# show router l2tp session

===============================================================================
L2TP Session Summary
===============================================================================
ID                 Control Conn ID     Tunnel-ID   Session-ID  State
-------------------------------------------------------------------------------
346586242          346554368           5288        31874       established
-------------------------------------------------------------------------------
No. of sessions: 1
===============================================================================
```

The L2TP endpoint/peer information shows there are two tunnels for tunnel endpoint 192.168.22.2.

```
[]
A:admin@LAC# show router l2tp peer
===============================================================================
L2TP Peers
===============================================================================
Peer IP                                          Port  Tun Active Ses Active
                                  Drain Reachability Tun Total  Ses Total
-------------------------------------------------------------------------------
192.168.22.2                                     1701  1          1
                                                       1          1
192.168.22.3                                     1701  0          0
                                                       0          0
-------------------------------------------------------------------------------
No. of peers: 2
===============================================================================
```

The following command gives a system overview of subscriber session related data. This system overview shows the current and peak values per session type (local PTA, LAC, LTS, LNS) and an overview of the

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

number of originated or terminated L2TP tunnels. Peak values can be cleared via the **clear subscriber-mgmt peakvalue-stats** command.

```
[]
A:admin@LAC# show subscriber-mgmt statistics system statistics-id session

===============================================================================
Subscriber Management Statistics for System
===============================================================================
       Type                                Current     Peak      Peak Timestamp
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
PPP Session Statistics
-------------------------------------------------------------------------------
Local  PPP Sessions     - PPPoE                 0        0
       PPP Sessions     - PPPoEoA               0        0
       PPP Sessions     - PPPoA                 0        0
       PPP Sessions     - L2TP (LNS)            0        0
-------------------------------------------------------------------------------
LAC    PPP Sessions     - PPPoE                 1        2 05/28/2019 13:16:38
       PPP Sessions     - PPPoEoA               0        0
       PPP Sessions     - PPPoA                 0        0
       PPP Sessions     - L2TP (LTS)            0        0
-------------------------------------------------------------------------------
Total  PPP Sessions     - established           1        2 05/28/2019 13:16:38
       PPP Sessions     - in setup              0        1 05/28/2019 15:04:49
       PPP Sessions     - local                 0        0
       PPP Sessions     - LAC                   1        2 05/28/2019 13:16:38
-------------------------------------------------------------------------------
L2TP   L2TP Tunnels     - originator            1        2 05/28/2019 15:02:39
       L2TP Tunnels     - receiver              0        0
       Total L2TP Tunnels                       1        2 05/28/2019 15:02:39
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
IPOE Session Statistics
-------------------------------------------------------------------------------
Total  IPOE Sessions    - established           0        0
       IPOE Sessions    - in setup              0        0
-------------------------------------------------------------------------------
===============================================================================
Peak values last reset at : n/a
```

VRF hosted LAC

This example returns VPRN 65536 as the L2TP service instance [26-6527-104 Alc-Tunnel-Serv-Id]. The VPRN 65536 interface system address is used as the L2TP source address since the attribute Tunnel-Client-Endpoint is not returned.

The IP address 192.168.33.1 (Tunnel-Server-Endpoint) needs to be routable in VRF 65536 over a SAP or to a remote PE. This example uses BGP/MPLS IP Virtual Private Networks (VPNs) (RFC 4364) to access the remote PE.

```
[]
A:admin@LAC# show router 65536 route-table
===============================================================================
Route Table (Service: 65536)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
---snip---
192.168.33.1/32                          Local   Local     23h44m06s  0
       system                                                    0
192.168.33.2/32                          Remote  BGP VPN   05h39m09s  170
       192.0.2.2 (tunneled)                                      10
---snip---
-------------------------------------------------------------------------------
No. of Routes: 4
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

Operational PPPoE session information for IES 1 (base instance) is shown using following command.

```
[]
A:admin@LAC# show service id 1 ppp session

===============================================================================
PPP sessions for service 1
===============================================================================
User-Name
  Descr.
          Up Time       Type  Termination     IP/L2TP-Id/Interface-Id MC-Stdby
-------------------------------------------------------------------------------
user1@wholesale.com
  svc:1 sap:1/1/3:100 mac:00:00:00:00:01:01 sid:1
          0d 00:03:55   oE    lac             643922850
-------------------------------------------------------------------------------
No. of PPP sessions: 1
===============================================================================
```

Operational tunnel information for VPRN 65536 is displayed as follows.

```
[]
A:admin@LAC# show router 65536 l2tp tunnel
===============================================================================
Conn ID    Loc-Tu-ID Rem-Tu-ID State              Blacklist-state   Ses Active
  Group                                                            Ses Total
    Assignment
-------------------------------------------------------------------------------
643891200  9825      9819      established        not-blacklisted   1
  RADIUS-returned-TG                                               1
    RADIUS-returned-TN
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================
```

Operational session information for VPRN 65536 is displayed using following command, and shows that the session is established.

```
[]
A:admin@LAC# show router 65536 l2tp session

===============================================================================
L2TP Session Summary
===============================================================================
ID                  Control Conn ID     Tunnel-ID   Session-ID  State
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
643922850           643891200           9825        31650       established
-------------------------------------------------------------------------------
No. of sessions: 1
===============================================================================
```

Scenario 2: Node-derived L2TP parameters

RADIUS returns L2TP group

This example returns VPRN 65536 as the L2TP service instance [26-6527-104] Alc-Tunnel-Serv-Id and an l2tp group-name wholesale.com [26-6527-46] Alc-Tunnel-Group.

```
12 2019/05/28 15:26:19.861 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Access-Accept(2) id 237 len 95 from 172.16.1.11:1812 vrid 1 pol rsp-radius-1
    VSA [26] 15 Nokia(6527)
      SUBSC PROF STR [12] 13 sub-profile-1
    VSA [26] 15 Nokia(6527)
      SLA PROF STR [13] 13 sla-profile-1
    VSA [26] 6 Nokia(6527)
      TUNNEL SERVICE ID [104] 4 65536
    VSA [26] 15 Nokia(6527)
      TUNNEL GROUP [46] 13 wholesale.com

"
```

For operational PPPoE session information in IES 1/base instance, use following command.

```
[]
A:admin@LAC# show service id 1 ppp session

===============================================================================
PPP sessions for service 1
===============================================================================
User-Name
  Descr.
          Up Time       Type  Termination     IP/L2TP-Id/Interface-Id MC-Stdby
-------------------------------------------------------------------------------
user1@wholesale.com
  svc:1 sap:1/1/3:100 mac:00:00:00:00:01:01 sid:1
          0d 00:02:38   oE    lac             181162829
-------------------------------------------------------------------------------
No. of PPP sessions: 1
===============================================================================
```

Operational tunnel information for VPRN 65536 shows the tunnel is in the established state.

```
[]
A:admin@LAC# show router 65536 l2tp tunnel
===============================================================================
Conn ID    Loc-Tu-ID Rem-Tu-ID State             Blacklist-state   Ses Active
  Group                                                            Ses Total
    Assignment
-------------------------------------------------------------------------------
181141504  2764      38        established       not-blacklisted   1
  wholesale.com                                                    1
    wholesale.com
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================
```

The operational session information for VPRN 65536 shows the session is in the established state.

```
[]
A:admin@LAC# show router 65536 l2tp session

===============================================================================
L2TP Session Summary
===============================================================================
ID                  Control Conn ID     Tunnel-ID   Session-ID  State
-------------------------------------------------------------------------------
181162829           181141504           2764        21325       established
-------------------------------------------------------------------------------
No. of sessions: 1
===============================================================================
```

LUDB returns L2TP group

This example returns VPRN 65536 as the L2TP service instance and l2tp group-name wholesale.com (LUDB l2tp group "wholesale.com" service-id 65536).

The **debug subscriber-mgmt local-user-db l2tp detail all** command shows the LUDB authentication access (The returned parameter details are not shown).

```
11 2019/05/28 15:32:48.859 CEST MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
  user-name:
    original:  user1@wholesale.com
    masked:    user1@wholesale.com

  Host wholesale.com found in user data base l2tp"
```

To show the operational data from LUDB *l2tp*, use the following command.

```
[]
A:admin@LAC# show subscriber-mgmt local-user-db "l2tp" ppp-host "wholesale.com"
                                        | match N/A invert-match
                                        | match none invert-match
===============================================================================
PPP Host "wholesale.com"
===============================================================================
Admin State          : Up
Last Mgmt Change     : 05/27/2019 15:26:46
Host Identification
 User Name           : wholesale.com (domain only)
Matched Objects      : userName
Password Type        : ignore
PADO Delay           : 0msec
Diameter app policy  : (Not Specified)
Diameter auth policy : (Not Specified)
Force IPv6CP         : Disabled
Ignore DF Bit        : Disabled
DHCPv6 lease times
 Renew timer         : > 9999 days
 Rebind timer        : > 9999 days
 Preferred lifetime  : 0d 00:00:00
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
 Valid lifetime      : 0d 00:00:00
Identification Strings (option 254)
 Subscriber Id       : user@wholesale.com
 SLA Profile String  : sla-profile-1
 Sub Profile String  : sub-profile-1
L2TP
 Service             : 65536
 Tunnel Group        : wholesale.com
MSAP defaults
Filter Overrules
Access loop info
===============================================================================
```

The **debug router l2tp event lac-session-setup** command shows the LAC tunnel selected for this example.

```
12 2019/05/28 15:32:48.860 CEST MINOR: DEBUG #2001 vprn65536 PPPoE 263411->L2TP
"PPPoE 263411->L2TP: UDP 192.168.33.1:1701 -> 192.168.33.2:1701
preference 50 tunnel wholesale.com:wholesale.com
    request to open new tunnel 7432"
```

```
13 2019/05/28 15:32:48.860 CEST MINOR: DEBUG #2001 vprn65536 PPPoE 263411->L2TP
"PPPoE 263411->L2TP: UDP 192.168.33.1:1701 -> 192.168.33.2:1701
preference 50 tunnel wholesale.com:wholesale.com
    create session 487068937"
```

For the operational PPPoE session information in IES 1/base instance, use the following command:

```
[]
A:admin@LAC# show service id 1 ppp session

===============================================================================
PPP sessions for service 1
===============================================================================
User-Name
  Descr.
          Up Time       Type  Termination     IP/L2TP-Id/Interface-Id MC-Stdby
-------------------------------------------------------------------------------
user1@wholesale.com
  svc:1 sap:1/1/3:100 mac:00:00:00:00:01:01 sid:1
          0d 00:02:56   oE    lac             487068937
-------------------------------------------------------------------------------
No. of PPP sessions: 1
===============================================================================
```

Operational tunnel information for VPRN 65536 can be obtained using following command.

```
[]
A:admin@LAC# show router 65536 l2tp tunnel
===============================================================================
Conn ID    Loc-Tu-ID Rem-Tu-ID State             Blacklist-state   Ses Active
  Group                                                            Ses Total
    Assignment
-------------------------------------------------------------------------------
487063552  7432      10509     established       not-blacklisted   1
  wholesale.com                                                    1
    wholesale.com
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The operational session information for VPRN 65536 shows the session is in the established state.

```
[]
A:admin@LAC# show router 65536 l2tp session

===============================================================================
L2TP Session Summary
===============================================================================
ID                  Control Conn ID     Tunnel-ID   Session-ID  State
-------------------------------------------------------------------------------
487068937           487063552           7432        5385        established
-------------------------------------------------------------------------------
No. of sessions: 1
===============================================================================
```

Advanced topics

Non-session-triggered L2TP tunnel setup

In addition to the ppp-session-triggered setup, an L2TP tunnel can also be set up via a tools command or an auto-establish command.

These non-session-triggers are useful, for example, during the initial configuration phase where the LAC-LNS tunnel setup can be tested without the need for a user to attempt and establish a PPPoE connection.

The PPPoE user still triggers the L2TP session-setup over this L2TP tunnel and RADIUS needs to return an l2tp group-name with the relevant name during authentication.

Auto-establish

Every minute, a check is performed to determine if tunnels need to be established (a process referred to as scan auto-establish). The tunnel state is establishedIdle when the tunnel is setup, and becomes established when user triggered sessions are set up over this tunnel.

```
service
    vprn 65536
        l2tp {
            admin-state enable
            group "wholesale.com" {
                admin-state enable
                tunnel "wholesale.com" {
                    admin-state enable
                    auto-establish true
                    peer 192.168.33.2
                    local-address 192.168.33.1
                    local-name "lac-pe1"
                }
            }
        }
    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

There is no difference in operational behavior for a tunnel set up via a session-trigger or an auto-establish command. Removing the auto-establish parameter has no impact on active tunnels (establishedIdle or established).

```
[]
A:admin@LAC# show router 65536 l2tp tunnel
===============================================================================
Conn ID    Loc-Tu-ID Rem-Tu-ID State              Blacklist-state   Ses Active
  Group                                                             Ses Total
    Assignment
-------------------------------------------------------------------------------
876150784  13369     11022     establishedIdle    not-blacklisted   0
  wholesale.com                                                     0
    wholesale.com
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================
```

Tools tunnel start

First revert to the original situation, without auto-establish, as follows:

```
configure exclusive
service vprn 65536 l2tp group "wholesale.com" tunnel "wholesale.com"
                                                      delete auto-establish
commit
```

Verify the tunnel does not exist anymore, as follows:

```
[]
A:admin@LAC# show router 65536 l2tp tunnel
No entries found.
```

Issue the tools command to manually establish the L2TP tunnel, as follows:

```
[]
A:admin@LAC# tools perform router 65536 l2tp group "wholesale.com" tunnel
                                                        "wholesale.com" start
```

Verify a new tunnel has been created, as follows:

```
[]
A:admin@LAC# show router 65536 l2tp tunnel
===============================================================================
Conn ID    Loc-Tu-ID Rem-Tu-ID State              Blacklist-state   Ses Active
  Group                                                             Ses Total
    Assignment
-------------------------------------------------------------------------------
45940736   701       8671      establishedIdle    not-blacklisted   0
  wholesale.com                                                     0
    wholesale.com
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

How long remains a tunnel idle before torn down?

An L2TP tunnel can be torn down automatically, after the expiration of an idle-timer, or manually through a tools command.

Idle-timeout

A persistent tunnel is a tunnel that remains available after the last session over that tunnel is closed. To create a persistent tunnel, the idle-timeout parameter must be set to infinite.

A non-persistent tunnel is torn down immediately (idle-timeout zero) after the last session over that tunnel is closed or after a configurable delay. The idle-timeout parameter is set via the RADIUS [26-6527-49] Alc-Tunnel-Idle-Timeout attribute or the corresponding node parameter. The default value for this parameter is infinite (persistent).

```
configure router l2tp | configure service vprn l2tp
       idle-timeout [0..3600] s
       ---snip---
       group <tunnel-group-name>
           idle-timeout [0..3600] s | infinite
           ---snip---
           tunnel  <tunnel-name>
              idle-timeout [0..3600] s | infinite
              ---snip---
```

The following shows an example of a non-persistent tunnel (idle-timeout 30 seconds). The tunnel changes state from established to establishedIdle when the last session is terminated. Idle-timeout seconds later, the session changes to the closed state. For the purpose of troubleshooting, the operational data stays available for destruct-timeout seconds (see later).

```
[]
A:admin@LAC# show router 65536 l2tp tunnel detail

===============================================================================
L2TP Tunnel 1068892160
===============================================================================

Connection ID: 1068892160
Protocol     : v2
State        : closed
IP           : 192.168.33.1

---snip---

Name         : lac-pe1
Remote Name  : LNS1
Assignment ID: wholesale.com
Group Name   : wholesale.com
Acct. Policy : N/A
Error Message: idle timeout (30 seconds) expired

                                       Remote Conn ID    : 192479232
Tunnel ID         : 16310              Remote Tunnel ID  : 2937
Preference        : 50                 Receive Window    : 64
Hello Interval (s): 60                 AVP Hiding        : never
Idle TO (s)       : 30                 Destruct TO (s)   : 60
Max Retr Estab    : 5                  Max Retr Not Estab: 5
Cfg'd Sess Limit  : unlimited          Oper Session Limit: 32767                    '
Transport Type    : udpIp              Challenge         : never
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
---snip---

No. of tunnels: 1
===============================================================================
```

The following shows an example of a persistent tunnel (idle-timeout infinite).

```
[]
A:admin@LAC# show router 65536 l2tp tunnel detail

===============================================================================
L2TP Tunnel 9240576
===============================================================================

Connection ID: 9240576
Protocol     : v2
State        : establishedIdle
IP           : 192.168.33.1

---snip---

Name         : lac-pe1
Remote Name  : LNS1
Assignment ID: wholesale.com
Group Name   : wholesale.com
Acct. Policy : N/A
Error Message: N/A

                                      Remote Conn ID    : 750714880
Tunnel ID         : 141               Remote Tunnel ID  : 11455
Preference        : 50                Receive Window    : 64
Hello Interval (s): 60                AVP Hiding        : never
Idle TO (s)       : infinite          Destruct TO (s)   : 60
Max Retr Estab    : 5                 Max Retr Not Estab: 5
Cfg'd Sess Limit  : unlimited         Oper Session Limit: 32767                '
Transport Type    : udpIp             Challenge         : never

---snip---

No. of tunnels: 1
===============================================================================
```

Tools tunnel stop

In addition to the idle-timeout used for tunnel termination, a tools stop command is also available that can be used to terminate persistent and non-persistent tunnels at any moment in time. Be aware that this command is very destructive and destroys all sessions carried over the closed tunnel.

Following command shows the tunnel is in the establishedIdle state.

```
[]
A:admin@LAC# show router 65536 l2tp tunnel
===============================================================================
Conn ID    Loc-Tu-ID Rem-Tu-ID State              Blacklist-state   Ses Active
  Group                                                            Ses Total
    Assignment
-------------------------------------------------------------------------------
331022336  5051      9053      establishedIdle    not-blacklisted   0
  wholesale.com                                                    0
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    wholesale.com
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================
```

The following command terminates the l2tp tunnel. The tunnel is aborted (the LAC sends StopCCN) using the <connection-id> or <tunnel-group-name>+<tunnel-name> as input. This StopCCN indicates "operator request" as the error reason.

```
[]
A:admin@LAC# tools perform router 65536 l2tp group wholesale.com
                                                        tunnel wholesale.com stop
INFO: CLI #2007: Info while processing command - INFO: CLI stopped 1 tunnels,
                                                            destructed 0 tunnels.
```

The following debug output shows the tunnel being aborted.

```
5 2019/05/28 16:01:14.824 CEST MINOR: DEBUG #2001 vprn65536 L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.168.33.1:1701 -> 192.168.33.2:1701
tunnel 9053 session 0, ns 2 nr 1, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        StopControlConnectionNotification(4)
    AVP ResultCode(0,1), flags: mandatory, reserved=0
        result-code: "generalRequestToClearControlConnection"(1),
                                error-code: "noGeneralError"(0)
        error-msg: "operator request"
    AVP AssignedTunnelId(0,9), flags: mandatory, reserved=0
        5051"
```

Alternatively, the tunnel can also be stopped with the following command. The effect would be the same.

```
[]
A:admin@LAC# tools perform router 65536 l2tp tunnel 964886528 stop
```

Keepalive - L2TP Hello

A keepalive mechanism is employed by L2TP in order to differentiate between tunnel outages and no control or data activity on a tunnel for an extended period. This is accomplished by injecting Hello control messages after a specified period of time has elapsed since the last data or control message (ZLB not included) was received on a tunnel. As for any other L2TP control message, if the Hello message is not reliably delivered, then the tunnel is declared down and reset, as defined in RFC 2661, *Layer Two Tunneling Protocol "L2TP"*. This means that SR OS does not send Hello packets if session control traffic is handled over this tunnel. The hello timer is reset if the system transmits any control packet over this tunnel (ZLB packets and data traffic are not taken into account).

The keepalive function is disabled (not recommended) using RADIUS [26-6527-50] Alc-Tunnel-Hello-Interval -1 or hello-interval infinite (default 300). The number of retries for unsuccessful Hello packet delivery equals RADIUS [26-6527-52] Alc-Tunnel-Max-Retries-Estab or node parameter max-retries-estab (default 5). The retry interval is initially set to 1 second and doubles on each retry with a maximum interval of 8 seconds. Using a max-retries-estab 7 results in a retry of [1,2,4,8,8,8,8 seconds].

```
configure router l2tp | configure service vprn l2tp]
    hello-interval [60..3600] s | infinite  # default 300 s
    max-retries-estab [2..7]                # default 5
    ---snip---
    group <tunnel-group-name>
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        hello-interval [60..3600] s | infinite
        max-retries-estab [2..7]
        ---snip---
        tunnel  <tunnel-name>
            hello-interval [60..3600] s | infinite
            max-retries-estab [2..7]
            ---snip---
```

For example, the LAC can be configured with an hello-timer of 1 minute and the LNS with an hello-timer of 2 minutes. The hello-timer interval for LAC and LNS do not have to be same because the keepalive mechanism works asynchronous. See Figure 356: L2TP keepalive mechanism.

```
[]
A:admin@LAC# show router 65536 l2tp tunnel
===============================================================================
Conn ID    Loc-Tu-ID Rem-Tu-ID State             Blacklist-state   Ses Active
  Group                                                            Ses Total
    Assignment
-------------------------------------------------------------------------------
1002438656 15296     15972     established       not-blacklisted   1
  wholesale.com                                                    1
    wholesale.com
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================
```

Figure 356: L2TP keepalive mechanism

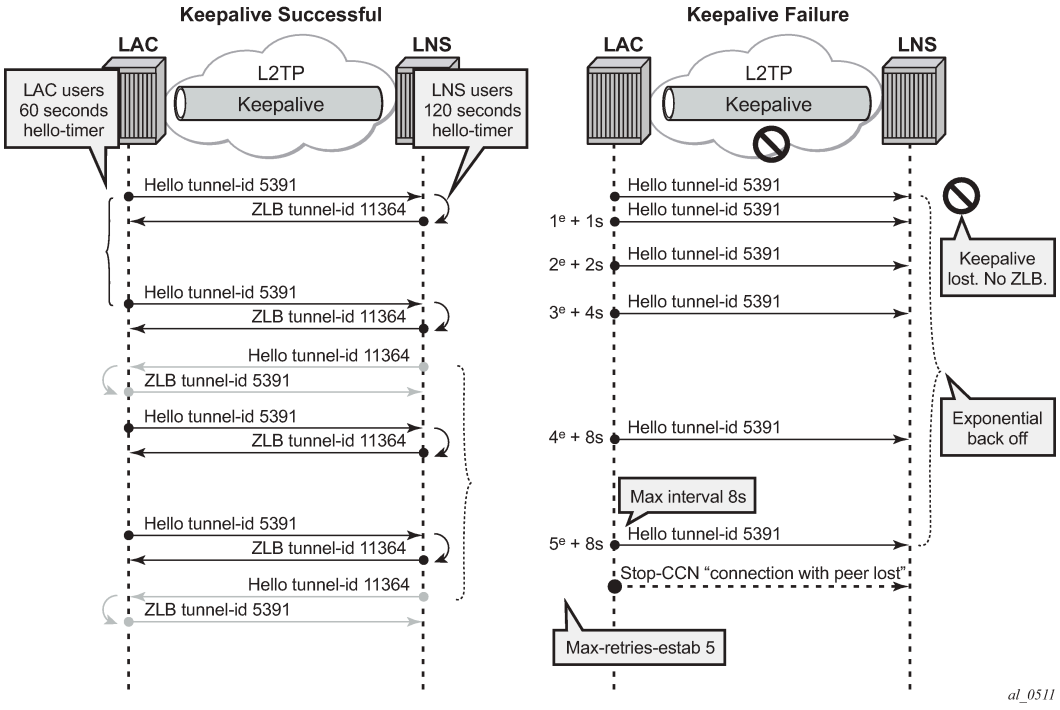

Figure 356: L2TP keepalive mechanism shows the tunnel being closed after 5 unsuccessful Hello deliveries with error-message connection with peer lost.

```
[]
A:admin@LAC# show router 65536 l2tp tunnel detail
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
L2TP Tunnel 1002438656
===============================================================================
Connection ID: 1002438656
Protocol     : v2

---snip---

Acct. Policy : N/A
Error Message: connection with peer lost

                                       Remote Conn ID    : 1046740992
Tunnel ID         : 15296              Remote Tunnel ID  : 15972
Preference        : 50                 Receive Window    : 64
Hello Interval (s): 60                 AVP Hiding        : never

---snip---

No. of tunnels: 1
===============================================================================
```

Keeping closed tunnel and session information

The destruct-timeout parameter (expressed in seconds) controls the period of time that the tunnel, or session data related to a closed (disconnected) tunnel, or session persists before being removed. The destruct_timeout is a debugging aid by keeping underlying memory structures after the tunnel or session is terminated. It is configured via the RADIUS [26-6527-51] Alc-Tunnel-Destruct-Timeout attribute or the corresponding node parameter. Default value for this parameter is 60 seconds.

```
configure router l2tp | configure service vprn l2tp
    destruct-timeout [60..86400]
    ---snip---
    group <tunnel-group-name>
        destruct-timeout [60..86400]
        ---snip---
        tunnel  <tunnel-name>
            destruct-timeout [60..86400]
```

The following output shows a session that is closed and the reason for it being terminated.

```
[]
A:admin@LAC# show router 65536 l2tp session detail

===============================================================================
L2TP Session 904095211
===============================================================================
Connection ID: 904095211
State        : closed
Tunnel Group : wholesale.com
Assignment ID: wholesale.com
Error Message: Terminated by PPPoE: Received PPPoE PADT

Control Conn ID   : 904069120           Rem Cntrl Conn ID : 716898304
Tunnel ID         : 13795               Remote Tunnel ID  : 10939
Session ID        : 26091               Remote Session ID : 20875
PW Type           : ppp                 Remote Conn ID    : 716919179
Time Started      : 05/28/2019 16:18:17
Time Established  : 05/28/2019 16:18:17 Time Closed       : 05/28/2019 16:18:20
CDN Result        : generalError        General Error     : vendorSpecific
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
No. of sessions: 1
===============================================================================
```

The following output shows a tunnel that is closed and the reason for it being closed.

```
[]
A:admin@LAC# show router 65536 l2tp tunnel detail

===============================================================================
L2TP Tunnel 904069120
===============================================================================

Connection ID: 904069120
Protocol     : v2
State        : closedByPeer
IP           : 192.168.33.1
UDP          : 1701
Peer IP      : 192.168.33.2
Peer UDP     : 1701
Tx dst-IP    : 192.168.33.2
Tx dst-UDP   : 1701
Rx src-IP    : 192.168.33.2
Rx src-UDP   : 1701
Name         : lac-pe1
Remote Name  : LNS1
Assignment ID: wholesale.com
Group Name   : wholesale.com
Acct. Policy : N/A
Error Message: idle timeout (60 seconds) expired

---snip---

No. of tunnels: 1
===============================================================================
```

When the Destruct TO expires the tunnel and session is deleted, as follows:

```
[]
A:admin@LAC# show router l2tp session detail
No entries found.

[]
A:admin@LAC# show router l2tp tunnel detail
No entries found.
```

Floating peers

A floating peer exists if the peer LNS address indicated in the source address of the SCCRP is different from the peer address known on the LAC. Floating peer allowance is configuration driven and is rejected by default.

The parameter peer-address-change-policy specifies whether the LAC accepts, ignores or rejects requests from a peer to change the destination IP address or UDP port.

```
configure router l2tp | configure service vprn l2tp
       peer-address-change-policy accept | ignore | reject
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- accept — Specifies that this system accepts any source IP address change for received L2TP control messages related to a locally originated tunnel in the state wait-reply and rejects any peer address change for other tunnels. In case the new peer IP address is accepted, it is learned and used as destination address in subsequent L2TP messages.

  Figure 357: Floating peers accept

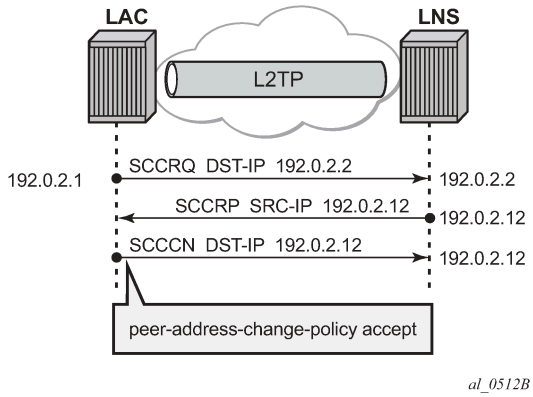

- Ignore — Specifies that this system ignores any source IP address change for received L2TP control messages, does not learn any new peer IP address and does not change the destination address in subsequent L2TP messages.

  Figure 358: Floating peers ignore

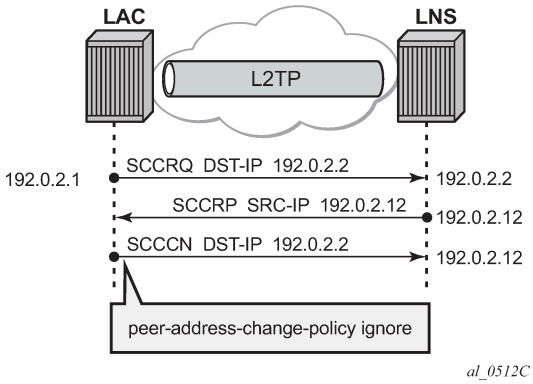

- Reject — Specifies that this system rejects any source IP address change for received L2TP control messages and drops those messages.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 359: Floating peers reject

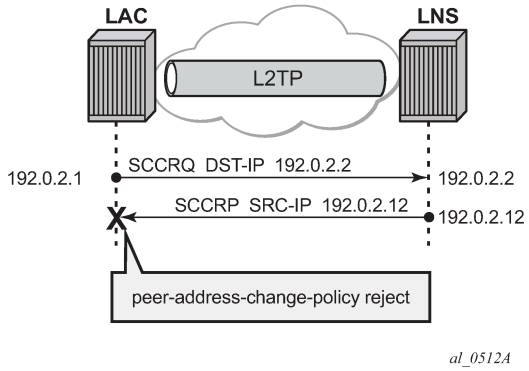

The values Peer IP, Tx dst-IP and Rx src-IP in the **show router l2tp tunnel detail** command indicates if floating peers are used or not.

An example of a floating peer (peer-address-change-policy accept) is as follows.

```
[]
A:admin@:LAC# show router l2tp tunnel detail
=======================================
L2TP Tunnel Status
=======================================
Connection ID: 897122304
State        : established
IP           : 192.0.2.1
UDP          : 1701
Peer IP      : 192.0.2.2 # (1) peer address used in SCCRQ
Peer UDP     : 1701
Tx dst-IP    : 192.0.2.12 # (3) peer address used in SCCCN
Tx dst-UDP   : 1701
Rx src-IP    : 192.0.2.12 # (2) SCCRP different IP received
Rx src-UDP   : 1701
---snip---
```

Tx/Rx connect speed - AVP 24/38

The connect speed (TX AVP 24 and RX AVP 38) is passed in the ICCN messages sent from the LAC to the LNS. The L2TP AVP 24 defines the (Tx) connect speed in bps from the perspective of traffic flowing from the LAC towards the subscriber (BNG downstream rate).The L2TP AVP 38 defines the (Rx) connect speed in bps from the perspective of traffic flowing from the subscriber towards the LAC (BNG upstream rate).

The report-rate configuration option indicates what rate is reported to the LNS when creating an L2TP session.

```
configure subscriber-mgmt sla-profile <sla-profile-name> ingress | egress
    report-rate agg-rate-limit|scheduler|pppoe-actual-rate|
                                    policer|rfc5515-actual-rate
```

- agg-rate-limit — Take the aggregate rate as received from the RADIUS Access-Accept message in VSA Alc-Subscriber-QoS-Override. When this RADIUS VSA is not present in the Access-Accept, or

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

when RADIUS is not used, then take the configured aggregate rate limit. In the case where this is not configured, then take the port rate.

- scheduler <scheduler-name> — Take the rate of the specified scheduler. In case the scheduler is not linked with the scheduler-policy from the subscriber-profile, then take the port rate.
- pppoe-actual-rate — Take the rate from the DSL-Forum Vendor-Specific PPPoE Tag when available, otherwise take the port rate.
- rfc5515-actual-rate — Put the same value as the transmitted Actual-Data-Rate-Upstream AVP in the Rx-Connect-Speed AVP, and the same value as the transmitted Actual-Data-Rate-Downstream AVP in the Tx-Connect-Speed AVP.

Calling number AVP 22 format

The format of AVP 22 Calling Number in the ICRQ message is configurable via the parameter calling-number-format. The default format is "%S<space>%s" and corresponds to the concatenation of system-name<space>sap-id. Available parameters are %S (system-name), %c (Agent Circuit Id), %r Agent Remote Id, %s (sap-id), %l (Logical Line ID) and fixed strings. A combination can be configured from any of these parameters, but the total configured format cannot exceed 255 characters.

Example 1: Default configuration.

```
configure exclusive
service vprn 65536 l2tp lac calling-number-format "%S %s"
```

```
18 2019/05/28 16:26:20.903 CEST MINOR: DEBUG #2001 vprn65536 L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.168.33.1:1701 -> 192.168.33.2:1701
tunnel 11703 session 0, ns 2 nr 1, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        IncomingCallRequest(10)
    AVP AssignedSessionId(0,14), flags: mandatory, reserved=0
        20683
    AVP CallSerialNumber(0,15), flags: mandatory, reserved=0
        283037
    AVP CallingNumber(0,22), flags: mandatory, reserved=0
        "LAC 1/1/3:100"
    AVP AgentCircuitId(3561,1), flags:, reserved=0
        "circuit0"
    AVP AgentRemoteId(3561,2), flags:, reserved=0
        "remote0"
    AVP ActDataRateUp(3561,129), flags:, reserved=0
        2000000
    AVP ActDataRateDown(3561,130), flags:, reserved=0
        4000000"
```

Example 2: Customized configuration and all parameters (%S %s %c) are available to construct the requested AVP 22.

```
configure exclusive
service vprn 65536 l2tp lac calling-number-format "start-%S###%s###%c-end"
```

```
18 2019/05/28 16:29:39.771 CEST MINOR: DEBUG #2001 vprn65536 L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.168.33.1:1701 -> 192.168.33.2:1701
tunnel 4838 session 0, ns 2 nr 1, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        IncomingCallRequest(10)
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    AVP AssignedSessionId(0,14), flags: mandatory, reserved=0
        16185
    AVP CallSerialNumber(0,15), flags: mandatory, reserved=0
        283038
    AVP CallingNumber(0,22), flags: mandatory, reserved=0
        "start-LAC###1/1/3:100###circuit0-end"
    AVP AgentCircuitId(3561,1), flags:, reserved=0
        "circuit0"
    AVP AgentRemoteId(3561,2), flags:, reserved=0
        "remote0"
    AVP ActDataRateUp(3561,129), flags:, reserved=0
        2000000
    AVP ActDataRateDown(3561,130), flags:, reserved=0
        4000000"
```

Example 3: Customized configuration and not all parameters are available to construct the requested AVP 22. Option-82 circuit-id (%c),remote-id (%r), and LLID (%l) information are lacking and therefore missing (skipped) in the formatted attribute.

```
configure exclusive
router vprn 65536 l2tp lac calling-number-format "%S#%c#%r#%l#%s"
```

```
18 2019/05/28 16:32:11.553 CEST MINOR: DEBUG #2001 vprn65536 L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.168.33.1:1701 -> 192.168.33.2:1701
tunnel 14364 session 0, ns 2 nr 1, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        IncomingCallRequest(10)
    AVP AssignedSessionId(0,14), flags: mandatory, reserved=0
        14710
    AVP CallSerialNumber(0,15), flags: mandatory, reserved=0
        283039
    AVP CallingNumber(0,22), flags: mandatory, reserved=0
        "LAC#circuit0#remote0##1/1/3:100"
    AVP AgentCircuitId(3561,1), flags:, reserved=0
        "circuit0"
    AVP AgentRemoteId(3561,2), flags:, reserved=0
        "remote0"
    AVP ActDataRateUp(3561,129), flags:, reserved=0
        2000000
    AVP ActDataRateDown(3561,130), flags:, reserved=0
        4000000"
```

Prevent LAC from transmitting calling number AVP 22 to LNS

By default, the LAC includes the Calling Number AVP 22 in the L2TP incoming-call-request (ICRQ) packets transmitted to LNS. This AVP identifies the interface that is connected to the customer in the access network. Network access interface information can be hidden by configuring the LAC not to send the Calling Number AVP to the LNS.

Use the following command to disable the sending of L2TP Calling Number AVP 22.

```
configure exclusive
router l2tp exclude-avps calling-number
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

AVP 100 - Cisco-Nas-Port

Interoperation with a Cisco LNS requires that the LAC communicates a NAS port type to the LNS via the L2TP ICRQ 'Cisco Nas Port Info AVP (100)'. This AVP (100) includes information that identifies the NAS port and indicates whether the port type is Ethernet or ATM and is configured via the cisco-nas-port parameter.

The Cisco AVP 100 format is as follows:

- First 5 bytes are NAS-Port-Type:
  - 0f10090203 (Ethernet)
  - 0f10090201 (ATM)
- Remaining 4 bytes corresponds with the configured cisco-nas-port value

Example:

- Ethernet 12b s-vlan-id; 10b c-vlan-id; 3b slot number; 2b MDA nbr; 5b port
- ATM 12b VPI; 10b VCI; 3b slot number; 2b MDA nbr; 5b port

```
configure exclusive
service vprn 65536 l2tp lac cisco-nas-port ethernet "*12o*10i*3s*2m*5p"
```

nas-port 1/1/3:100 corresponds to 102563 (000000000000 0001100100 001 01 00011).

```
22 2019/05/28 16:37:35.879 CEST MINOR: DEBUG #2001 vprn65536 L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.168.33.1:1701 -> 192.168.33.2:1701
tunnel 6288 session 0, ns 2 nr 1, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        IncomingCallRequest(10)
    AVP CiscoNasPort(9,100), flags:, reserved=0
        102563 type=ethernet(0f:10:09:02:03)
    AVP AssignedSessionId(0,14), flags: mandatory, reserved=0
        28354
    AVP CallSerialNumber(0,15), flags: mandatory, reserved=0
        283041
    AVP CallingNumber(0,22), flags: mandatory, reserved=0
        "LAC 1/1/3:100"
    AVP AgentCircuitId(3561,1), flags:, reserved=0
        "circuit0"
    AVP AgentRemoteId(3561,2), flags:, reserved=0
        "remote0"
    AVP ActDataRateUp(3561,129), flags:, reserved=0
        2000000
    AVP ActDataRateDown(3561,130), flags:, reserved=0
        4000000"
```

L2TP group/peer/tunnel draining

When the LAC has established sessions, the LAC can avoid the creation of new sessions for a specific group, peer, or tunnel, via the **drain** command.

No new sessions are created for a group, peer or tunnel that is being drained (draining state) but the current sessions are left intact.

After the **drain** command is issued, the group, peer, or tunnel moves from a draining to drained state when the last session is closed. A drained group, peer, or tunnel can then be managed (reconfigured, deleted) without any user impact.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Be aware that a group, peer, or tunnel in a draining or drained state is skipped in the tunnel selection process. The next example shows a tunnel draining; group and peer draining works according in the same way.

A tunnel has 1 session and is in established state.

```
[]
A:admin@LAC# show router 65536 l2tp tunnel
===============================================================================
Conn ID    Loc-Tu-ID Rem-Tu-ID State              Blacklist-state   Ses Active
  Group                                                             Ses Total
    Assignment
-------------------------------------------------------------------------------
285540352  4357      7237      established        not-blacklisted   1
  wholesale.com                                                     1
    wholesale.com
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================
```

The following tools **drain** command puts the tunnel in a draining state and leaves the sessions intact.

```
[]
A:admin@LAC# tools perform router 65536 l2tp tunnel 1023868928 drain
```

Initially the tunnel is in the draining state.

```
[]
A:admin@LAC# show router 65536 l2tp tunnel
===============================================================================
Conn ID    Loc-Tu-ID Rem-Tu-ID State              Blacklist-state   Ses Active
  Group                                                             Ses Total
    Assignment
-------------------------------------------------------------------------------
285540352  4357      7237      draining           not-blacklisted   1
  wholesale.com                                                     1
    wholesale.com
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================
```

The tunnel moves to the drained state at the moment the last session is closed. Debugging shows that a drained tunnel is also not used as last resort and is skipped during the tunnel selection process.

```
[]
A:admin@LAC# show router 65536 l2tp tunnel
===============================================================================
Conn ID    Loc-Tu-ID Rem-Tu-ID State              Blacklist-state   Ses Active
  Group                                                             Ses Total
    Assignment
-------------------------------------------------------------------------------
285540352  4357      7237      drained            not-blacklisted   0
  wholesale.com                                                     1
    wholesale.com
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The following output shows new sessions cannot select a drained tunnel.

```
19821 2019/05/28 16:46:52.549 CEST MINOR: DEBUG #2001 vprn65536 PPPoE 282635->L2TP
"PPPoE 282635->L2TP: UDP 192.168.33.1:1701 -> 192.168.33.2:1701
preference 50 tunnel wholesale.com:wholesale.com
    no additional session can be created in tunnel 4357"
```

```
19822 2019/05/28 16:46:52.549 CEST MINOR: DEBUG #2001 vprn65536 PPPoE 282635->L2TP
"PPPoE 282635->L2TP:
    stop: no more tunnels can be tried"
```

The drained tunnel can then be closed without user impact.

```
[]
A:admin@LAC# tools perform router "65536" l2tp tunnel 285540352 stop
```

```
19846 2019/05/28 16:48:57.338 CEST MINOR: DEBUG #2001 vprn65536 L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.168.33.1:1701 -> 192.168.33.2:1701
tunnel 7237 session 0, ns 8 nr 2, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        StopControlConnectionNotification(4)
    AVP ResultCode(0,1), flags: mandatory, reserved=0
        result-code: "generalRequestToClearControlConnection"(1),
                                error-code: "noGeneralError"(0)
        error-msg: "operator request"
    AVP AssignedTunnelId(0,9), flags: mandatory, reserved=0
        4357"
```

For draining and undraining for example a group, following commands can be used.

```
tools perform router 65536 l2tp group "wholesale.com" drain
tools perform router 65536 l2tp group "wholesale.com" no drain
```

Conclusion

This example provides the LAC L2TP access server configuration and troubleshooting commands for the LAA architecture (tunneled-access) model.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Vport-Based Load Balancing on a LAG

This chapter describes Vport-based load balancing on a LAG.

Topics in this chapter include:

- Applicability
- Overview
- Configuration
- Conclusion
- Appendix with configuration files

Applicability

The information and MD-CLI configuration in this chapter is based on SR OS Release 22.10.R1.

Vport-based load balancing on a LAG applies to all service router (SR) systems based on FP4 and higher. This functionality is not supported on Virtualized Service Router (VSR) because VSR does not support Vports.

The target audience for this chapter is subscriber management users who are already familiar with subscriber management and EVPN VPWS. In this chapter, the emphasis is on the LAG configuration.

Overview

In enhanced subscriber management (ESM), a Vport is a representation of a downstream access node that hosts subscribers. A Vport is used to control bandwidth toward the access node. When deployed on a link aggregation group (LAG), Vports must be configured on every LAG member port, as shown in the figure Figure 360: Vport concept.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 360: Vport concept

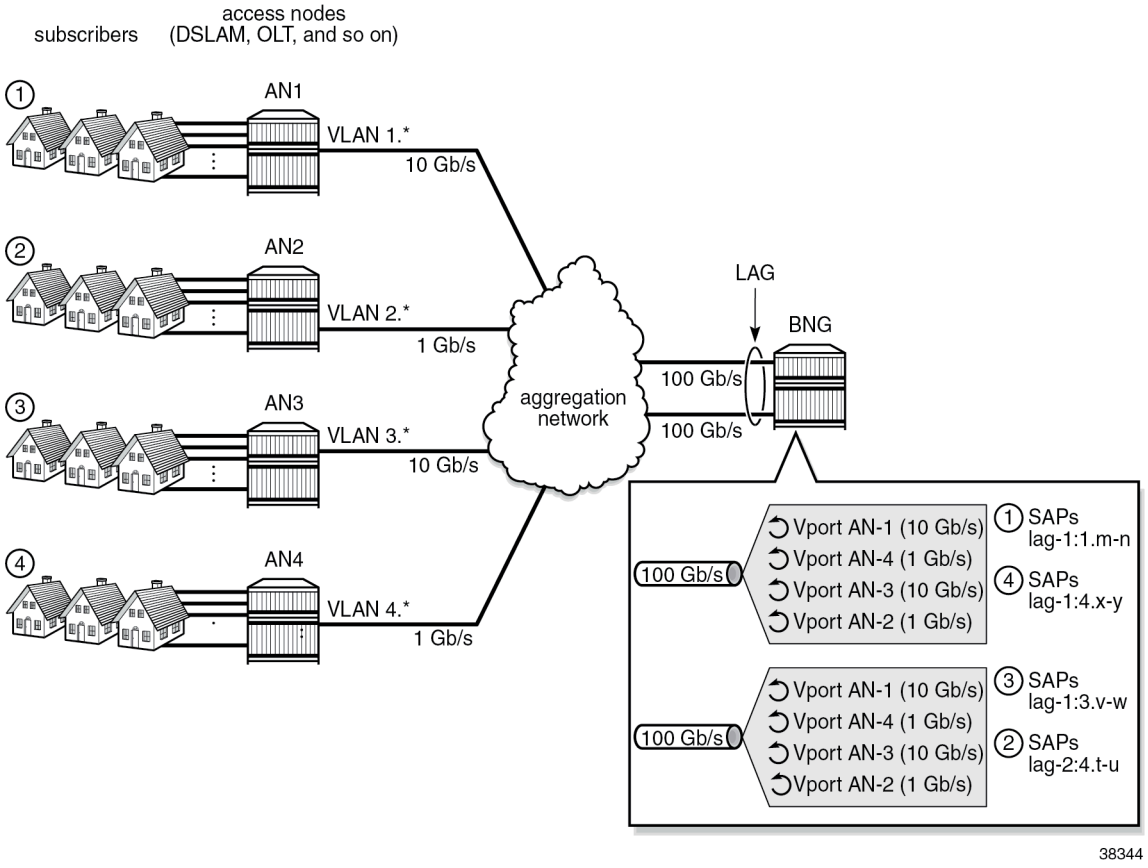

The Vport bandwidth in a LAG can be configured in two ways:

- The configured Vport bandwidth is assigned to every LAG member port.
- The Vport on each link gets only a fraction of the total bandwidth of the access node that it represents.

In both cases, the configured Vport bandwidth corresponds to the full capacity of the access node that it controls (a bandwidth cap or ceiling per access node). The difference is in how this bandwidth is assigned by the system to each LAG member port, which is controlled by the **lag access adapt-qos mode link | port-fair** command.

When the configured Vport bandwidth is assigned to every LAG member port, oversubscription of the access node must be avoided by directing all traffic associated with a Vport to a single LAG member link. This configuration requires load balancing per Vport and is the focus of this chapter. Per-Vport load balancing is suitable in environments with a large variation in bandwidth consumption between subscribers within a Vport.

When the Vport on each link gets only a fraction of the total bandwidth of the access node, the configured Vport bandwidth is automatically distributed over the LAG member links: the Vport bandwidth per LAG member port is the Vport bandwidth divided by the number of LAG member ports.

In other words, each LAG member port gets an equal share of the configured bandwidth. This scenario calls for load balancing per subscriber, where traffic toward the same access node is distributed over the member links. This scenario is suitable for environments with subscribers with similar bandwidth

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

requirements. Otherwise, a few high-bandwidth-consuming subscribers within a Vport can use the same link, causing the aggregate bandwidth to exceed the Vport cap per LAG member port. This issue can cause unnecessary congestion and packet drops, while the instance of the same Vport on the other links remain underutilized.

In Releases earlier than Release 22.10, the implementation of load balancing per Vport, where the configured Vport bandwidth is assigned to every member, implies the following:

- The subscriber SAPs must be instantiated on all LAG member ports, which has a negative effect on the subscriber scalability that is driven by the finite number of SAPs per forwarding complex.
- All Vports are accounted with equal weight in the hashing algorithm, regardless of the configured bandwidth of the Vports, which could lead to suboptimal load balancing when Vports with different bandwidths are present on the LAG.

Release 22.10 expands load balancing per Vport to include the class and weight per Vport and the support for a single subscriber SAP instance per forwarding complex.

The focus of this chapter is on the class- and weight-based Vport load balancing where a single copy of a SAP is created per forwarding complex for each subscriber.

The figure Figure 361: Example topology shows the example topology used in this chapter.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 361: Example topology

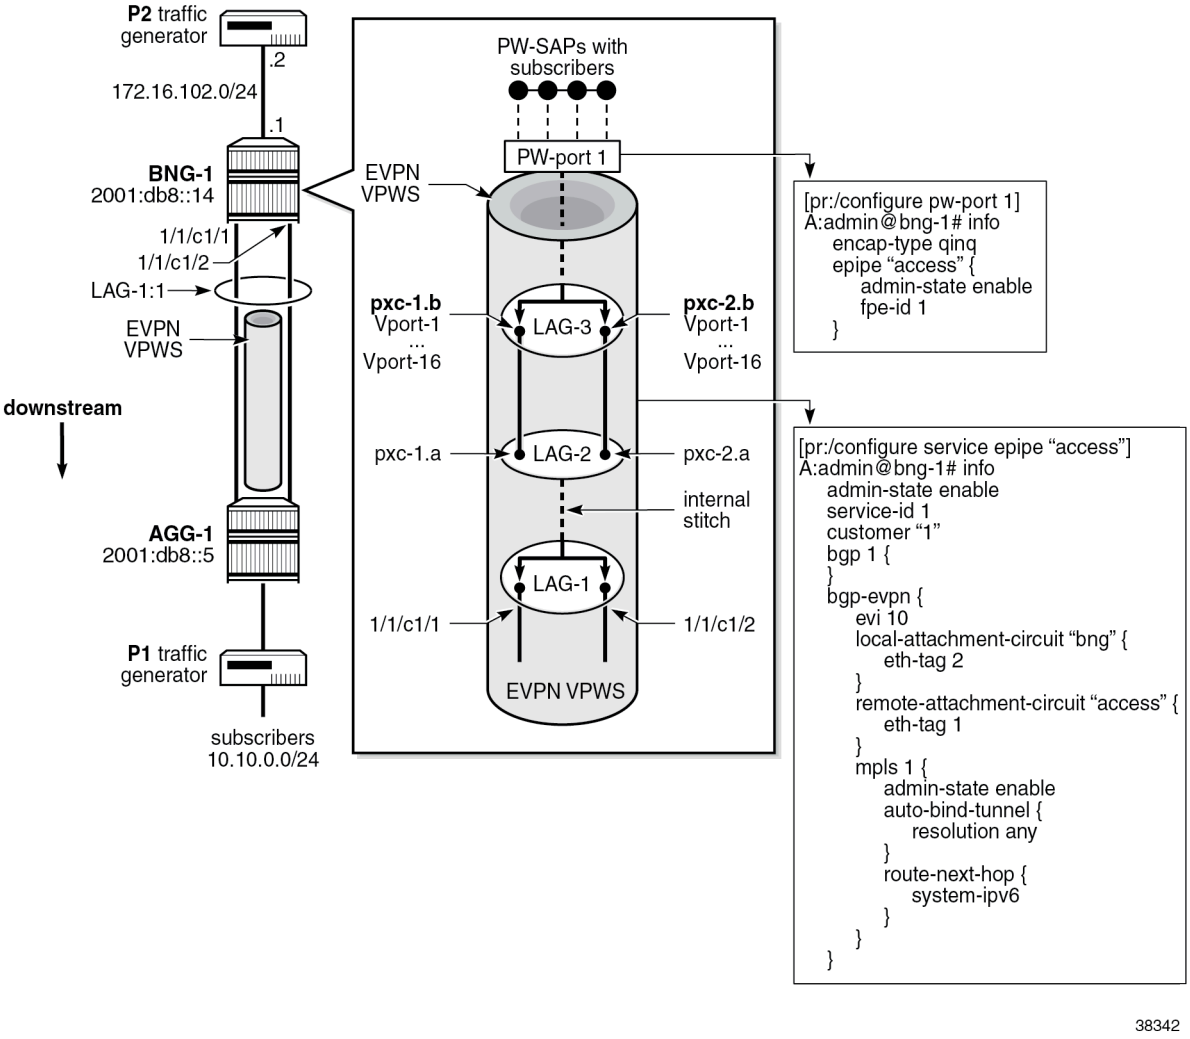

In this setup:

- The subscribers are instantiated in 1:1 mode of operation (one subscriber per SAP) on a pseudowire (PW) port in the Broadband Network Gateway (BNG) "BNG-1".
- The PW-port is the termination point for the EVPN VPWS, which connects the BNG to an aggregation node "AGG-1" in the access network.
- The PW-port is based on the Forwarding Path Extension (FPE), with the LAG "lag-3" containing two PXC member ports, pxc-1.b and pxc-2.b.
- There are 16 Vports configured on each of the PXC ports in the LAG "lag-3".
- 32 IPv4 subscribers are associated with these Vports. The individual subscriber association with Vports is shown in Table 1.
- The EVPN VPWS connecting the BNG to "AGG-1" is configured over the LAG "lag-1" with two member ports on the BNG (ports 1/1/c1/1 and 1/1/c1/2).

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- Traffic is sent in the downstream direction, from traffic generator port "p2" toward the simulated subscribers on traffic generator port "p1".
- The traffic flow and traffic load balancing in the BNG is examined at two points:
  - on the "lag-3" underlaying the PW-port, where traffic is load-balanced per Vport in the ESM context
  - on the "lag-1" where traffic is load-balanced in the service context, outside of the ESM
- Bandwidth utilization on the LAG links is examined via **show** and **monitor** commands at the BNG.

Traffic load balancing on the LAGs "lag-1" and "lag-3" are independent of each other. Traffic on the LAG "lag-3" is load-balanced in the subscriber context and is Vport-aware, while traffic on the LAG "lag-1" is disjoined from the subscriber context. On the LAG "lag-1", service hashing mechanisms apply; on the LAG "lag-3", subscriber management hashing mechanisms apply.

Note: An example topology with FPE-based PW ports might give the impression that load balancing on the external links connecting the BNG to the AGG node is performed in the subscriber context, based on subscriber ID or Vport. This is not the case.

To demonstrate the two-stage load balancing, first on "lag-3" and then on "lag-1", an example topology with subscribers instantiated on a FPE based PW-port is chosen.

Although the focus of this chapter is class- and weight-based Vport load balancing, using the same traffic patterns, in addition, the following two other types of load balancing are evaluated:

- load balancing based on Vports without class and weights
- load balancing based on subscribers (subscriber IDs that are internally assigned identifiers for each subscriber)

Comparing the three types of load balancing helps to form a better overall picture about load balancing traffic options in a subscriber management context.

Association between subscribers and Vports

Subscribers are associated with Vports during the subscriber authentication phase. In this example, RADIUS is used for authentication.

The table Table 24: Subscriber association with Vports shows the mapping between subscribers and Vport names, Vport configured rates (bandwidth), classes, weights, and transmission (Tx) rates in frames per second (FPS) of generated traffic (offered traffic).

Table 24: Subscriber association with Vports

Class	Vport	Weight	Vport rate in kb/s	Tx rate FPS	Subscriber
1	vport-1	15	9 000 000	900	1
1	vport-2	5	3 000 000	300	2, 3
1	vport-3	5	3 000 000	300	4, 5
1	vport-4	5	3 000 000	300	6, 7
2	vport-5	15	3 000 000	300	8
2	vport-6	5	1 000 000	100	9, 10

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Class	Vport	Weight	Vport rate in kb/s	Tx rate FPS	Subscriber
2	vport-7	5	1 000 000	100	11, 12
2	vport-8	5	1 000 000	100	13, 14
3	vport-9	7	1 400 000	140	15
3	vport-10	1	200 000	20	16, 17
3	vport-11	1	200 000	20	18, 19
3	vport-12	1	200 000	20	20, 21
3	vport-13	1	200 000	20	22, 23
3	vport-14	1	200 000	20	24, 25
3	vport-15	1	200 000	20	26, 27
3	vport-16	1	200 000	20	28, 29, 30, 31, 32

The Vports in SR have names from "vport-1" to "vport-16".

The class represents a bandwidth tier. Vports with similarly configured bandwidth values can be grouped into classes and can be load-balanced independently of Vports in other classes. For example, Vports in class 1 are distributed over the LAG member ports independently of the Vports in classes 2 and 3.

The weight is configured for each Vport. Within each class, the hashing algorithm tries to equalize the sum of weights across the LAG member ports.

The traffic generator sends a traffic stream for each subscriber. For example, the traffic generator sends 900 fps toward subscriber 1 associated with Vport "vport-1", while toward the subscribers 28, 29, 30, 31, and 32 associated with "vport-16", the traffic generator sends an aggregated rate of 20 fps, or 4 fps toward each subscriber. The packet size is uniform (200 bytes) across the subscribers.

The selected traffic rates represent the Vport bandwidth. In other words, the ratio of the traffic rates within each Vport is the same as the ratio between the configured Vport bandwidths, which allows for a more intuitive interpretation of the results.

Because this setup is run on PC-based sims, the selected offered rates are low (tens of fps rather than millions of fps). However, the rates are high enough for the purpose of this chapter.

For brevity, the subscribers in this table are represented by numerical values. In SR, the subscriber numbers are prepended with the string "ipoe -", for example, subscriber 1 in the table represents subscriber "ipoe-1" in the SR.

The table Table 24: Subscriber association with Vports shows that the configured Vport bandwidth and, consequently, the transmission rates are selected in a very unbalanced way. These unbalanced values are used to allow for a clearer interpretation of the results that are presented later in this chapter.

For example, traffic for "vport-9" is seven times higher than traffic in each of the Vports from 10 to 16. If the eight Vports from 9 to 16 were distributed without regarding their configured Vport bandwidth, then the most likely outcome would be that four Vports would be hashed to one link and four to the other. This outcome would result in a very uneven load balancing with traffic utilization on one link being more than twice the utilization on the other link.

Vport classes and weights ensure that uneven load balancing does not happen and that traffic is more evenly split across the two LAG member ports.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The figure Figure 362: Vport bandwidth distribution is a graphical representation of the Vport bandwidth in the table Table 24: Subscriber association with Vports, where the size of each rectangle represents the Vport bandwidth configured rate. Vports in different classes are represented in different colors. Vports are divided into smaller areas, where each area represents the bandwidth of subscribers associated with the Vport. For example, the size of Vport 1 (9 Gb/s) is three times the size of each of the Vports 2, 3, 4, and 5 (3 Gb/s), or 45 times the size of any of the Vports from 10 to 16 (200 Mb/s). The SR independently attempts to equalize traffic in class 1, class 2, and class 3.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Figure 362: Vport bandwidth distribution

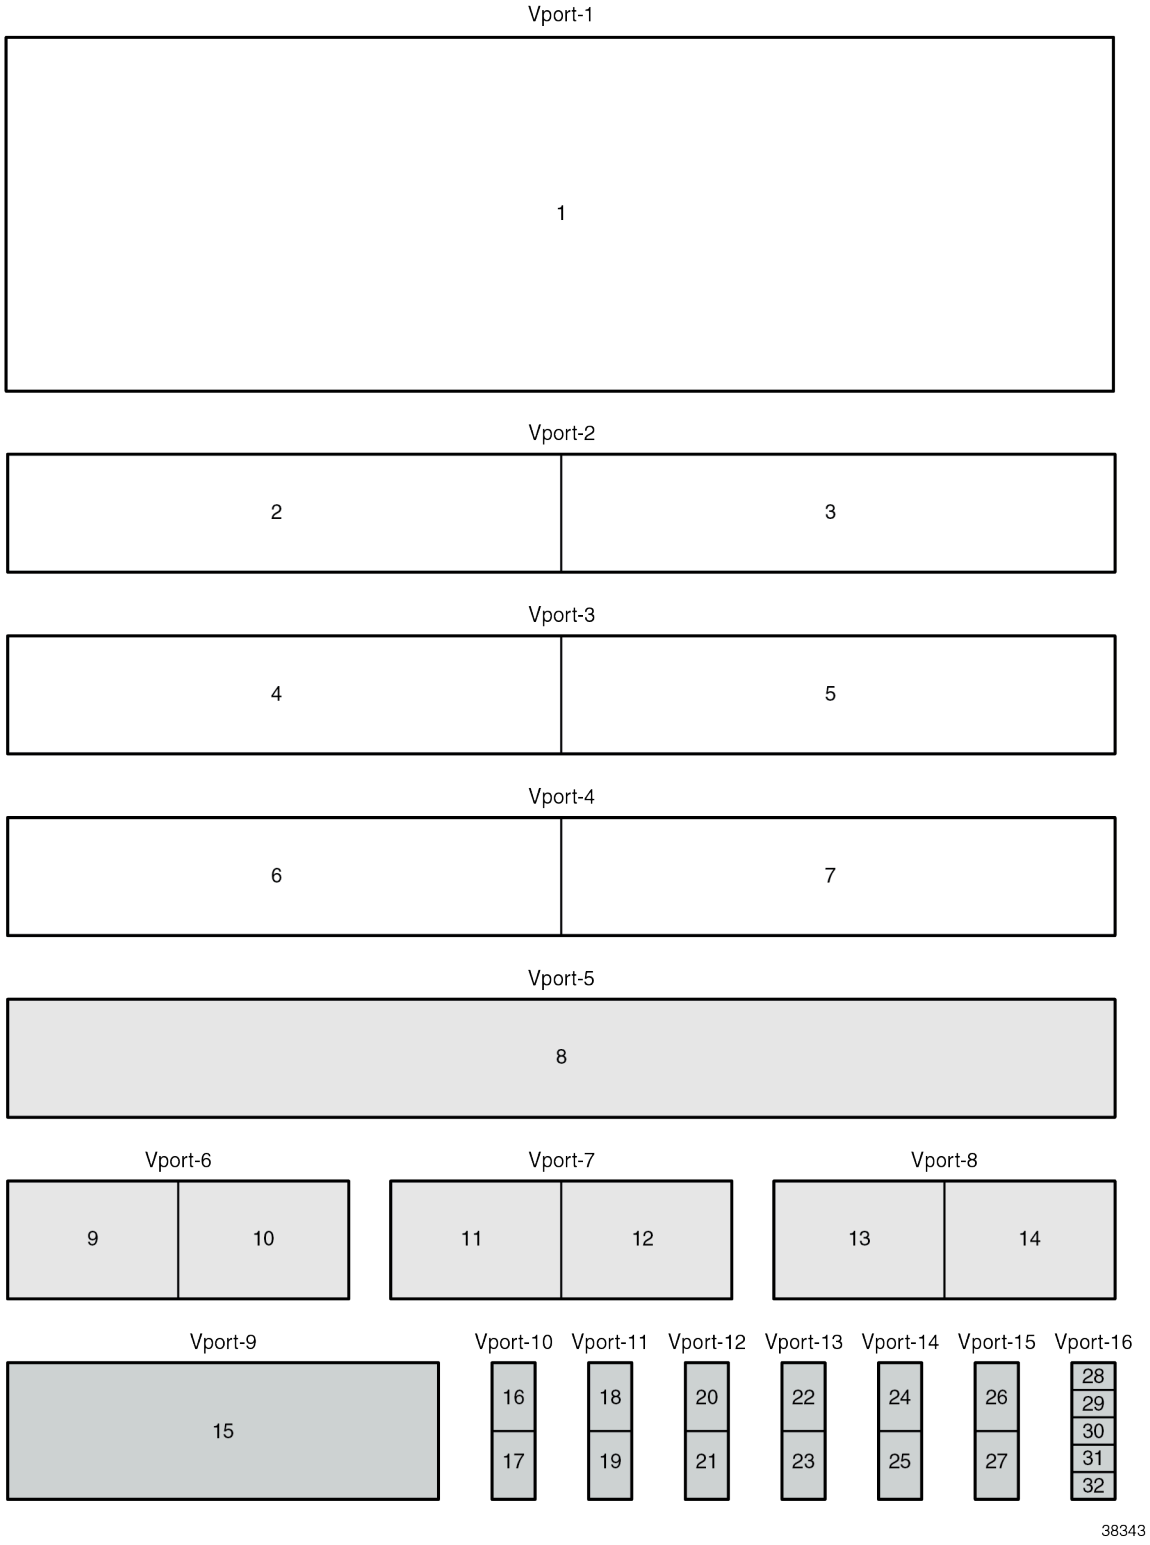

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Configuration

Complete configuration files for the AGG-1 and the BNG-1 node in this topology are provided in the appendix.

This section describes key configuration blocks and **show** output on the BNG.

The configuration section is split into the three following sections:

- FPE
- Vports
- LAG

FPE

The setup on the BNG requires FPE-based PXCs for the PW-port. An arbitrary choice is made to select a port-based PXC instead of the MAC- or internal-based PXC. The following two PXCs will become part of the PXC-based LAG on which the subscribers are instantiated.

```
[pr:/configure port-xc]
A:admin@bng-1#

    pxc 1 {
        admin-state enable
        description "pw-port lag 2,3"
        port-id 1/1/c4/2
    }
    pxc 2 {
        admin-state enable
        description "pw-port lag 2,3"
        port-id 1/1/c4/3
    }


A:admin@bng-1# /show port-xc

===============================================================================
Port Cross-Connect Information
===============================================================================
PXC    Admin     Oper      Port          Description
Id     State     State     Id
-------------------------------------------------------------------------------
1      Up        Up        1/1/c4/2      pw-port lag 2,3
2      Up        Up        1/1/c4/3      pw-port lag 2,3
-------------------------------------------------------------------------------
No. of PXCs: 2
===============================================================================
```

The logical PXC ports must be explicitly configured and enabled by the operator:

```
[pr:/configure]
A:admin@bng-1#
    port pxc-1.a {
        admin-state enable
    }
    port pxc-1.b {
        admin-state enable
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    }
    port pxc-2.a {
        admin-state enable
    }
    port pxc-2.b {
        admin-state enable
    }
```

The FPE configuration that is required for the PW-port is as follows:

```
[pr:/configure]
A:admin@bng-1#
    fwd-path-ext {
        sdp-id-range {
            start 17500
            end 17600
        }
        fpe 1 {
            description "pw-port on a lag"
            path {
                xc-lag-a "lag-2"
                xc-lag-b "lag-3"
            }
            application {
                pw-port-extension {
                }
            }
        }
    }

A:admin@bng-1#/show  fwd-path-ext fpe 1

===============================================================================
FPE Id: 1
===============================================================================
Description        : pw-port on a lag
Multi-Path         : Disabled
Path               : lag-2, lag-3
Pw Port Extension  : Enabled                          Oper     : up
Sub Mgmt Extension : Disabled                         Oper     : N/A
Vxlan              : Disabled                         Oper     : down
Segment-Routing V6 : Disabled
If-A Qos Policy    : default
If-B Qos Policy    : default
===============================================================================
```

Vports

The egress QoS hierarchy for the subscriber is the following:

- subscriber queues with their own rates and buffering capabilities
- the aggregate rate of the subscriber – the total traffic from all the queues of a subscriber cannot exceed this rate
- the aggregate rate of a group of subscribers connected to the same access node or a PON – this aggregate rate is controlled by a Vport
- the aggregate rate of subscribers connected to the same local port – this aggregate rate is controlled by a port-scheduler-policy

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Vports are configured on every member port of "lag-3". The LAG underlies the PW-port on which the subscribers are terminated. The member ports of "lag-3" are pxc-1.b and pxc-2.b.

The Vports are attached to the port scheduler via a port scheduler policy named "lag-3-pxc", which is shown in the following configuration.

Only the configuration of the pxc-1.b is shown. The configuration of port-2.b is identical to that for port pxc-1.b.

Each Vport is configured with an aggregate rate, the class, and the weight as described in the table Table 24: Subscriber association with Vports.

The mapping between the Vport and the subscriber is based on the `int-dest-id` string returned by RADIUS during subscriber authentication. For example, the RADIUS-returned string for subscriber "ipoe-1" is carried in the Nokia VSA `Alc-Int-Dest-Id-Str` with value "`vport-1`".

The Vport configuration and port-scheduler policy reference under pxc-1.b (a member port of "lag-3") is as follows:

```
pr:/configure port pxc-1.b]
A:admin@bng-1#
    admin-state enable
    description "PXC lag-3, pw-port termination"
    ethernet {
        access {
            egress {
                virtual-port "vport-1" {
                    aggregate-rate {
                        rate 9000000
                    }
                    host-match {
                        int-dest-id "vport-1" { }
                    }
                    lag-per-link-hash {
                        class 1
                        weight 15
                    }
                }
                virtual-port "vport-2" {
                    aggregate-rate {
                        rate 3000000
                    }
                    host-match {
                        int-dest-id "vport-2" { }
                    }
                    lag-per-link-hash {
                        class 1
                        weight 5
                    }
                }
                virtual-port "vport-3" {
                    aggregate-rate {
                        rate 3000000
                    }
                    host-match {
                        int-dest-id "vport-3" { }
                    }
                    lag-per-link-hash {
                        class 1
                        weight 5
                    }
                }
                virtual-port "vport-4" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    aggregate-rate {
                        rate 3000000
                    }
                    host-match {
                        int-dest-id "vport-4" { }
                    }
                    lag-per-link-hash {
                        class 1
                        weight 5
                    }
                }
                virtual-port "vport-5" {
                    aggregate-rate {
                        rate 3000000
                    }
                    host-match {
                        int-dest-id "vport-5" { }
                    }
                    lag-per-link-hash {
                        class 2
                        weight 15
                    }
                }
                virtual-port "vport-6" {
                    aggregate-rate {
                        rate 1000000
                    }
                    host-match {
                        int-dest-id "vport-6" { }
                    }
                    lag-per-link-hash {
                        class 2
                        weight 5
                    }
                }
                virtual-port "vport-7" {
                    aggregate-rate {
                        rate 1000000
                    }
                    host-match {
                        int-dest-id "vport-7" { }
                    }
                    lag-per-link-hash {
                        class 2
                        weight 5
                    }
                }
                virtual-port "vport-8" {
                    aggregate-rate {
                        rate 1000000
                    }
                    host-match {
                        int-dest-id "vport-8" { }
                    }
                    lag-per-link-hash {
                        class 2
                        weight 5
                    }
                }
                virtual-port "vport-9" {
                    aggregate-rate {
                        rate 1400000
                    }
                    host-match {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        int-dest-id "vport-9" { }
                    }
                    lag-per-link-hash {
                        class 3
                        weight 7
                    }
                }
                virtual-port "vport-10" {
                    aggregate-rate {
                        rate 200000
                    }
                    host-match {
                        int-dest-id "vport-10" { }
                    }
                    lag-per-link-hash {
                        class 3
                        weight 1
                    }
                }
                virtual-port "vport-11" {
                    aggregate-rate {
                        rate 200000
                    }
                    host-match {
                        int-dest-id "vport-11" { }
                    }
                    lag-per-link-hash {
                        class 3
                        weight 1
                    }
                }
                virtual-port "vport-12" {
                    aggregate-rate {
                        rate 200000
                    }
                    host-match {
                        int-dest-id "vport-12" { }
                    }
                    lag-per-link-hash {
                        class 3
                        weight 1
                    }
                }
                virtual-port "vport-13" {
                    aggregate-rate {
                        rate 200000
                    }
                    host-match {
                        int-dest-id "vport-13" { }
                    }
                    lag-per-link-hash {
                        class 3
                        weight 1
                    }
                }
                virtual-port "vport-14" {
                    aggregate-rate {
                        rate 200000
                    }
                    host-match {
                        int-dest-id "vport-14" { }
                    }
                    lag-per-link-hash {
                        class 3
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    weight 1
                }
            }
            virtual-port "vport-15" {
                aggregate-rate {
                    rate 200000
                }
                host-match {
                    int-dest-id "vport-15" { }
                }
                lag-per-link-hash {
                    class 3
                    weight 1
                }
            }
            virtual-port "vport-16" {
                aggregate-rate {
                    rate 200000
                }
                host-match {
                    int-dest-id "vport-16" { }
                }
                lag-per-link-hash {
                    class 3
                    weight 1
                }
            }
        }
    }
    egress {
        port-scheduler-policy {
            policy-name "lag-3-pxc"
        }
    }
}
```

The following command shows the Vport instances created under port pxc-1.b:

```
A:admin@bng-1# show port pxc-1.b vport

===============================================================================
Port pxc-1.b Access Egress vport
===============================================================================
VPort Name         : vport-1
Description        : (Not Specified)
Port Sched Policy  : None
Sched Policy       : None
Rate Limit         : 9000000
Limit Unused BW    : Disabled
Rate Modify        : disabled
Modify delta       : 0
Monitor Port Sched : Disabled
Lag PLHW class     : 1
Lag PLHW weight    : 15

Host-Matches
-------------------------------------------------------------------------------
Dest: vport-1
-------------------------------------------------------------------------------

VPort Name         : vport-2
Description        : (Not Specified)
Port Sched Policy  : None
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Sched Policy       : None
Rate Limit         : 3000000
Limit Unused BW    : Disabled
Rate Modify        : disabled
Modify delta       : 0
Monitor Port Sched : Disabled
Lag PLHW class     : 1
Lag PLHW weight    : 5

Host-Matches
-------------------------------------------------------------------------------
Dest: vport-2
-------------------------------------------------------------------------------

VPort Name         : vport-3
Description        : (Not Specified)
Port Sched Policy  : None
Sched Policy       : None
Rate Limit         : 3000000
Limit Unused BW    : Disabled
Rate Modify        : disabled
Modify delta       : 0
Monitor Port Sched : Disabled
Lag PLHW class     : 1
Lag PLHW weight    : 5

Host-Matches
-------------------------------------------------------------------------------
Dest: vport-3
-------------------------------------------------------------------------------
---snip---
```

The following command shows the association between the Vports and the subscribers for port pxc-1.b:

```
A:admin@bng-1# show port pxc-1.b vport associations

===============================================================================
Port pxc-1.b Access Egress vport
===============================================================================

-------------------------------------------------------------------------------
VPort "vport-1"
-------------------------------------------------------------------------------
svc-id : 10
 sap   : pw-1:1.1
 subscr: ipoe-1
 ip    : 10.10.0.16
 mac   : 00:14:01:00:00:01  pppoe-sid: N/A


-------------------------------------------------------------------------------
VPort "vport-2"
-------------------------------------------------------------------------------
svc-id : 10
 sap   : pw-1:1.2
 subscr: ipoe-2
 ip    : 10.10.0.21
 mac   : 00:14:01:00:00:02  pppoe-sid: N/A
svc-id : 10
 sap   : pw-1:1.3
 subscr: ipoe-3
 ip    : 10.10.0.41
 mac   : 00:14:01:00:00:03  pppoe-sid: N/A
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
VPort "vport-3"
-------------------------------------------------------------------------------
svc-id : 10
 sap   : pw-1:1.4
 subscr: ipoe-4
 ip    : 10.10.0.24
 mac   : 00:14:01:00:00:04  pppoe-sid: N/A
svc-id : 10
 sap   : pw-1:1.5
 subscr: ipoe-5
 ip    : 10.10.0.38
 mac   : 00:14:01:00:00:05  pppoe-sid: N/A
---snip---
```

The following configuration shows that each subscriber has only one egress queue, which is port-parented. This queue is configured in the egress SAP policy, which is then referenced in the SLA profile assigned to the subscriber during the authentication phase.

Port-schedule policy "lag-3-pxc" is also configured under the QoS hierarchy, using default values in this case (which is why there is nothing explicitly configured under **port-scheduler-policy**).

```
[pr:/configure qos]
A:admin@bng-1#
    sap-egress "sap-egess-1" {
        policy-id 2
        queue 1 {
            port-parent {
            }
        }
    }
    port-scheduler-policy "lag-3-pxc" {
    }
```

The SLA profile for the subscriber must include **vport** as the port-parent location, otherwise the subscriber would not be associated with the Vport:

```
[pr:/configure subscriber-mgmt sla-profile "sla-profile-1"]
A:admin@bng-1#
    egress {
        qos {
            sap-egress {
                policy-name "sap-egess-1"
                port-parent-location vport

            }
        }
    }
```

There are 32 subscribers instantiated in the BNG:

```
A:admin@bng-1# show service active-subscribers summary

===============================================================================
Active Subscriber table summary
===============================================================================
Total Count       : 32
===============================================================================
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The following **show** command displays the subscriber hierarchy:

```
A:admin@bng-1# show service  active-subscribers hierarchy

===============================================================================
Active Subscribers Hierarchy
===============================================================================
-- ipoe-1
   (sub-profile-1)
   |
   +-- sap:[pw-1:1.1] - sla:sla-profile-1
       |
       +-- IPOE-session - mac:00:14:01:00:00:01 - svc:10
           |
           +-- 10.10.0.25 - DHCP

-- ipoe-2
   (sub-profile-1)
   |
   +-- sap:[pw-1:1.2] - sla:sla-profile-1
       |
       +-- IPOE-session - mac:00:14:01:00:00:02 - svc:10
           |
           +-- 10.10.0.26 - DHCP

-- ipoe-3
   (sub-profile-1)
   |
   +-- sap:[pw-1:1.3] - sla:sla-profile-1
       |
       +-- IPOE-session - mac:00:14:01:00:00:03 - svc:10
           |
           +-- 10.10.0.27 - DHCP

---snip---

-------------------------------------------------------------------------------
Number of active subscribers : 32
Flags: (N) = the host or the managed route is in non-forwarding state
===============================================================================
---snip---
```

Some of the additional **show** commands that can be used to display subscriber states are:

```
show service active-subscribers detail
show service id "vport-hashing" subscriber-hosts detail
show service id "vport-hashing" ipoe session detail
show service id "vport-hashing" dhcp lease-state detail
```

The following **show** commands display information related to the QoS hierarchy:

```
show qos scheduler-hierarchy port pxc-1.b detail
show qos agg-rate port pxc-1.b vport "vport-1" detail
```

LAG

The two LAGs examined in this chapter are:

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

- PXC LAG "lag-3", which underlies the PW-port where the subscribers are terminated. This LAG is implicitly applied to the ESM context(directly via FPE, indirectly via PW-port and PW capture SAP) where load balancing is performed based on Vports.
- LAG "lag-1" connects the BNG to the aggregation node "AGG-1". This LAG is applied in the service context. Load balancing on this LAG is performed based on the service context.

This chapter examines load balancing over member links for both LAGs.

PXC LAG "lag-3" (ESM)

The configuration for "lag-3" is the following:

```
[pr:/configure lag "lag-3"]
A:admin@bng-1#
    admin-state enable
    description "fpe pw-port pxc lag - termination side"
    mode hybrid
    max-ports 64
    access {
        per-fp-ing-queuing true
        per-fp-egr-queuing true
        per-fp-sap-instance true
        adapt-qos {
            mode link
        }
    }
    per-link-hash {
        weighted {
            subscriber-hash-mode vport
        }
    }
    port pxc-1.b {
    }
    port pxc-2.b {
    }
```

The **per-link-hash** command enables weighted Vport hashing in the subscriber context.

If classes and weights for Vports are not configured, hashing is still performed per Vport, in which case all Vports are treated equally, regardless of the bandwidth.

With the **per-fp-sap-instance** command enabled, the system allocates only one SAP instance per subscriber per forwarding complex.

For example, if a LAG has two member ports on the same forwarding complex, then for each subscriber, only one SAP instance is allocated (as opposed to one SAP instance on each member port). The result of this allocation is improved scaling because the number of SAPs per forwarding complex is finite. The scaling improvement can be verified with the following commands.

Before any subscriber is instantiated, the number of subscriber hosts and SLA profile instances is 0. A number of SAPs are already allocated (three in this case), related to non-subscriber interfaces (EVPN VPWS, regular interfaces, and so on).

Note: The setup runs on PC-based simulators where total scaling figures are reduced.

```
A:admin@bng-1# tools dump resource-usage card "1" fp 1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
===============================================================================
Resource Usage Information for Card Slot #1 FP #1
===============================================================================
                                                  Total  Allocated       Free
-------------------------------------------------------------------------------

                            Subscriber Hosts -     262143          0     262143

            Subscriber SLA Profile Instances |     131071          0     131071
                               SAP Instances |      98303          3      98300
===============================================================================
```

After 32 subscribers are instantiated on the PW-port and on "lag-3", only 32 SAP instances are allocated:

35 allocated SAP instances − 3 residual SAP instances = 32 SAP instances (for 32 subscriber hosts)

The following command shows that 32 subscriber hosts and 35 SAP instances are allocated:

```
A:admin@bng-1#  /tools dump resource-usage card 1 fp 1

===============================================================================
Resource Usage Information for Card Slot #1 FP #1
===============================================================================
                                                  Total  Allocated       Free
-------------------------------------------------------------------------------
                            Subscriber Hosts -     262143         32     262111

            Subscriber SLA Profile Instances |     131071         32     131039
                               SAP Instances |      98303         35      98268

===============================================================================
```

The **per-fp-ing-queuing** and **per-fp-egr-queuing** commands have similar effect for queues, although the **per-fp-egr-queuing true** command is the only mode of operation in subscriber management on a LAG.

The **access adapt-qos mode link** command instructs the system to apply the configured Vport bandwidth to every member port of a LAG. This command is required in per-Vport load balancing.

The alternative mode is **port-fair**, where the configured Vport bandwidth is divided across the member ports in the LAG (each member port gets a share of the configured bandwidths). Per-Vport load balancing is not recommended in this case due to likely overshoot of the Vport assigned bandwidth to each member port, which would result in unnecessary packet drops.

Examining Vport load balancing effects

For a better understanding of what per-Vport load balancing means, the results for three cases are compared:

- traffic utilization on LAG member links with weighted (for class and weight) per-Vport load balancing
- traffic utilization on LAG member links with Vport load balancing but without classes and weights configured
- traffic utilization on LAG member links with per-subscriber load balancing (Vport load balancing is disabled)

The expectation is that there is a higher degree of load balancing in the first case than in the next two cases. This different degree for load balancing is because of the unequal nature of the generated traffic,

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

where some subscribers send significantly more traffic than others, and the inability of the latter two cases to deal with this lack of balance.

Load balancing per Vport with classes and weights

The configuration file for BNG-1 can be found here.

In this setup, classes and weights configured under Vports are considered in the hashing algorithm. The following command shows how Vports are distributed over the LAG member ports. Although all Vports are instantiated on both member ports of "lag-3", this command shows how the actual traffic for each Vport is distributed over the member links of "lag-3":

```
A:admin@bng-1# show lag "lag-3" associations per-link-hash vport

=====================================================
VPort Associations
=====================================================
Vport Name                        Active Link
-----------------------------------------------------
vport-1                           pxc-1.b
vport-2                           pxc-2.b
vport-3                           pxc-2.b
vport-4                           pxc-2.b
vport-5                           pxc-1.b
vport-6                           pxc-2.b
vport-7                           pxc-2.b
vport-8                           pxc-2.b
vport-9                           pxc-2.b
vport-10                          pxc-1.b
vport-11                          pxc-1.b
vport-12                          pxc-1.b
vport-13                          pxc-1.b
vport-14                          pxc-1.b
vport-15                          pxc-1.b
vport-16                          pxc-1.b
=====================================================
Number of VPort associations: 16
=====================================================
```

The following command shows how subscriber traffic is distributed over the member ports, based on each subscriber's association with a Vport. For example, the subscriber "ipoe-1" is associated with "vport-1". The output of the previous command shows that "vport-1" is hashed to port pxc-1.b and, therefore, the LAG active link (as shown below) for the subscriber "ipoe-1" is pxc-1.b.

```
A:admin@bng-1# show service  id "vport-hashing" subscriber-hosts detail

===============================================================================
Subscriber Host table
===============================================================================
Sap
  IP Address
    MAC Address                PPPoE-SID       Origin          Fwding State
      Subscriber
-------------------------------------------------------------------------------
[pw-1:1.1]
  10.10.0.42
    00:14:01:00:00:01          N/A             DHCP            Fwding
      ipoe-1
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : vport-1
LAG Active Link       : pxc-1.b
Acct-Session-Id       : B496FF00000070639A19D0
Acct-Q-Inst-Session-Id: B496FF00000071639A19D0
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:1.2]
  10.10.0.43
    00:14:01:00:00:02          N/A             DHCP            Fwding
      ipoe-2
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : vport-2
LAG Active Link       : pxc-2.b
Acct-Session-Id       : B496FF00000072639A19D0
Acct-Q-Inst-Session-Id: B496FF00000073639A19D0
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
---snip---
```

The table Table 25: Vport distribution over member ports repeats the Vport distribution shown in the table Table 24: Subscriber association with Vports, but adds a column with the active member link for each Vport:

Table 25: Vport distribution over member ports

Class	Vport	Weight	Vport rate in kb/s	Tx rate FPS	Subscriber	Member link
1	vport-1	15	9 000 000	900	1	**pxc-1.b**
1	vport-2	5	3 000 000	300	2, 3	**pxc-2.b**
1	vport-3	5	3 000 000	300	4, 5	**pxc-2.b**
1	vport-4	5	3 000 000	300	6, 7	**pxc-2.b**
2	vport-5	15	3 000 000	300	8	**pxc-1.b**

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Class	Vport	Weight	Vport rate in kb/s	Tx rate FPS	Subscriber	Member link
2	vport-6	5	1 000 000	100	9, 10	**pxc-2.b**
2	vport-7	5	1 000 000	100	11, 12	**pxc-2.b**
2	vport-8	5	1 000 000	100	13, 14	**pxc-2.b**
3	vport-9	7	1 400 000	140	15	**pxc-2.b**
3	vport-10	1	200 000	20	16, 17	**pxc-1.b**
3	vport-11	1	200 000	20	18, 19	**pxc-1.b**
3	vport-12	1	200 000	20	20, 21	**pxc-1.b**
3	vport-13	1	200 000	20	22, 23	**pxc-1.b**
3	vport-14	1	200 000	20	24, 25	**pxc-1.b**
3	vport-15	1	200 000	20	26, 27	**pxc-1.b**
3	vport-16	1	200 000	20	28, 29, 30, 31, 32	**pxc-1.b**

To verify the Vport distribution over member ports, the traffic on the member ports of "lag-3" can be monitored.

An efficient way to observe utilization on ports in an SR is to run a **monitor port** command that periodically collects the port data (forwarded octets and packets) and displays them in various forms (raw counts or as rates). This **monitor port** command is used on regular ports on "lag-1" in the next section. However, the **monitor port** command cannot be used on PXC ports.

Because the **monitor port** command is not supported on PXC ports, a dedicated **monitor-lag** command is developed via pysros. The output of this command shows the distribution of traffic per LAG member port, per Vport, and per subscriber. Rates are displayed in percentages of the total downstream traffic sent toward all subscribers. During each command execution, the data is, by default, collected three times in 10 second intervals (this frequency is configurable).

This example demonstrates that bandwidth is evenly distributed over the two LAG member ports (approximately 50% on pxc-1.b port and 50% on pxc-2.b port). Subscriber rates and subscriber to Vport mappings are also consistent with the expectations according to the table Table 25: Vport distribution over member ports .

The output shows that in each class, the high-bandwidth Vports ("vport-1" in class 1, "vport-5" in class 2, and "vport-9" in class 3) are counterbalanced with lower-bandwidth Vports from the same class that are hashed to the opposite link. In this way, load balancing is distributed in 50/50 ratio over the two member ports on the "lag-3": both member ports pxc-1.b and pxc-2.b get approximately 50% of the load.

```
 A:admin@bng-1# monitor-lag -l lag-3 -i 10 -r 3


-------------------------------------------------------------------------
                                  Run 1
-------------------------------------------------------------------------
               Object               |           BW % of Total
-------------------------------------------------------------------------


pxc-1.b-----------------------------------49.41
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
   |--------vport-1-----------------------------------34.41
   |           |--------ipoe-1-----------------------------------34.41
   |--------vport-2-----------------------------------0.00
   |--------vport-3-----------------------------------0.00
   |--------vport-4-----------------------------------0.00
   |--------vport-5-----------------------------------11.47
   |           |--------ipoe-8-----------------------------------11.47
   |--------vport-6-----------------------------------0.00
   |--------vport-7-----------------------------------0.00
   |--------vport-8-----------------------------------0.00
   |--------vport-9-----------------------------------0.00
   |--------vport-10----------------------------------0.59
   |           |--------ipoe-16----------------------------------0.29
   |           |--------ipoe-17----------------------------------0.29
   |--------vport-11----------------------------------0.59
   |           |--------ipoe-19----------------------------------0.29
   |           |--------ipoe-18----------------------------------0.29
   |--------vport-12----------------------------------0.59
   |           |--------ipoe-20----------------------------------0.29
   |           |--------ipoe-21----------------------------------0.29
   |--------vport-13----------------------------------0.59
   |           |--------ipoe-23----------------------------------0.29
   |           |--------ipoe-22----------------------------------0.29
   |--------vport-14----------------------------------0.59
   |           |--------ipoe-25----------------------------------0.29
   |           |--------ipoe-24----------------------------------0.29
   |--------vport-15----------------------------------0.59
   |           |--------ipoe-27----------------------------------0.29
   |           |--------ipoe-26----------------------------------0.29
   |--------vport-16----------------------------------0.00
               |--------ipoe-32----------------------------------0.00
               |--------ipoe-30----------------------------------0.00
               |--------ipoe-31----------------------------------0.00
               |--------ipoe-29----------------------------------0.00
               |--------ipoe-28----------------------------------0.00


pxc-2.b------------------------------------50.59
   |--------vport-1-----------------------------------0.00
   |--------vport-2-----------------------------------11.47
   |           |--------ipoe-3-----------------------------------5.88
   |           |--------ipoe-2-----------------------------------5.59
   |--------vport-3-----------------------------------11.47
   |           |--------ipoe-4-----------------------------------5.88
   |           |--------ipoe-5-----------------------------------5.59
   |--------vport-4-----------------------------------11.18
   |           |--------ipoe-7-----------------------------------5.59
   |           |--------ipoe-6-----------------------------------5.59
   |--------vport-5-----------------------------------0.00
   |--------vport-6-----------------------------------3.53
   |           |--------ipoe-10----------------------------------1.76
   |           |--------ipoe-9-----------------------------------1.76
   |--------vport-7-----------------------------------3.53
   |           |--------ipoe-12----------------------------------1.76
   |           |--------ipoe-11----------------------------------1.76
   |--------vport-8-----------------------------------4.12
   |           |--------ipoe-14----------------------------------2.06
   |           |--------ipoe-13----------------------------------2.06
   |--------vport-9-----------------------------------5.29
   |           |--------ipoe-15----------------------------------5.29
   |--------vport-10----------------------------------0.00
   |--------vport-11----------------------------------0.00
   |--------vport-12----------------------------------0.00
   |--------vport-13----------------------------------0.00
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
   |--------vport-14-----------------------------------0.00
   |--------vport-15-----------------------------------0.00
   |--------vport-16-----------------------------------0.00


-------------------------------------------------------------------------
                                Run 2
-------------------------------------------------------------------------
              Object               |          BW % of Total
-------------------------------------------------------------------------


pxc-1.b-------------------------------------49.99
   |--------vport-1------------------------------------33.58
   |           |--------ipoe-1-------------------------------------33.58
   |--------vport-2------------------------------------0.00
   |--------vport-3------------------------------------0.00
   |--------vport-4------------------------------------0.00
   |--------vport-5------------------------------------11.20
   |           |--------ipoe-8-------------------------------------11.20
   |--------vport-6------------------------------------0.00
   |--------vport-7------------------------------------0.00
   |--------vport-8------------------------------------0.00
   |--------vport-9------------------------------------0.00
   |--------vport-10-----------------------------------0.74
   |           |--------ipoe-16------------------------------------0.37
   |           |--------ipoe-17------------------------------------0.37
   |--------vport-11-----------------------------------0.74
   |           |--------ipoe-19------------------------------------0.37
   |           |--------ipoe-18------------------------------------0.37
   |--------vport-12-----------------------------------0.74
   |           |--------ipoe-20------------------------------------0.37
   |           |--------ipoe-21------------------------------------0.37
   |--------vport-13-----------------------------------0.74
   |           |--------ipoe-23------------------------------------0.37
   |           |--------ipoe-22------------------------------------0.37
   |--------vport-14-----------------------------------0.74
   |           |--------ipoe-25------------------------------------0.37
   |           |--------ipoe-24------------------------------------0.37
   |--------vport-15-----------------------------------0.74
   |           |--------ipoe-27------------------------------------0.37
   |           |--------ipoe-26------------------------------------0.37
   |--------vport-16-----------------------------------0.75
               |--------ipoe-32------------------------------------0.15
               |--------ipoe-30------------------------------------0.15
               |--------ipoe-31------------------------------------0.15
               |--------ipoe-29------------------------------------0.15
               |--------ipoe-28------------------------------------0.15


pxc-2.b-------------------------------------50.01
   |--------vport-1------------------------------------0.00
   |--------vport-2------------------------------------11.19
   |           |--------ipoe-3-------------------------------------5.60
   |           |--------ipoe-2-------------------------------------5.59
   |--------vport-3------------------------------------11.20
   |           |--------ipoe-4-------------------------------------5.60
   |           |--------ipoe-5-------------------------------------5.60
   |--------vport-4------------------------------------11.20
   |           |--------ipoe-7-------------------------------------5.60
   |           |--------ipoe-6-------------------------------------5.59
   |--------vport-5------------------------------------0.00
   |--------vport-6------------------------------------3.73
   |           |--------ipoe-10------------------------------------1.87
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
   |           |--------ipoe-9-----------------------------------1.87
   |--------vport-7-------------------------------------3.73
   |           |--------ipoe-12----------------------------------1.87
   |           |--------ipoe-11----------------------------------1.87
   |--------vport-8-------------------------------------3.73
   |           |--------ipoe-14----------------------------------1.86
   |           |--------ipoe-13----------------------------------1.86
   |--------vport-9-------------------------------------5.22
   |           |--------ipoe-15----------------------------------5.22
   |--------vport-10------------------------------------0.00
   |--------vport-11------------------------------------0.00
   |--------vport-12------------------------------------0.00
   |--------vport-13------------------------------------0.00
   |--------vport-14------------------------------------0.00
   |--------vport-15------------------------------------0.00
   |--------vport-16------------------------------------0.00


-------------------------------------------------------------------------
                               Run 3
-------------------------------------------------------------------------
              Object               |           BW % of Total
-------------------------------------------------------------------------


pxc-1.b-------------------------------------**50.01**
   |--------vport-1-------------------------------------33.58
   |           |--------ipoe-1-----------------------------------33.58
   |--------vport-2-------------------------------------0.00
   |--------vport-3-------------------------------------0.00
   |--------vport-4-------------------------------------0.00
   |--------vport-5-------------------------------------11.19
   |           |--------ipoe-8-----------------------------------11.19
   |--------vport-6-------------------------------------0.00
   |--------vport-7-------------------------------------0.00
   |--------vport-8-------------------------------------0.00
   |--------vport-9-------------------------------------0.00
   |--------vport-10------------------------------------0.74
   |           |--------ipoe-16----------------------------------0.37
   |           |--------ipoe-17----------------------------------0.37
   |--------vport-11------------------------------------0.74
   |           |--------ipoe-19----------------------------------0.37
   |           |--------ipoe-18----------------------------------0.37
   |--------vport-12------------------------------------0.75
   |           |--------ipoe-20----------------------------------0.38
   |           |--------ipoe-21----------------------------------0.38
   |--------vport-13------------------------------------0.75
   |           |--------ipoe-23----------------------------------0.38
   |           |--------ipoe-22----------------------------------0.38
   |--------vport-14------------------------------------0.75
   |           |--------ipoe-25----------------------------------0.38
   |           |--------ipoe-24----------------------------------0.38
   |--------vport-15------------------------------------0.75
   |           |--------ipoe-27----------------------------------0.38
   |           |--------ipoe-26----------------------------------0.38
   |--------vport-16------------------------------------0.74
               |--------ipoe-32----------------------------------0.15
               |--------ipoe-30----------------------------------0.15
               |--------ipoe-31----------------------------------0.15
               |--------ipoe-29----------------------------------0.15
               |--------ipoe-28----------------------------------0.15


pxc-2.b-------------------------------------**49.99**
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
   |--------vport-1-------------------------------------0.00
   |--------vport-2-------------------------------------11.19
   |           |--------ipoe-3-------------------------------------5.60
   |           |--------ipoe-2-------------------------------------5.60
   |--------vport-3-------------------------------------11.19
   |           |--------ipoe-4-------------------------------------5.60
   |           |--------ipoe-5-------------------------------------5.60
   |--------vport-4-------------------------------------11.19
   |           |--------ipoe-7-------------------------------------5.60
   |           |--------ipoe-6-------------------------------------5.60
   |--------vport-5-------------------------------------0.00
   |--------vport-6-------------------------------------3.73
   |           |--------ipoe-10------------------------------------1.86
   |           |--------ipoe-9-------------------------------------1.86
   |--------vport-7-------------------------------------3.73
   |           |--------ipoe-12------------------------------------1.87
   |           |--------ipoe-11------------------------------------1.86
   |--------vport-8-------------------------------------3.74
   |           |--------ipoe-14------------------------------------1.87
   |           |--------ipoe-13------------------------------------1.87
   |--------vport-9-------------------------------------5.22
   |           |--------ipoe-15------------------------------------5.22
   |--------vport-10------------------------------------0.00
   |--------vport-11------------------------------------0.00
   |--------vport-12------------------------------------0.00
   |--------vport-13------------------------------------0.00
   |--------vport-14------------------------------------0.00
   |--------vport-15------------------------------------0.00
   |--------vport-16------------------------------------0.00
```

Another way to verify the Vport distribution over member ports is to run the following command:

```
show port "pxc-1.b" vport statistics
```

The output of this command shows the number of forwarded octets and packets. By running this command periodically, and with the help of some arithmetic, it can be deduced how much traffic is forwarded by each port.

Load balancing per Vport without classes and weights

In this example, Vport classes and weights are all set to the same value. This scenario is the same as one where classes and weight are not used at all. Traffic is still load-balanced per Vport but with no regard to the bandwidth of each Vport. In other words, all Vports are treated equally, contrary to the fact that their configured bandwidth is very different.

The configuration file for BNG-1 can be found here .

The same commands are used as in the previous example. The output of the following command shows that the ports are hashed very differently than before, with Vports evenly distributed between the two member ports (eight Vports hashed to pxc-1.b and eight Vports hashed to pxc-2.b).

```
A:admin@bng-1# show lag "lag-3" associations per-link-hash vport

=====================================================
VPort Associations
=====================================================
Vport Name                      Active Link
-----------------------------------------------------
vport-1                         pxc-1.b
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
vport-2                          pxc-2.b
vport-3                          pxc-1.b
vport-4                          pxc-2.b
vport-5                          pxc-2.b
vport-6                          pxc-1.b
vport-7                          pxc-1.b
vport-8                          pxc-2.b
vport-9                          pxc-1.b
vport-10                         pxc-2.b
vport-11                         pxc-1.b
vport-12                         pxc-2.b
vport-13                         pxc-1.b
vport-14                         pxc-2.b
vport-15                         pxc-1.b
vport-16                         pxc-2.b
=====================================================
Number of VPort associations: 16
=====================================================
```

The subscriber association with Vports and member ports is consistent with the output of the preceding command:

```
A:admin@bng-1# show service  id "vport-hashing" subscriber-hosts detail

===============================================================================
Subscriber Host table
===============================================================================
Sap
  IP Address
    MAC Address                PPPoE-SID       Origin          Fwding State
      Subscriber
-------------------------------------------------------------------------------
[pw-1:1.1]
  10.10.0.74
    00:14:01:00:00:01          N/A             DHCP            Fwding
      ipoe-1
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : vport-1
LAG Active Link       : pxc-1.b
Acct-Session-Id       : B496FF000000D0639A24FD
Acct-Q-Inst-Session-Id: B496FF000000D1639A24FD
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
[pw-1:1.2]
  10.10.0.75
    00:14:01:00:00:02          N/A             DHCP            Fwding
      ipoe-2
-------------------------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : vport-2
LAG Active Link       : pxc-2.b
Acct-Session-Id       : B496FF000000D4639A24FD
Acct-Q-Inst-Session-Id: B496FF000000D5639A24FD
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-------------------------------------------------------------------------------
---snip---
```

The table Table 26: Vport distribution over member ports repeats the Vport distribution over member ports from the table Table 24: Subscriber association with Vports while adding a column with the active member link for each Vport.

Table 26: Vport distribution over member ports

Class	Vport	Weight	Vport rate in kb/s	Tx rate FPS	Subscriber	Member link
1	Vport-1	15	9 000 000	900	1	pxc-1.b
1	Vport-2	5	3 000 000	300	2, 3	pxc-2.b
1	Vport-3	5	3 000 000	300	4, 5	pxc-1.b
1	Vport-4	5	3 000 000	300	6, 7	pxc-2.b
2	Vport-5	15	3 000 000	300	8	pxc-2.b
2	Vport-6	5	1 000 000	100	9, 10	pxc-1.b
2	Vport-7	5	1 000 000	100	11, 12	pxc-1.b
2	Vport-8	5	1 000 000	100	13, 14	pxc-2.b
3	Vport-9	7	1 400 000	140	15	pxc-1.b
3	Vport-10	1	200 000	20	16, 17	pxc-2.b
3	Vport-11	1	200 000	20	18, 19	pxc-1.b
3	Vport-12	1	200 000	20	20, 21	pxc-2.b
3	Vport-13	1	200 000	20	22, 23	pxc-1.b
3	Vport-14	1	200 000	20	24, 25	pxc-2.b
3	Vport-15	1	200 000	20	26, 27	pxc-2.b
3	Vport-16	1	200 000	20	28, 29, 30, 31, 32	pxc-2.b

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

The output of the **monitor-lag** command for "lag-3" shows unequal load balancing: member port pxc-1.b receives approximately 60% of the load and member port pxc-2.b only 40%.

```
A:admin@bng-1# monitor-lag -l lag-3 -i 10 -r 3


-------------------------------------------------------------------------
                                Run 1
-------------------------------------------------------------------------
              Object               |           BW % of Total
-------------------------------------------------------------------------


pxc-1.b-----------------------------------**60.29**
   |--------vport-1-----------------------------------34.49
   |           |--------ipoe-1-------------------------------------34.49
   |--------vport-2-----------------------------------0.00
   |--------vport-3-----------------------------------11.30
   |           |--------ipoe-4-------------------------------------5.51
   |           |--------ipoe-5-------------------------------------5.80
   |--------vport-4-----------------------------------0.00
   |--------vport-5-----------------------------------0.00
   |--------vport-6-----------------------------------4.06
   |           |--------ipoe-10------------------------------------2.03
   |           |--------ipoe-9-------------------------------------2.03
   |--------vport-7-----------------------------------3.48
   |           |--------ipoe-12------------------------------------1.74
   |           |--------ipoe-11------------------------------------1.74
   |--------vport-8-----------------------------------0.00
   |--------vport-9-----------------------------------5.22
   |           |--------ipoe-15------------------------------------5.22
   |--------vport-10----------------------------------0.00
   |--------vport-11----------------------------------0.58
   |           |--------ipoe-19------------------------------------0.29
   |           |--------ipoe-18------------------------------------0.29
   |--------vport-12----------------------------------0.00
   |--------vport-13----------------------------------0.58
   |           |--------ipoe-23------------------------------------0.29
   |           |--------ipoe-22------------------------------------0.29
   |--------vport-14----------------------------------0.00
   |--------vport-15----------------------------------0.58
   |           |--------ipoe-27------------------------------------0.29
   |           |--------ipoe-26------------------------------------0.29
   |--------vport-16----------------------------------0.00


pxc-2.b-----------------------------------**39.71**
   |--------vport-1-----------------------------------0.00
   |--------vport-2-----------------------------------11.30
   |           |--------ipoe-3-------------------------------------5.51
   |           |--------ipoe-2-------------------------------------5.80
   |--------vport-3-----------------------------------0.00
   |--------vport-4-----------------------------------11.30
   |           |--------ipoe-7-------------------------------------5.51
   |           |--------ipoe-6-------------------------------------5.80
   |--------vport-5-----------------------------------11.30
   |           |--------ipoe-8-------------------------------------11.30
   |--------vport-6-----------------------------------0.00
   |--------vport-7-----------------------------------0.00
   |--------vport-8-----------------------------------3.77
   |           |--------ipoe-14------------------------------------2.03
   |           |--------ipoe-13------------------------------------1.74
   |--------vport-9-----------------------------------0.00
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
   |--------vport-10-----------------------------------0.87
   |           |--------ipoe-16-----------------------------------0.58
   |           |--------ipoe-17-----------------------------------0.29
   |--------vport-11-----------------------------------0.00
   |--------vport-12-----------------------------------0.58
   |           |--------ipoe-20-----------------------------------0.29
   |           |--------ipoe-21-----------------------------------0.29
   |--------vport-13-----------------------------------0.00
   |--------vport-14-----------------------------------0.58
   |           |--------ipoe-25-----------------------------------0.29
   |           |--------ipoe-24-----------------------------------0.29
   |--------vport-15-----------------------------------0.00
   |--------vport-16-----------------------------------0.00
               |--------ipoe-32-----------------------------------0.00
               |--------ipoe-30-----------------------------------0.00
               |--------ipoe-31-----------------------------------0.00
               |--------ipoe-29-----------------------------------0.00
               |--------ipoe-28-----------------------------------0.00


------------------------------------------------------------------------
                                Run 2
------------------------------------------------------------------------
              Object               |           BW % of Total
------------------------------------------------------------------------


pxc-1.b-------------------------------------**59.70**
   |--------vport-1------------------------------------33.58
   |           |--------ipoe-1------------------------------------33.58
   |--------vport-2------------------------------------0.00
   |--------vport-3------------------------------------11.20
   |           |--------ipoe-4------------------------------------5.60
   |           |--------ipoe-5------------------------------------5.60
   |--------vport-4------------------------------------0.00
   |--------vport-5------------------------------------0.00
   |--------vport-6------------------------------------3.73
   |           |--------ipoe-10-----------------------------------1.87
   |           |--------ipoe-9------------------------------------1.87
   |--------vport-7------------------------------------3.73
   |           |--------ipoe-12-----------------------------------1.87
   |           |--------ipoe-11-----------------------------------1.87
   |--------vport-8------------------------------------0.00
   |--------vport-9------------------------------------5.23
   |           |--------ipoe-15-----------------------------------5.23
   |--------vport-10-----------------------------------0.00
   |--------vport-11-----------------------------------0.74
   |           |--------ipoe-19-----------------------------------0.37
   |           |--------ipoe-18-----------------------------------0.37
   |--------vport-12-----------------------------------0.00
   |--------vport-13-----------------------------------0.74
   |           |--------ipoe-23-----------------------------------0.37
   |           |--------ipoe-22-----------------------------------0.37
   |--------vport-14-----------------------------------0.00
   |--------vport-15-----------------------------------0.74
   |           |--------ipoe-27-----------------------------------0.37
   |           |--------ipoe-26-----------------------------------0.37
   |--------vport-16-----------------------------------0.00


pxc-2.b-------------------------------------**40.30**
   |--------vport-1------------------------------------0.00
   |--------vport-2------------------------------------11.19
   |           |--------ipoe-3------------------------------------5.60
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
   |           |--------ipoe-2-------------------------------------5.59
   |--------vport-3-----------------------------------0.00
   |--------vport-4-----------------------------------11.20
   |           |--------ipoe-7-------------------------------------5.60
   |           |--------ipoe-6-------------------------------------5.60
   |--------vport-5-----------------------------------11.20
   |           |--------ipoe-8-------------------------------------11.20
   |--------vport-6-----------------------------------0.00
   |--------vport-7-----------------------------------0.00
   |--------vport-8-----------------------------------3.73
   |           |--------ipoe-14------------------------------------1.86
   |           |--------ipoe-13------------------------------------1.87
   |--------vport-9-----------------------------------0.00
   |--------vport-10----------------------------------0.74
   |           |--------ipoe-16------------------------------------0.37
   |           |--------ipoe-17------------------------------------0.37
   |--------vport-11----------------------------------0.00
   |--------vport-12----------------------------------0.74
   |           |--------ipoe-20------------------------------------0.37
   |           |--------ipoe-21------------------------------------0.37
   |--------vport-13----------------------------------0.00
   |--------vport-14----------------------------------0.74
   |           |--------ipoe-25------------------------------------0.37
   |           |--------ipoe-24------------------------------------0.37
   |--------vport-15----------------------------------0.00
   |--------vport-16----------------------------------0.75
               |--------ipoe-32------------------------------------0.15
               |--------ipoe-30------------------------------------0.15
               |--------ipoe-31------------------------------------0.15
               |--------ipoe-29------------------------------------0.15
               |--------ipoe-28------------------------------------0.15


-------------------------------------------------------------------------
                                Run 3
-------------------------------------------------------------------------
              Object               |           BW % of Total
-------------------------------------------------------------------------


pxc-1.b------------------------------------**59.70**
   |--------vport-1-----------------------------------33.58
   |           |--------ipoe-1-------------------------------------33.58
   |--------vport-2-----------------------------------0.00
   |--------vport-3-----------------------------------11.19
   |           |--------ipoe-4-------------------------------------5.60
   |           |--------ipoe-5-------------------------------------5.59
   |--------vport-4-----------------------------------0.00
   |--------vport-5-----------------------------------0.00
   |--------vport-6-----------------------------------3.73
   |           |--------ipoe-10------------------------------------1.86
   |           |--------ipoe-9-------------------------------------1.86
   |--------vport-7-----------------------------------3.73
   |           |--------ipoe-12------------------------------------1.87
   |           |--------ipoe-11------------------------------------1.86
   |--------vport-8-----------------------------------0.00
   |--------vport-9-----------------------------------5.22
   |           |--------ipoe-15------------------------------------5.22
   |--------vport-10----------------------------------0.00
   |--------vport-11----------------------------------0.75
   |           |--------ipoe-19------------------------------------0.38
   |           |--------ipoe-18------------------------------------0.38
   |--------vport-12----------------------------------0.00
   |--------vport-13----------------------------------0.75
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
   |           |--------ipoe-23-------------------------------------0.38
   |           |--------ipoe-22-------------------------------------0.38
   |--------vport-14-----------------------------------0.00
   |--------vport-15-----------------------------------0.75
   |           |--------ipoe-27-------------------------------------0.38
   |           |--------ipoe-26-------------------------------------0.38
   |--------vport-16-----------------------------------0.00


pxc-2.b-------------------------------------40.30
   |--------vport-1------------------------------------0.00
   |--------vport-2------------------------------------11.19
   |           |--------ipoe-3--------------------------------------5.59
   |           |--------ipoe-2--------------------------------------5.59
   |--------vport-3------------------------------------0.00
   |--------vport-4------------------------------------11.20
   |           |--------ipoe-7--------------------------------------5.60
   |           |--------ipoe-6--------------------------------------5.60
   |--------vport-5------------------------------------11.20
   |           |--------ipoe-8--------------------------------------11.20
   |--------vport-6------------------------------------0.00
   |--------vport-7------------------------------------0.00
   |--------vport-8------------------------------------3.73
   |           |--------ipoe-14-------------------------------------1.87
   |           |--------ipoe-13-------------------------------------1.87
   |--------vport-9------------------------------------0.00
   |--------vport-10-----------------------------------0.74
   |           |--------ipoe-16-------------------------------------0.37
   |           |--------ipoe-17-------------------------------------0.37
   |--------vport-11-----------------------------------0.00
   |--------vport-12-----------------------------------0.75
   |           |--------ipoe-20-------------------------------------0.38
   |           |--------ipoe-21-------------------------------------0.38
   |--------vport-13-----------------------------------0.00
   |--------vport-14-----------------------------------0.75
   |           |--------ipoe-25-------------------------------------0.38
   |           |--------ipoe-24-------------------------------------0.38
   |--------vport-15-----------------------------------0.00
   |--------vport-16-----------------------------------0.74
               |--------ipoe-32-------------------------------------0.15
               |--------ipoe-30-------------------------------------0.15
               |--------ipoe-31-------------------------------------0.15
               |--------ipoe-29-------------------------------------0.15
               |--------ipoe-28-------------------------------------0.15
```

Per-subscriber load balancing

In this example, Vports are not considered in hashing, due to the removal of the **subscriber-hash-mode vport** command:

```
    per-link-hash {
        weighted {
            delete subscriber-hash-mode
```

The configuration file for BNG-1 can be found here .

The "lag-3" configuration is now as follows:

```
    lag "lag-3" {
        admin-state enable
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        description "fpe pw-port pxc lag - termination side"
        mode hybrid
        max-ports 64
        access {
            per-fp-ing-queuing false
            per-fp-egr-queuing false
            per-fp-sap-instance false
            adapt-qos {
                mode port-fair
            }
        }
        port pxc-1.b {
        }
        port pxc-2.b {
        }
    }
```

Because Vports are not used, a simpler **monitor** command is used to collect the data that is going through the port scheduler without visibility into subscribers. The output of this command is:

```
A:admin@bng-1# monitor-scheduler -l lag-3 -i 10 -r 3


-------------------------------------------------------------------------
                    Run 1 (interval = 10 seconds)
-------------------------------------------------------------------------
         Lag Member Port           |           BW % of Total
-------------------------------------------------------------------------
             pxc-1.b               |               36.27
             pxc-2.b               |               63.73


-------------------------------------------------------------------------
                    Run 2 (interval = 10 seconds)
-------------------------------------------------------------------------
         Lag Member Port           |           BW % of Total
-------------------------------------------------------------------------
             pxc-1.b               |               36.27
             pxc-2.b               |               63.73


-------------------------------------------------------------------------
                    Run 3 (interval = 10 seconds)
-------------------------------------------------------------------------
         Lag Member Port           |           BW % of Total
-------------------------------------------------------------------------
             pxc-1.b               |               36.27
             pxc-2.b               |               63.73
```

The output shows that the per-subscriber hashing also produces unequal load balancing due to the significant difference in rates between the subscribers.

This unequal load balancing can be corrected by configuring classes and weights per subscriber (defined in the sub-profile), but this configuration is beyond the scope of this chapter.

© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

LAG "lag-1"

This section describes the load balancing on the LAG that connects the BNG to the aggregation node "AGG-1". LAG "lag-1" contains regular ports (no PXCs) and is referenced under the L3 interfaces that connect the BNG to the aggregation node.

```
[pr:/configure router "Base"]
A:admin@bng-1#
    autonomous-system 64500
    router-id 192.0.2.20
    interface "int-1-bng-1-p-1" {
        port lag-1:1
        ipv6 {
            address 2001:db8::501 {
                prefix-length 120
            }
        }
    }
```

Although not explicitly referenced in the Epipe (EVPN VPWS), this LAG is used to reach the network destination in EVPN, which is the BGP neighbor in the following configuration:

```
[pr:/configure service epipe "access"]
A:admin@bng-1#
    admin-state enable
    service-id 1
    customer "1"
    bgp 1 {
    }
    bgp-evpn {
        evi 10
        local-attachment-circuit "bng" {
            eth-tag 2
        }
        remote-attachment-circuit "access" {
            eth-tag 1
        }
        mpls 1 {
            admin-state enable
            send-tunnel-encap {
                mpls-over-udp false
            }
            auto-bind-tunnel {
                resolution any
            }
            route-next-hop {
                system-ipv6
            }
        }
    }

[pr:/configure router "Base" bgp]
A:admin@bng-1#
    admin-state enable
    vpn-apply-export true
    vpn-apply-import true
    rapid-withdrawal true
    rapid-update {
        evpn true
    }
    group "evpn" {
        peer-as 64500
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        local-address 2001:db8::14
        family {
            evpn true
        }
    }
    neighbor "2001:db8::5" {
        group "evpn"
    }
```

LAG "lag-1" is configured as follows with the member ports shown in the figure Figure 361: Example topology:

```
[pr:/configure lag "lag-1"]
A:admin@bng-1#
    admin-state enable
    description "lag to p-1"
    mode hybrid
    max-ports 64
    port 1/1/c1/1 {
    }
    port 1/1/c1/2 {
    }
```

This section describes two ways of traffic load balancing toward the network destination (in this case, the system IPv6 address, 2001:db8::5, on the aggregation node) in the Epipe (EVPN VPWS).

By default, traffic is load-balanced per source and destination IP address pair. Traffic is sent from the same source IP address on the traffic generator toward each individual subscriber, each subscriber with its own IP address. Because every <source, destination> IP address pair in the setup represents a unique subscriber, one could conclude that this would be equivalent to per-subscriber load balancing. However, the hashing keys here are the source and destination IP address, while in per-subscriber hashing, the hashing key is the subscriber ID (and internal SR number). As a result, the outcome of traffic distribution is different in "lag-1" and "lag-3" in the preceding **monitor-scheduler** command with load-balancing per subscriber.

Optimally, L4 ports can be added to the source and destination IP address as the hashing input.

The output of the **monitor lag** command when load balancing per source and destination IP address is enabled is shown below.

Note: The **monitor lag** command is a native SR command, not to be confused with the **monitor-lag** command using pysros.

The distribution of this traffic is 30/70 (811 output packet on port 1/1/c1/1 versus 1864 output packets on port 1/1/c1/2; all packets are the same size of 200 bytes).

This is a very different distribution from the output in the preceding **monitor-scheduler** command with load balancing per subscriber where the traffic distribution over the two ports was approximately 36/63.

```
A:admin@bng-1# monitor lag 1 interval 3 repeat 999 rate packets

===============================================================================
Monitor statistics for LAG ID 1
===============================================================================
Port-id         Input packets                   Output packets
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
At time t = 3 sec (Mode: Rate)
-------------------------------------------------------------------------------
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
1/1/c1/1        0                                   811
1/1/c1/2        0                                   1864
-------------------------------------------------------------------------------
Totals          0                                   2675

-------------------------------------------------------------------------------
At time t = 6 sec (Mode: Rate)
-------------------------------------------------------------------------------
1/1/c1/1        0                                   812
1/1/c1/2        0                                   1868
-------------------------------------------------------------------------------
Totals          0                                   2680

-------------------------------------------------------------------------------
At time t = 9 sec (Mode: Rate)
-------------------------------------------------------------------------------
1/1/c1/1        2                                   813
1/1/c1/2        1                                   1869
-------------------------------------------------------------------------------
Totals          3                                   2682
---snip---
```

A more even distribution is achieved by enabling additional fields for hashing purposes. In the following example, in addition to the source and destination IP addresses, the hashing algorithm also considers L4 ports, which in this example, are randomly changing in each packet.

```
[pr:/configure system load-balancing]
A:admin@bng-1#
    l4-load-balancing true
```

The result of such hashing is more even load balancing (approximately 50/50), as follows:

```
A:admin@bng-1# monitor lag 1 interval 3 repeat 5 rate packets

===============================================================================
Monitor statistics for LAG ID 1
===============================================================================
Port-id         Input packets                   Output packets
-------------------------------------------------------------------------------
At time t = 3 sec (Mode: Rate)
-------------------------------------------------------------------------------
1/1/c1/1        1                                   1354
1/1/c1/2        0                                   1319
-------------------------------------------------------------------------------
Totals          1                                   2673

-------------------------------------------------------------------------------
At time t = 6 sec (Mode: Rate)
-------------------------------------------------------------------------------
1/1/c1/1        1                                   1341
1/1/c1/2        1                                   1339
-------------------------------------------------------------------------------
Totals          2                                   2680

-------------------------------------------------------------------------------
At time t = 9 sec (Mode: Rate)
-------------------------------------------------------------------------------
1/1/c1/1        0                                   1326
1/1/c1/2        0                                   1354
-------------------------------------------------------------------------------
Totals          0                                   2680
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
-------------------------------------------------------------------------------
At time t = 12 sec (Mode: Rate)
-------------------------------------------------------------------------------
1/1/c1/1        0                                1326
1/1/c1/2        0                                1353
-------------------------------------------------------------------------------
Totals          0                                2679

-------------------------------------------------------------------------------
At time t = 15 sec (Mode: Rate)
-------------------------------------------------------------------------------
1/1/c1/1        2                                1360
1/1/c1/2        1                                1321
-------------------------------------------------------------------------------
Totals          3                                2681
```

Conclusion

This chapter provides fundamentals for understanding and configuring the class- and weight-based Vport traffic load balancing on a LAG in ESM. This is used in deployments where not only the bandwidth service offerings for subscribers vary greatly, but also where there is a significant variation in configured bandwidth between the Vports.

The configuration setup relies on FPE-based PW-ports, with two stage LAGs. The first stage LAG contains PXC based member ports as the subscriber termination points. This stage is the focus of this chapter. The second stage LAG is applied in the EVPN VPWS service context with the regular faceplate ports connecting the BNG to the access nodes. The distinct role of each LAG is explained, and the outcome of the load balancing techniques is each stage is explored.

Traffic utilization on the LAG member links is analyzed and compared between different hashing algorithms (per Vport with classes and weights, per Vport without classes and weighs, and per subscriber).

Appendix

The following files are used in this chapter:

- configuration of AGG-1
- configuration of BNG-1 with per-Vport load balancing with classes and weights
- configuration of BNG-1 with per-Vport load balancing without classes and weights
- configuration of BNG-1 with per-subscriber load balancing
- RADIUS users

agg-1

```
# TiMOS-B-22.10.R1 both/x86_64 Nokia 7750 SR Copyright (c) 2000-2022 Nokia.
# All rights reserved. All use subject to applicable license agreements.
# Built on Sun Oct 30 14:49:55 PDT 2022 by builder in /builds/c/2210B/R1/panos/main/sros
# Configuration format version 22.10 revision 0

# Generated 2022-11-18T08:49:28.3Z by admin from 135.231.208.32
# Commit ID 1
#   Committed 2022-11-17T11:34:34.8Z by system (MD-CLI) from Console
#   Log       "System booted version B-22.10.R1."
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
configure {
    card 1 {
        card-type iom-1
        mda 1 {
            mda-type me6-100gb-qsfp28
        }
        mda 2 {
            mda-type me6-100gb-qsfp28
            xconnect {
                mac 1 {
                    loopback 1 {
                    }
                }
            }
        }
        fp 1 {
        }
    }
    log {
        filter "1001" {
            named-entry "10" {
                description "Collect only events of major severity or higher"
                action forward
                match {
                    severity {
                        gte major
                    }
                }
            }
        }
        log-id "100" {
            description "Default Serious Errors Log"
            filter "1001"
            source {
                main true
            }
            destination {
                memory {
                    max-entries 500
                }
            }
        }
        log-id "99" {
            description "Default System Log"
            source {
                main true
            }
            destination {
                memory {
                    max-entries 500
                }
            }
        }
    }
    port 1/1/c1 {
        admin-state enable
        connector {
            breakout c10-10g
        }
    }
    port 1/1/c1/1 {
        admin-state enable
        ethernet {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            mode hybrid
        }
    }
    port 1/1/c1/2 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/3 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/4 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/5 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/6 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/7 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/8 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/9 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/10 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2 {
        admin-state enable
        connector {
            breakout c10-10g
        }
    }
    port 1/1/c2/1 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/2 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/3 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/4 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/5 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/6 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/7 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/8 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/9 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2/10 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3 {
        admin-state enable
        connector {
            breakout c10-10g
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    }
    port 1/1/c3/1 {
        admin-state enable
        ethernet {
            mode hybrid
            encap-type qinq
        }
    }
    port 1/1/c3/2 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/3 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/4 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/5 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/6 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/7 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/8 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/9 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c3/10 {
        admin-state enable
        ethernet {
            mode hybrid
            encap-type dot1q
        }
    }
    port 1/1/c4 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        admin-state enable
        connector {
            breakout c4-10g
        }
    }
    port 1/1/c4/1 {
        admin-state enable
        ethernet {
            mode hybrid
            dot1x {
                tunneling true
            }
        }
    }
    port 1/1/c5 {
    }
    port 1/1/c6 {
    }
    port 1/2/c1 {
    }
    port 1/2/c2 {
    }
    port 1/2/c3 {
    }
    port 1/2/c4 {
    }
    port 1/2/c5 {
    }
    port 1/2/c6 {
    }
    port 1/2/m1/1 {
    }
    router "Base" {
        autonomous-system 64500
        router-id 192.0.2.5
        interface "int-1-pe-1-p-1" {
            port 1/1/c1/1:1
            ipv6 {
                address 2001:db8::101 {
                    prefix-length 120
                }
            }
        }
        interface "system" {
            ipv4 {
                primary {
                    address 192.0.2.5
                    prefix-length 32
                }
            }
            ipv6 {
                address 2001:db8::5 {
                    prefix-length 128
                }
            }
        }
        interface "to-ixia" {
            port 1/1/c3/1:1.1
            ipv4 {
                primary {
                    address 172.16.100.1
                    prefix-length 24
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            ipv6 {
                address 2001:db8::6401 {
                    prefix-length 120
                }
            }
        }
        interface "to-radius" {
            port 1/1/c3/10:114
            ipv4 {
                primary {
                    address 192.168.114.5
                    prefix-length 24
                }
            }
        }
        bgp {
            admin-state enable
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            rapid-update {
                evpn true
            }
            advertise-ipv6-next-hops {
                evpn true
            }
            group "evpn" {
                peer-as 64500
                local-address 2001:db8::5
                family {
                    evpn true
                }
            }
            neighbor "2001:db8::14" {
                group "evpn"
            }
        }
        isis 0 {
            admin-state enable
            advertise-passive-only false
            advertise-router-capability area
            ipv6-routing native
            level-capability 2
            traffic-engineering true
            area-address [49.0001]
            interface "int-1-pe-1-p-1" {
            }
            interface "system" {
            }
            interface "to-ixia" {
            }
        }
        ldp {
            interface-parameters {
                interface "int-1-pe-1-p-1" {
                    ipv6 {
                        admin-state enable
                    }
                }
            }
        }
    }
    service {
        system {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            extended-default-qinq-sap-lookup true
        }
        epipe "access" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
            }
            sap 1/1/c3/1:*.* {
            }
            bgp-evpn {
                evi 10
                local-attachment-circuit "access" {
                    eth-tag 1
                }
                remote-attachment-circuit "bng" {
                    eth-tag 2
                }
                mpls 1 {
                    admin-state enable
                    send-tunnel-encap {
                        mpls-over-udp true
                    }
                    auto-bind-tunnel {
                        resolution any
                    }
                    route-next-hop {
                        system-ipv6
                    }
                }
            }
        }
    }
    system {
        name "AGG-1"
        management-interface {
            configuration-mode model-driven
            cli {
                cli-engine [md-cli classic-cli]
            }
            yang-modules {
                nokia-submodules true
                nokia-combined-modules false
            }
            snmp {
                admin-state disable
            }
        }
        ip {
            allow-qinq-network-interface true
        }
        bluetooth {
            advertising-timeout 30
        }
        login-control {
            idle-timeout none
        }
        security {
            aaa {
                local-profiles {
                    profile "administrative" {
                        default-action permit-all
                        entry 10 {
                            match "configure system security"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        action permit
                    }
                    entry 20 {
                        match "show system security"
                        action permit
                    }
                    entry 30 {
                        match "tools perform security"
                        action permit
                    }
                    entry 40 {
                        match "tools dump security"
                        action permit
                    }
                    entry 50 {
                        match "admin system security"
                        action permit
                    }
                    entry 100 {
                        match "configure li"
                        action deny
                    }
                    entry 110 {
                        match "show li"
                        action deny
                    }
                    entry 111 {
                        match "clear li"
                        action deny
                    }
                    entry 112 {
                        match "tools dump li"
                        action deny
                    }
                    netconf {
                        base-op-authorization {
                            action true
                            cancel-commit true
                            close-session true
                            commit true
                            copy-config true
                            create-subscription true
                            delete-config true
                            discard-changes true
                            edit-config true
                            get true
                            get-config true
                            get-data true
                            get-schema true
                            kill-session true
                            lock true
                            validate true
                        }
                    }
                }
                profile "default" {
                    entry 10 {
                        match "exec"
                        action permit
                    }
                    entry 20 {
                        match "exit"
                        action permit
                    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        entry 30 {
                            match "help"
                            action permit
                        }
                        entry 40 {
                            match "logout"
                            action permit
                        }
                        entry 50 {
                            match "password"
                            action permit
                        }
                        entry 60 {
                            match "show config"
                            action deny
                        }
                        entry 65 {
                            match "show li"
                            action deny
                        }
                        entry 66 {
                            match "clear li"
                            action deny
                        }
                        entry 67 {
                            match "tools dump li"
                            action deny
                        }
                        entry 68 {
                            match "state li"
                            action deny
                        }
                        entry 70 {
                            match "show"
                            action permit
                        }
                        entry 75 {
                            match "state"
                            action permit
                        }
                        entry 80 {
                            match "enable-admin"
                            action permit
                        }
                        entry 90 {
                            match "enable"
                            action permit
                        }
                        entry 100 {
                            match "configure li"
                            action deny
                        }
                        netconf {
                            base-op-authorization {
                                action true
                                cancel-commit true
                                close-session true
                                commit true
                                copy-config true
                                create-subscription true
                                delete-config true
                                discard-changes true
                                edit-config true
                                get true
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                            get-config true
                            get-data true
                            get-schema true
                            validate true
                        }
                    }
                }
            }
        }
        ssh {
            server-cipher-list-v2 {
                cipher 190 {
                    name aes256-ctr
                }
                cipher 192 {
                    name aes192-ctr
                }
                cipher 194 {
                    name aes128-ctr
                }
                cipher 200 {
                    name aes128-cbc
                }
                cipher 205 {
                    name 3des-cbc
                }
                cipher 225 {
                    name aes192-cbc
                }
                cipher 230 {
                    name aes256-cbc
                }
            }
            client-cipher-list-v2 {
                cipher 190 {
                    name aes256-ctr
                }
                cipher 192 {
                    name aes192-ctr
                }
                cipher 194 {
                    name aes128-ctr
                }
                cipher 200 {
                    name aes128-cbc
                }
                cipher 205 {
                    name 3des-cbc
                }
                cipher 225 {
                    name aes192-cbc
                }
                cipher 230 {
                    name aes256-cbc
                }
            }
            server-mac-list-v2 {
                mac 200 {
                    name hmac-sha2-512
                }
                mac 210 {
                    name hmac-sha2-256
                }
                mac 215 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        name hmac-sha1
                    }
                    mac 220 {
                        name hmac-sha1-96
                    }
                    mac 225 {
                        name hmac-md5
                    }
                    mac 240 {
                        name hmac-md5-96
                    }
                }
                client-mac-list-v2 {
                    mac 200 {
                        name hmac-sha2-512
                    }
                    mac 210 {
                        name hmac-sha2-256
                    }
                    mac 215 {
                        name hmac-sha1
                    }
                    mac 220 {
                        name hmac-sha1-96
                    }
                    mac 225 {
                        name hmac-md5
                    }
                    mac 240 {
                        name hmac-md5-96
                    }
                }
            }
            user-params {
                local-user {
                    user "admin" {
                        password "$2y$10$TQrZlpBDra86.qoexZUzQeBXDY1FcdDhGWdD9lLxMuFyPVSm0OGy6"
                        access {
                            console true
                        }
                        console {
                            member ["administrative"]
                        }
                    }
                }
            }
        }
    }
}


# Finished 2022-11-18T08:49:28.3Z
```

bng-per-vport-balancing-with-classes-weights

```
# TiMOS-C-22.10.R1 cpm/x86_64 Nokia 7750 SR Copyright (c) 2000-2022 Nokia.
# All rights reserved. All use subject to applicable license agreements.
# Built on Sun Oct 30 14:49:55 PDT 2022 by builder in /builds/c/2210B/R1/panos/main/sros
# Configuration format version 22.10 revision 0

# Generated 2022-12-14T13:31:35.4-06:00 by admin from 135.231.208.32
# Commit ID 2
#   Committed 2022-12-14T12:45:28.3-06:00 by admin (MD-CLI) from 135.231.208.32
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
# Commit ID 1
#   Committed 2022-12-14T12:38:03.7-06:00 by system (MD-CLI) from Console
#   Log       "System booted version C-22.10.R1."

configure {
    aaa {
        radius {
            server-policy "radius-server-1" {
                servers {
                    router-instance "Base"
                    server 1 {
                        server-name "free-radius-1"
                    }
                }
            }
        }
    }
    card 1 {
        card-type iom5-e
        mda 1 {
            mda-type me6-100gb-qsfp28
        }
        mda 2 {
            mda-type me6-100gb-qsfp28
        }
    }
    fwd-path-ext {
        sdp-id-range {
            start 17500
            end 17600
        }
        fpe 1 {
            description "pw-port on a lag"
            path {
                xc-lag-a "lag-2"
                xc-lag-b "lag-3"
            }
            application {
                pw-port-extension {
                }
            }
        }
    }
    lag "lag-1" {
        admin-state enable
        description "lag to p-1"
        mode hybrid
        max-ports 64
        port 1/1/c1/1 {
        }
        port 1/1/c1/2 {
        }
    }
    lag "lag-2" {
        admin-state enable
        description "fpe pw-port pxc lag - transit side"
        mode hybrid
        max-ports 64
        port pxc-1.a {
        }
        port pxc-2.a {
        }
    }
    lag "lag-3" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        admin-state enable
        description "fpe pw-port pxc lag - termination side"
        mode hybrid
        max-ports 64
        access {
            per-fp-ing-queuing true
            per-fp-egr-queuing true
            per-fp-sap-instance true
            adapt-qos {
                mode link
            }
        }
        per-link-hash {
            weighted {
                subscriber-hash-mode vport
            }
        }
        port pxc-1.b {
        }
        port pxc-2.b {
        }
    }
    log {
        filter "1001" {
            named-entry "10" {
                description "Collect only events of major severity or higher"
                action forward
                match {
                    severity {
                        gte major
                    }
                }
            }
        }
        log-id "100" {
            description "Default Serious Errors Log"
            filter "1001"
            source {
                main true
            }
            destination {
                memory {
                    max-entries 500
                }
            }
        }
        log-id "99" {
            description "Default System Log"
            source {
                main true
            }
            destination {
                memory {
                    max-entries 500
                }
            }
        }
    }
    port 1/1/c1 {
        admin-state enable
        connector {
            breakout c10-10g
        }
    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    port 1/1/c1/1 {
        admin-state enable
        description "lag-1 to p-1 access"
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/2 {
        admin-state enable
        description "lag-1 to p-1 access"
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/3 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/4 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/5 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/6 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/7 {
        admin-state enable
        description "to pe-2 network"
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/8 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/9 {
        admin-state enable
        ethernet {
            mode hybrid
            dot1x {
                tunneling true
            }
        }
    }
    port 1/1/c1/10 {
        admin-state enable
        description "RADIUS and mirroring"
        ethernet {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            mode hybrid
        }
    }
    port 1/1/c2 {
        admin-state enable
        connector {
            breakout c1-100g
        }
    }
    port 1/1/c2/1 {
        admin-state enable
        ethernet {
            mode hybrid
            encap-type dot1q
            dot1x {
                tunneling true
            }
        }
    }
    port 1/1/c4 {
        admin-state enable
        connector {
            breakout c4-10g
        }
    }
    port 1/1/c4/1 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c4/2 {
        admin-state enable
        description "PXC 1; pw-port lag 2,3"
        ethernet {
            mode hybrid
            dot1x {
                tunneling true
            }
        }
    }
    port 1/1/c4/3 {
        admin-state enable
        description "PXC 2; pw-port lag 2,3"
        ethernet {
            mode hybrid
            dot1x {
                tunneling true
            }
        }
    }
    port 1/1/c4/4 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port pxc-1.a {
        admin-state enable
    }
    port pxc-1.b {
        admin-state enable
        description "PXC lag-3, pw-port termination"
        ethernet {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            access {
                egress {
                    virtual-port "vport-1" {
                        aggregate-rate {
                            rate 9000000
                        }
                        host-match {
                            int-dest-id "vport-1" { }
                        }
                        lag-per-link-hash {
                            class 1
                            weight 15
                        }
                    }
                    virtual-port "vport-10" {
                        aggregate-rate {
                            rate 200000
                        }
                        host-match {
                            int-dest-id "vport-10" { }
                        }
                        lag-per-link-hash {
                            class 3
                            weight 1
                        }
                    }
                    virtual-port "vport-11" {
                        aggregate-rate {
                            rate 200000
                        }
                        host-match {
                            int-dest-id "vport-11" { }
                        }
                        lag-per-link-hash {
                            class 3
                            weight 1
                        }
                    }
                    virtual-port "vport-12" {
                        aggregate-rate {
                            rate 200000
                        }
                        host-match {
                            int-dest-id "vport-12" { }
                        }
                        lag-per-link-hash {
                            class 3
                            weight 1
                        }
                    }
                    virtual-port "vport-13" {
                        aggregate-rate {
                            rate 200000
                        }
                        host-match {
                            int-dest-id "vport-13" { }
                        }
                        lag-per-link-hash {
                            class 3
                            weight 1
                        }
                    }
                    virtual-port "vport-14" {
                        aggregate-rate {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        rate 200000
                    }
                    host-match {
                        int-dest-id "vport-14" { }
                    }
                    lag-per-link-hash {
                        class 3
                        weight 1
                    }
                }
                virtual-port "vport-15" {
                    aggregate-rate {
                        rate 200000
                    }
                    host-match {
                        int-dest-id "vport-15" { }
                    }
                    lag-per-link-hash {
                        class 3
                        weight 1
                    }
                }
                virtual-port "vport-16" {
                    aggregate-rate {
                        rate 200000
                    }
                    host-match {
                        int-dest-id "vport-16" { }
                    }
                    lag-per-link-hash {
                        class 3
                        weight 1
                    }
                }
                virtual-port "vport-2" {
                    aggregate-rate {
                        rate 3000000
                    }
                    host-match {
                        int-dest-id "vport-2" { }
                    }
                    lag-per-link-hash {
                        class 1
                        weight 5
                    }
                }
                virtual-port "vport-3" {
                    aggregate-rate {
                        rate 3000000
                    }
                    host-match {
                        int-dest-id "vport-3" { }
                    }
                    lag-per-link-hash {
                        class 1
                        weight 5
                    }
                }
                virtual-port "vport-4" {
                    aggregate-rate {
                        rate 3000000
                    }
                    host-match {
                        int-dest-id "vport-4" { }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    }
                    lag-per-link-hash {
                        class 1
                        weight 5
                    }
                }
                virtual-port "vport-5" {
                    aggregate-rate {
                        rate 3000000
                    }
                    host-match {
                        int-dest-id "vport-5" { }
                    }
                    lag-per-link-hash {
                        class 2
                        weight 15
                    }
                }
                virtual-port "vport-6" {
                    aggregate-rate {
                        rate 1000000
                    }
                    host-match {
                        int-dest-id "vport-6" { }
                    }
                    lag-per-link-hash {
                        class 2
                        weight 5
                    }
                }
                virtual-port "vport-7" {
                    aggregate-rate {
                        rate 1000000
                    }
                    host-match {
                        int-dest-id "vport-7" { }
                    }
                    lag-per-link-hash {
                        class 2
                        weight 5
                    }
                }
                virtual-port "vport-8" {
                    aggregate-rate {
                        rate 1000000
                    }
                    host-match {
                        int-dest-id "vport-8" { }
                    }
                    lag-per-link-hash {
                        class 2
                        weight 5
                    }
                }
                virtual-port "vport-9" {
                    aggregate-rate {
                        rate 1400000
                    }
                    host-match {
                        int-dest-id "vport-9" { }
                    }
                    lag-per-link-hash {
                        class 3
                        weight 7
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    }
                }
            }
        }
        egress {
            port-scheduler-policy {
                policy-name "lag-3-pxc"
            }
        }
    }
}
port pxc-2.a {
    admin-state enable
}
port pxc-2.b {
    admin-state enable
    description "PXC lag-3, pw-port termination"
    ethernet {
        access {
            egress {
                virtual-port "vport-1" {
                    aggregate-rate {
                        rate 9000000
                    }
                    host-match {
                        int-dest-id "vport-1" { }
                    }
                    lag-per-link-hash {
                        class 1
                        weight 15
                    }
                }
                virtual-port "vport-10" {
                    aggregate-rate {
                        rate 200000
                    }
                    host-match {
                        int-dest-id "vport-10" { }
                    }
                    lag-per-link-hash {
                        class 3
                        weight 1
                    }
                }
                virtual-port "vport-11" {
                    aggregate-rate {
                        rate 200000
                    }
                    host-match {
                        int-dest-id "vport-11" { }
                    }
                    lag-per-link-hash {
                        class 3
                        weight 1
                    }
                }
                virtual-port "vport-12" {
                    aggregate-rate {
                        rate 200000
                    }
                    host-match {
                        int-dest-id "vport-12" { }
                    }
                    lag-per-link-hash {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        class 3
                        weight 1
                    }
                }
                virtual-port "vport-13" {
                    aggregate-rate {
                        rate 200000
                    }
                    host-match {
                        int-dest-id "vport-13" { }
                    }
                    lag-per-link-hash {
                        class 3
                        weight 1
                    }
                }
                virtual-port "vport-14" {
                    aggregate-rate {
                        rate 200000
                    }
                    host-match {
                        int-dest-id "vport-14" { }
                    }
                    lag-per-link-hash {
                        class 3
                        weight 1
                    }
                }
                virtual-port "vport-15" {
                    aggregate-rate {
                        rate 200000
                    }
                    host-match {
                        int-dest-id "vport-15" { }
                    }
                    lag-per-link-hash {
                        class 3
                        weight 1
                    }
                }
                virtual-port "vport-16" {
                    aggregate-rate {
                        rate 200000
                    }
                    host-match {
                        int-dest-id "vport-16" { }
                    }
                    lag-per-link-hash {
                        class 3
                        weight 1
                    }
                }
                virtual-port "vport-2" {
                    aggregate-rate {
                        rate 3000000
                    }
                    host-match {
                        int-dest-id "vport-2" { }
                    }
                    lag-per-link-hash {
                        class 1
                        weight 5
                    }
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    virtual-port "vport-3" {
                        aggregate-rate {
                            rate 3000000
                        }
                        host-match {
                            int-dest-id "vport-3" { }
                        }
                        lag-per-link-hash {
                            class 1
                            weight 5
                        }
                    }
                    virtual-port "vport-4" {
                        aggregate-rate {
                            rate 3000000
                        }
                        host-match {
                            int-dest-id "vport-4" { }
                        }
                        lag-per-link-hash {
                            class 1
                            weight 5
                        }
                    }
                    virtual-port "vport-5" {
                        aggregate-rate {
                            rate 3000000
                        }
                        host-match {
                            int-dest-id "vport-5" { }
                        }
                        lag-per-link-hash {
                            class 2
                            weight 15
                        }
                    }
                    virtual-port "vport-6" {
                        aggregate-rate {
                            rate 1000000
                        }
                        host-match {
                            int-dest-id "vport-6" { }
                        }
                        lag-per-link-hash {
                            class 2
                            weight 5
                        }
                    }
                    virtual-port "vport-7" {
                        aggregate-rate {
                            rate 1000000
                        }
                        host-match {
                            int-dest-id "vport-7" { }
                        }
                        lag-per-link-hash {
                            class 2
                            weight 5
                        }
                    }
                    virtual-port "vport-8" {
                        aggregate-rate {
                            rate 1000000
                        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        host-match {
                            int-dest-id "vport-8" { }
                        }
                        lag-per-link-hash {
                            class 2
                            weight 5
                        }
                    }
                    virtual-port "vport-9" {
                        aggregate-rate {
                            rate 1400000
                        }
                        host-match {
                            int-dest-id "vport-9" { }
                        }
                        lag-per-link-hash {
                            class 3
                            weight 7
                        }
                    }
                }
            }
            egress {
                port-scheduler-policy {
                    policy-name "lag-3-pxc"
                }
            }
        }
    }
    port-xc {
        pxc 1 {
            admin-state enable
            description "pw-port lag 2,3"
            port-id 1/1/c4/2
        }
        pxc 2 {
            admin-state enable
            description "pw-port lag 2,3"
            port-id 1/1/c4/3
        }
    }
    pw-port 1 {
        encap-type qinq
        epipe "access" {
            admin-state enable
            fpe-id 1
        }
    }
    python {
        python-script "monitor" {
            admin-state enable
            urls ["ftp://135.231.216.68/pub/configs/alu/SIMS/ACGs/hashing-per-vport-on-pxc/
monitor_lag_sr.py"]
            version python3
        }
        python-script "monitor-sched" {
            admin-state enable
            urls ["ftp://135.231.216.68/pub/configs/alu/SIMS/ACGs/hashing-per-vport-on-pxc/
monitor_lag_port_scheduler_sr.py"]
            version python3
        }
        python-script "monitor-sched-clear" {
            admin-state enable
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            urls ["ftp://135.231.216.68/pub/configs/alu/SIMS/ACGs/hashing-per-vport-on-pxc/
monitor_lag_port_scheduler_clear_sr.py"]
            version python3
        }
    }
    qos {
        sap-egress "sap-egess-1" {
            policy-id 2
            queue 1 {
                port-parent {
                }
            }
        }
        port-scheduler-policy "lag-3-pxc" {
        }
    }
    router "Base" {
        autonomous-system 64500
        router-id 192.0.2.20
        interface "int-1-bng-1-p-1" {
            port lag-1:1
            ipv6 {
                address 2001:db8::501 {
                    prefix-length 120
                }
            }
        }
        interface "system" {
            ipv4 {
                primary {
                    address 192.0.2.20
                    prefix-length 32
                }
            }
            ipv6 {
                address 2001:db8::14 {
                    prefix-length 128
                }
            }
        }
        interface "to-radius" {
            port 1/1/c1/10:114
            ipv4 {
                primary {
                    address 192.168.114.20
                    prefix-length 24
                }
            }
        }
        bgp {
            admin-state enable
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            rapid-update {
                evpn true
            }
            group "evpn" {
                peer-as 64500
                local-address 2001:db8::14
                family {
                    evpn true
                }
            }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            neighbor "2001:db8::5" {
                group "evpn"
            }
        }
        isis 0 {
            admin-state enable
            advertise-passive-only false
            advertise-router-capability area
            ipv6-routing native
            level-capability 2
            traffic-engineering true
            area-address [49.0001]
            interface "int-1-bng-1-p-1" {
            }
            interface "system" {
            }
        }
        ldp {
            interface-parameters {
                interface "int-1-bng-1-p-1" {
                    ipv6 {
                        admin-state enable
                    }
                }
            }
        }
        radius {
            server "free-radius-1" {
                address 192.168.114.2
                secret "HDqTwZYSvEu934VNhUQy/pubZxTKpSDzvHg=" hash2
                accept-coa true
            }
        }
    }
    service {
        epipe "access" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
            }
            bgp-evpn {
                evi 10
                local-attachment-circuit "bng" {
                    eth-tag 2
                }
                remote-attachment-circuit "access" {
                    eth-tag 1
                }
                mpls 1 {
                    admin-state enable
                    send-tunnel-encap {
                        mpls-over-udp false
                    }
                    auto-bind-tunnel {
                        resolution any
                    }
                    route-next-hop {
                        system-ipv6
                    }
                }
            }
        }
        vpls "capture-sap" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            admin-state enable
            service-id 5
            customer "1"
            load-balancing {
                per-service-hashing true
            }
            capture-sap pw-1:*.* {
                radius-auth-policy "radius-1"
                trigger-packet {
                    dhcp true
                }
                ipoe-session {
                    admin-state enable
                    ipoe-session-policy "ipoe-session-policy-1"
                }
            }
        }
        vprn "vport-hashing" {
            admin-state enable
            service-id 10
            customer "1"
            interface "loopback-1" {
                admin-state enable
                loopback true
                ipv4 {
                    local-dhcp-server "dhcpv4"
                    primary {
                        address 192.168.0.1
                        prefix-length 32
                    }
                    neighbor-discovery {
                        local-proxy-arp false
                        remote-proxy-arp false
                    }
                    dhcp {
                        admin-state enable
                    }
                }
            }
            interface "to-ixia" {
                ipv4 {
                    primary {
                        address 172.16.102.1
                        prefix-length 24
                    }
                }
                sap 1/1/c1/7:2 {
                }
                ipv6 {
                    address 2001:db8::6601 {
                        prefix-length 120
                    }
                }
            }
            ipv6 {
                router-advertisement {
                    interface "to-ixia" {
                        admin-state enable
                        max-advertisement-interval 15
                        min-advertisement-interval 10
                    }
                }
            }
            dhcp-server {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                dhcpv4 "dhcpv4" {
                    admin-state enable
                    pool-selection {
                        use-gi-address {
                            scope pool
                        }
                        use-pool-from-client {
                        }
                    }
                    pool "dhcpv4-1" {
                        max-lease-time 1200
                        subnet 10.10.0.0/24 {
                            address-range 10.10.0.10 end 10.10.0.100 {
                                failover-control-type access-driven
                            }
                        }
                    }
                }
            }
            subscriber-interface "sub-int-1" {
                admin-state enable
                ipv4 {
                    address 10.10.0.254 {
                        prefix-length 24
                    }
                }
                group-interface "group-int-1" {
                    admin-state enable
                    radius-auth-policy "radius-1"
                    oper-up-while-empty true
                    ipv4 {
                        neighbor-discovery {
                            remote-proxy-arp true
                            populate false
                        }
                        dhcp {
                            admin-state enable
                            server [192.168.0.1]
                            trusted true
                            gi-address 10.10.0.254
                            match-circuit-id true
                            option-82 {
                                action keep
                                vendor-specific-option {
                                    pool-name true
                                }
                            }
                            lease-populate {
                                max-leases 100
                            }
                            client-applications {
                                dhcp true
                            }
                        }
                    }
                    ipoe-session {
                        admin-state enable
                        ipoe-session-policy "ipoe-session-policy-1"
                        sap-session-limit 100
                        force-auth {
                            cid-change false
                            rid-change false
                        }
                    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
        }
    }
}
sfm 1 {
    sfm-type m-sfm6-7/12
}
subscriber-mgmt {
    ipoe-session-policy "ipoe-session-policy-1" {
    }
    sub-profile "sub-profile-1" {
        egress {
            qos {
                agg-rate {
                    rate 1000
                }
            }
        }
    }
    sla-profile "sla-profile-1" {
        egress {
            qos {
                sap-egress {
                    policy-name "sap-egess-1"
                    port-parent-location vport
                    overrides {
                        queue 1 {
                            stat-mode v4-v6
                        }
                    }
                }
            }
        }
    }
    sub-ident-policy "sub-ident-policy-1" {
        sla-profile-map {
            use-direct-map-as-default true
        }
        sub-profile-map {
            use-direct-map-as-default true
        }
    }
    radius-authentication-policy "radius-1" {
        password "ncd8qyrNUMhYfa2SfrUqHMDZ9IXn3sVSmYBzbw==" hash2
        pppoe-access-method pap-chap
        radius-server-policy "radius-server-1"
        user-name {
            format circuit-id
        }
        include-radius-attribute {
            circuit-id true
            nas-identifier true
        }
    }
    msap-policy "msap-policy-1" {
        sub-sla-mgmt {
            subscriber-limit 1
            sub-ident-policy "sub-ident-policy-1"
            defaults {
                subscriber-id {
                    sap-id
                }
            }
            single-sub-parameters {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                profiled-traffic-only true
            }
        }
        ies-vprn-only-sap-parameters {
            anti-spoof next-hop-ip-and-mac-addr
            ingress {
                qos {
                    queuing-type service
                }
            }
        }
    }
}
system {
    name "bng-1"
    management-interface {
        configuration-mode model-driven
        cli {
            cli-engine [md-cli classic-cli]
            md-cli {
                environment {
                    more false
                    command-alias {
                        alias "monitor-lag" {
                            admin-state enable
                            python-script "monitor"
                            mount-point global { }
                        }
                        alias "monitor-schd-clear" {
                            admin-state enable
                            python-script "monitor-sched-clear"
                            mount-point global { }
                        }
                        alias "monitor-scheduler" {
                            admin-state enable
                            python-script "monitor-sched"
                            mount-point global { }
                        }
                    }
                }
            }
        }
        netconf {
            admin-state enable
        }
        yang-modules {
            nokia-submodules false
            nokia-combined-modules true
        }
        snmp {
            admin-state disable
        }
    }
    login-control {
        idle-timeout none
    }
    security {
        aaa {
            local-profiles {
                profile "administrative" {
                    default-action permit-all
                    entry 10 {
                        match "configure system security"
                        action permit
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    }
                    entry 20 {
                        match "show system security"
                        action permit
                    }
                    entry 30 {
                        match "tools perform security"
                        action permit
                    }
                    entry 40 {
                        match "tools dump security"
                        action permit
                    }
                    entry 50 {
                        match "admin system security"
                        action permit
                    }
                    entry 100 {
                        match "configure li"
                        action deny
                    }
                    entry 110 {
                        match "show li"
                        action deny
                    }
                    entry 111 {
                        match "clear li"
                        action deny
                    }
                    entry 112 {
                        match "tools dump li"
                        action deny
                    }
                    netconf {
                        base-op-authorization {
                            action true
                            cancel-commit true
                            close-session true
                            commit true
                            copy-config true
                            create-subscription true
                            delete-config true
                            discard-changes true
                            edit-config true
                            get true
                            get-config true
                            get-data true
                            get-schema true
                            kill-session true
                            lock true
                            validate true
                        }
                    }
                }
                profile "default" {
                    entry 10 {
                        match "exec"
                        action permit
                    }
                    entry 20 {
                        match "exit"
                        action permit
                    }
                    entry 30 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        match "help"
                        action permit
                    }
                    entry 40 {
                        match "logout"
                        action permit
                    }
                    entry 50 {
                        match "password"
                        action permit
                    }
                    entry 60 {
                        match "show config"
                        action deny
                    }
                    entry 65 {
                        match "show li"
                        action deny
                    }
                    entry 66 {
                        match "clear li"
                        action deny
                    }
                    entry 67 {
                        match "tools dump li"
                        action deny
                    }
                    entry 68 {
                        match "state li"
                        action deny
                    }
                    entry 70 {
                        match "show"
                        action permit
                    }
                    entry 75 {
                        match "state"
                        action permit
                    }
                    entry 80 {
                        match "enable-admin"
                        action permit
                    }
                    entry 90 {
                        match "enable"
                        action permit
                    }
                    entry 100 {
                        match "configure li"
                        action deny
                    }
                }
            }
        }
        ssh {
            server-admin-state enable
            server-cipher-list-v2 {
                cipher 190 {
                    name aes256-ctr
                }
                cipher 192 {
                    name aes192-ctr
                }
                cipher 194 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        name aes128-ctr
                    }
                    cipher 200 {
                        name aes128-cbc
                    }
                    cipher 205 {
                        name 3des-cbc
                    }
                    cipher 225 {
                        name aes192-cbc
                    }
                    cipher 230 {
                        name aes256-cbc
                    }
                }
                client-cipher-list-v2 {
                    cipher 190 {
                        name aes256-ctr
                    }
                    cipher 192 {
                        name aes192-ctr
                    }
                    cipher 194 {
                        name aes128-ctr
                    }
                    cipher 200 {
                        name aes128-cbc
                    }
                    cipher 205 {
                        name 3des-cbc
                    }
                    cipher 225 {
                        name aes192-cbc
                    }
                    cipher 230 {
                        name aes256-cbc
                    }
                }
                server-mac-list-v2 {
                    mac 200 {
                        name hmac-sha2-512
                    }
                    mac 210 {
                        name hmac-sha2-256
                    }
                    mac 215 {
                        name hmac-sha1
                    }
                    mac 220 {
                        name hmac-sha1-96
                    }
                    mac 225 {
                        name hmac-md5
                    }
                    mac 240 {
                        name hmac-md5-96
                    }
                }
                client-mac-list-v2 {
                    mac 200 {
                        name hmac-sha2-512
                    }
                    mac 210 {
                        name hmac-sha2-256
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    }
                    mac 215 {
                        name hmac-sha1
                    }
                    mac 220 {
                        name hmac-sha1-96
                    }
                    mac 225 {
                        name hmac-md5
                    }
                    mac 240 {
                        name hmac-md5-96
                    }
                }
            }
            user-params {
                local-user {
                    user "admin" {
                        password "$2y$10$TQrZlpBDra86.qoexZUzQeBXDY1FcdDhGWdD9lLxMuFyPVSm0OGy6"
                        access {
                            console true
                            netconf true
                        }
                        console {
                            member ["administrative"]
                        }
                    }
                }
            }
        }
        time {
            prefer-local-time true
            zone {
                standard {
                    name cst
                }
            }
            dst-zone "CDT" {
                end {
                    day sunday
                    month november
                    hours-minutes "02:00"
                }
                start {
                    day sunday
                    month march
                    hours-minutes "02:00"
                }
            }
            sntp {
                admin-state enable
                server 135.227.160.253 {
                }
            }
        }
    }
}


# Finished 2022-12-14T13:31:35.4-06:00
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

bng-per-vport-balancing-wo-classes-weights

```
# TiMOS-C-22.10.R1 cpm/x86_64 Nokia 7750 SR Copyright (c) 2000-2022 Nokia.
# All rights reserved. All use subject to applicable license agreements.
# Built on Sun Oct 30 14:49:55 PDT 2022 by builder in /builds/c/2210B/R1/panos/main/sros
# Configuration format version 22.10 revision 0

# Generated 2022-12-14T13:38:31.1-06:00 by admin from 135.231.208.32
# Commit ID 3
#   Committed 2022-12-14T13:33:11.3-06:00 by admin (MD-CLI) from 135.231.208.32

configure {
    aaa {
        radius {
            server-policy "radius-server-1" {
                servers {
                    router-instance "Base"
                    server 1 {
                        server-name "free-radius-1"
                    }
                }
            }
        }
    }
    card 1 {
        card-type iom5-e
        mda 1 {
            mda-type me6-100gb-qsfp28
        }
        mda 2 {
            mda-type me6-100gb-qsfp28
        }
    }
    fwd-path-ext {
        sdp-id-range {
            start 17500
            end 17600
        }
        fpe 1 {
            description "pw-port on a lag"
            path {
                xc-lag-a "lag-2"
                xc-lag-b "lag-3"
            }
            application {
                pw-port-extension {
                }
            }
        }
    }
    lag "lag-1" {
        admin-state enable
        description "lag to p-1"
        mode hybrid
        max-ports 64
        port 1/1/c1/1 {
        }
        port 1/1/c1/2 {
        }
    }
    lag "lag-2" {
        admin-state enable
        description "fpe pw-port pxc lag - transit side"
        mode hybrid
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        max-ports 64
        port pxc-1.a {
        }
        port pxc-2.a {
        }
    }
    lag "lag-3" {
        admin-state enable
        description "fpe pw-port pxc lag - termination side"
        mode hybrid
        max-ports 64
        access {
            per-fp-ing-queuing true
            per-fp-egr-queuing true
            per-fp-sap-instance true
            adapt-qos {
                mode link
            }
        }
        per-link-hash {
            weighted {
                subscriber-hash-mode vport
            }
        }
        port pxc-1.b {
        }
        port pxc-2.b {
        }
    }
    log {
        filter "1001" {
            named-entry "10" {
                description "Collect only events of major severity or higher"
                action forward
                match {
                    severity {
                        gte major
                    }
                }
            }
        }
        log-id "100" {
            description "Default Serious Errors Log"
            filter "1001"
            source {
                main true
            }
            destination {
                memory {
                    max-entries 500
                }
            }
        }
        log-id "99" {
            description "Default System Log"
            source {
                main true
            }
            destination {
                memory {
                    max-entries 500
                }
            }
        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    }
    port 1/1/c1 {
        admin-state enable
        connector {
            breakout c10-10g
        }
    }
    port 1/1/c1/1 {
        admin-state enable
        description "lag-1 to p-1 access"
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/2 {
        admin-state enable
        description "lag-1 to p-1 access"
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/3 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/4 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/5 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/6 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/7 {
        admin-state enable
        description "to pe-2 network"
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/8 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c1/9 {
        admin-state enable
        ethernet {
            mode hybrid
            dot1x {
                tunneling true
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
        }
    }
    port 1/1/c1/10 {
        admin-state enable
        description "RADIUS and mirroring"
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2 {
        admin-state enable
        connector {
            breakout c1-100g
        }
    }
    port 1/1/c2/1 {
        admin-state enable
        ethernet {
            mode hybrid
            encap-type dot1q
            dot1x {
                tunneling true
            }
        }
    }
    port 1/1/c4 {
        admin-state enable
        connector {
            breakout c4-10g
        }
    }
    port 1/1/c4/1 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c4/2 {
        admin-state enable
        description "PXC 1; pw-port lag 2,3"
        ethernet {
            mode hybrid
            dot1x {
                tunneling true
            }
        }
    }
    port 1/1/c4/3 {
        admin-state enable
        description "PXC 2; pw-port lag 2,3"
        ethernet {
            mode hybrid
            dot1x {
                tunneling true
            }
        }
    }
    port 1/1/c4/4 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    port pxc-1.a {
        admin-state enable
    }
    port pxc-1.b {
        admin-state enable
        description "PXC lag-3, pw-port termination"
        ethernet {
            access {
                egress {
                    virtual-port "vport-1" {
                        aggregate-rate {
                            rate 9000000
                        }
                        host-match {
                            int-dest-id "vport-1" { }
                        }
                    }
                    virtual-port "vport-10" {
                        aggregate-rate {
                            rate 200000
                        }
                        host-match {
                            int-dest-id "vport-10" { }
                        }
                    }
                    virtual-port "vport-11" {
                        aggregate-rate {
                            rate 200000
                        }
                        host-match {
                            int-dest-id "vport-11" { }
                        }
                    }
                    virtual-port "vport-12" {
                        aggregate-rate {
                            rate 200000
                        }
                        host-match {
                            int-dest-id "vport-12" { }
                        }
                    }
                    virtual-port "vport-13" {
                        aggregate-rate {
                            rate 200000
                        }
                        host-match {
                            int-dest-id "vport-13" { }
                        }
                    }
                    virtual-port "vport-14" {
                        aggregate-rate {
                            rate 200000
                        }
                        host-match {
                            int-dest-id "vport-14" { }
                        }
                    }
                    virtual-port "vport-15" {
                        aggregate-rate {
                            rate 200000
                        }
                        host-match {
                            int-dest-id "vport-15" { }
                        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    }
                    virtual-port "vport-16" {
                        aggregate-rate {
                            rate 200000
                        }
                        host-match {
                            int-dest-id "vport-16" { }
                        }
                    }
                    virtual-port "vport-2" {
                        aggregate-rate {
                            rate 3000000
                        }
                        host-match {
                            int-dest-id "vport-2" { }
                        }
                    }
                    virtual-port "vport-3" {
                        aggregate-rate {
                            rate 3000000
                        }
                        host-match {
                            int-dest-id "vport-3" { }
                        }
                    }
                    virtual-port "vport-4" {
                        aggregate-rate {
                            rate 3000000
                        }
                        host-match {
                            int-dest-id "vport-4" { }
                        }
                    }
                    virtual-port "vport-5" {
                        aggregate-rate {
                            rate 3000000
                        }
                        host-match {
                            int-dest-id "vport-5" { }
                        }
                    }
                    virtual-port "vport-6" {
                        aggregate-rate {
                            rate 1000000
                        }
                        host-match {
                            int-dest-id "vport-6" { }
                        }
                    }
                    virtual-port "vport-7" {
                        aggregate-rate {
                            rate 1000000
                        }
                        host-match {
                            int-dest-id "vport-7" { }
                        }
                    }
                    virtual-port "vport-8" {
                        aggregate-rate {
                            rate 1000000
                        }
                        host-match {
                            int-dest-id "vport-8" { }
                        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                }
                virtual-port "vport-9" {
                    aggregate-rate {
                        rate 1400000
                    }
                    host-match {
                        int-dest-id "vport-9" { }
                    }
                }
            }
        }
        egress {
            port-scheduler-policy {
                policy-name "lag-3-pxc"
            }
        }
    }
}
port pxc-2.a {
    admin-state enable
}
port pxc-2.b {
    admin-state enable
    description "PXC lag-3, pw-port termination"
    ethernet {
        access {
            egress {
                virtual-port "vport-1" {
                    aggregate-rate {
                        rate 9000000
                    }
                    host-match {
                        int-dest-id "vport-1" { }
                    }
                }
                virtual-port "vport-10" {
                    aggregate-rate {
                        rate 200000
                    }
                    host-match {
                        int-dest-id "vport-10" { }
                    }
                }
                virtual-port "vport-11" {
                    aggregate-rate {
                        rate 200000
                    }
                    host-match {
                        int-dest-id "vport-11" { }
                    }
                }
                virtual-port "vport-12" {
                    aggregate-rate {
                        rate 200000
                    }
                    host-match {
                        int-dest-id "vport-12" { }
                    }
                }
                virtual-port "vport-13" {
                    aggregate-rate {
                        rate 200000
                    }
                    host-match {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        int-dest-id "vport-13" { }
                    }
                }
                virtual-port "vport-14" {
                    aggregate-rate {
                        rate 200000
                    }
                    host-match {
                        int-dest-id "vport-14" { }
                    }
                }
                virtual-port "vport-15" {
                    aggregate-rate {
                        rate 200000
                    }
                    host-match {
                        int-dest-id "vport-15" { }
                    }
                }
                virtual-port "vport-16" {
                    aggregate-rate {
                        rate 200000
                    }
                    host-match {
                        int-dest-id "vport-16" { }
                    }
                }
                virtual-port "vport-2" {
                    aggregate-rate {
                        rate 3000000
                    }
                    host-match {
                        int-dest-id "vport-2" { }
                    }
                }
                virtual-port "vport-3" {
                    aggregate-rate {
                        rate 3000000
                    }
                    host-match {
                        int-dest-id "vport-3" { }
                    }
                }
                virtual-port "vport-4" {
                    aggregate-rate {
                        rate 3000000
                    }
                    host-match {
                        int-dest-id "vport-4" { }
                    }
                }
                virtual-port "vport-5" {
                    aggregate-rate {
                        rate 3000000
                    }
                    host-match {
                        int-dest-id "vport-5" { }
                    }
                }
                virtual-port "vport-6" {
                    aggregate-rate {
                        rate 1000000
                    }
                    host-match {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                            int-dest-id "vport-6" { }
                        }
                    }
                    virtual-port "vport-7" {
                        aggregate-rate {
                            rate 1000000
                        }
                        host-match {
                            int-dest-id "vport-7" { }
                        }
                    }
                    virtual-port "vport-8" {
                        aggregate-rate {
                            rate 1000000
                        }
                        host-match {
                            int-dest-id "vport-8" { }
                        }
                    }
                    virtual-port "vport-9" {
                        aggregate-rate {
                            rate 1400000
                        }
                        host-match {
                            int-dest-id "vport-9" { }
                        }
                    }
                }
            }
            egress {
                port-scheduler-policy {
                    policy-name "lag-3-pxc"
                }
            }
        }
    }
    port-xc {
        pxc 1 {
            admin-state enable
            description "pw-port lag 2,3"
            port-id 1/1/c4/2
        }
        pxc 2 {
            admin-state enable
            description "pw-port lag 2,3"
            port-id 1/1/c4/3
        }
    }
    pw-port 1 {
        encap-type qinq
        epipe "access" {
            admin-state enable
            fpe-id 1
        }
    }
    python {
        python-script "monitor" {
            admin-state enable
            urls ["ftp://135.231.216.68/pub/configs/alu/SIMS/ACGs/hashing-per-vport-on-pxc/
monitor_lag_sr.py"]
            version python3
        }
        python-script "monitor-sched" {
            admin-state enable
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            urls ["ftp://135.231.216.68/pub/configs/alu/SIMS/ACGs/hashing-per-vport-on-pxc/
monitor_lag_port_scheduler_sr.py"]
            version python3
        }
        python-script "monitor-sched-clear" {
            admin-state enable
            urls ["ftp://135.231.216.68/pub/configs/alu/SIMS/ACGs/hashing-per-vport-on-pxc/
monitor_lag_port_scheduler_clear_sr.py"]
            version python3
        }
    }
    qos {
        sap-egress "sap-egess-1" {
            policy-id 2
            queue 1 {
                port-parent {
                }
            }
        }
        port-scheduler-policy "lag-3-pxc" {
        }
    }
    router "Base" {
        autonomous-system 64500
        router-id 192.0.2.20
        interface "int-1-bng-1-p-1" {
            port lag-1:1
            ipv6 {
                address 2001:db8::501 {
                    prefix-length 120
                }
            }
        }
        interface "system" {
            ipv4 {
                primary {
                    address 192.0.2.20
                    prefix-length 32
                }
            }
            ipv6 {
                address 2001:db8::14 {
                    prefix-length 128
                }
            }
        }
        interface "to-radius" {
            port 1/1/c1/10:114
            ipv4 {
                primary {
                    address 192.168.114.20
                    prefix-length 24
                }
            }
        }
        bgp {
            admin-state enable
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            rapid-update {
                evpn true
            }
            group "evpn" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                peer-as 64500
                local-address 2001:db8::14
                family {
                    evpn true
                }
            }
            neighbor "2001:db8::5" {
                group "evpn"
            }
        }
        isis 0 {
            admin-state enable
            advertise-passive-only false
            advertise-router-capability area
            ipv6-routing native
            level-capability 2
            traffic-engineering true
            area-address [49.0001]
            interface "int-1-bng-1-p-1" {
            }
            interface "system" {
            }
        }
        ldp {
            interface-parameters {
                interface "int-1-bng-1-p-1" {
                    ipv6 {
                        admin-state enable
                    }
                }
            }
        }
        radius {
            server "free-radius-1" {
                address 192.168.114.2
                secret "HDqTwZYSvEu934VNhUQy/pubZxTKpSDzvHg=" hash2
                accept-coa true
            }
        }
    }
    service {
        epipe "access" {
            admin-state enable
            service-id 1
            customer "1"
            bgp 1 {
            }
            bgp-evpn {
                evi 10
                local-attachment-circuit "bng" {
                    eth-tag 2
                }
                remote-attachment-circuit "access" {
                    eth-tag 1
                }
                mpls 1 {
                    admin-state enable
                    send-tunnel-encap {
                        mpls-over-udp false
                    }
                    auto-bind-tunnel {
                        resolution any
                    }
                    route-next-hop {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    system-ipv6
                }
            }
        }
    }
    vpls "capture-sap" {
        admin-state enable
        service-id 5
        customer "1"
        load-balancing {
            per-service-hashing true
        }
        capture-sap pw-1:*.* {
            radius-auth-policy "radius-1"
            trigger-packet {
                dhcp true
            }
            ipoe-session {
                admin-state enable
                ipoe-session-policy "ipoe-session-policy-1"
            }
        }
    }
    vprn "vport-hashing" {
        admin-state enable
        service-id 10
        customer "1"
        interface "loopback-1" {
            admin-state enable
            loopback true
            ipv4 {
                local-dhcp-server "dhcpv4"
                primary {
                    address 192.168.0.1
                    prefix-length 32
                }
                neighbor-discovery {
                    local-proxy-arp false
                    remote-proxy-arp false
                }
                dhcp {
                    admin-state enable
                }
            }
        }
        interface "to-ixia" {
            ipv4 {
                primary {
                    address 172.16.102.1
                    prefix-length 24
                }
            }
            sap 1/1/c1/7:2 {
            }
            ipv6 {
                address 2001:db8::6601 {
                    prefix-length 120
                }
            }
        }
        ipv6 {
            router-advertisement {
                interface "to-ixia" {
                    admin-state enable
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    max-advertisement-interval 15
                    min-advertisement-interval 10
                }
            }
        }
        dhcp-server {
            dhcpv4 "dhcpv4" {
                admin-state enable
                pool-selection {
                    use-gi-address {
                        scope pool
                    }
                    use-pool-from-client {
                    }
                }
                pool "dhcpv4-1" {
                    max-lease-time 1200
                    subnet 10.10.0.0/24 {
                        address-range 10.10.0.10 end 10.10.0.100 {
                            failover-control-type access-driven
                        }
                    }
                }
            }
        }
        subscriber-interface "sub-int-1" {
            admin-state enable
            ipv4 {
                address 10.10.0.254 {
                    prefix-length 24
                }
            }
            group-interface "group-int-1" {
                admin-state enable
                radius-auth-policy "radius-1"
                oper-up-while-empty true
                ipv4 {
                    neighbor-discovery {
                        remote-proxy-arp true
                        populate false
                    }
                    dhcp {
                        admin-state enable
                        server [192.168.0.1]
                        trusted true
                        gi-address 10.10.0.254
                        match-circuit-id true
                        option-82 {
                            action keep
                            vendor-specific-option {
                                pool-name true
                            }
                        }
                        lease-populate {
                            max-leases 100
                        }
                        client-applications {
                            dhcp true
                        }
                    }
                }
                ipoe-session {
                    admin-state enable
                    ipoe-session-policy "ipoe-session-policy-1"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    sap-session-limit 100
                    force-auth {
                        cid-change false
                        rid-change false
                    }
                }
            }
        }
    }
    sfm 1 {
        sfm-type m-sfm6-7/12
    }
    subscriber-mgmt {
        ipoe-session-policy "ipoe-session-policy-1" {
        }
        sub-profile "sub-profile-1" {
            egress {
                qos {
                    agg-rate {
                        rate 1000
                    }
                }
            }
        }
        sla-profile "sla-profile-1" {
            egress {
                qos {
                    sap-egress {
                        policy-name "sap-egess-1"
                        port-parent-location vport
                        overrides {
                            queue 1 {
                                stat-mode v4-v6
                            }
                        }
                    }
                }
            }
        }
        sub-ident-policy "sub-ident-policy-1" {
            sla-profile-map {
                use-direct-map-as-default true
            }
            sub-profile-map {
                use-direct-map-as-default true
            }
        }
        radius-authentication-policy "radius-1" {
            password "ncd8qyrNUMhYfa2SfrUqHMDZ9IXn3sVSmYBzbw==" hash2
            pppoe-access-method pap-chap
            radius-server-policy "radius-server-1"
            user-name {
                format circuit-id
            }
            include-radius-attribute {
                circuit-id true
                nas-identifier true
            }
        }
        msap-policy "msap-policy-1" {
            sub-sla-mgmt {
                subscriber-limit 1
                sub-ident-policy "sub-ident-policy-1"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                defaults {
                    subscriber-id {
                        sap-id
                    }
                }
                single-sub-parameters {
                    profiled-traffic-only true
                }
            }
            ies-vprn-only-sap-parameters {
                anti-spoof next-hop-ip-and-mac-addr
                ingress {
                    qos {
                        queuing-type service
                    }
                }
            }
        }
    }
    system {
        name "bng-1"
        management-interface {
            configuration-mode model-driven
            cli {
                cli-engine [md-cli classic-cli]
                md-cli {
                    environment {
                        more false
                        command-alias {
                            alias "monitor-lag" {
                                admin-state enable
                                python-script "monitor"
                                mount-point global { }
                            }
                            alias "monitor-schd-clear" {
                                admin-state enable
                                python-script "monitor-sched-clear"
                                mount-point global { }
                            }
                            alias "monitor-scheduler" {
                                admin-state enable
                                python-script "monitor-sched"
                                mount-point global { }
                            }
                        }
                    }
                }
            }
            netconf {
                admin-state enable
            }
            yang-modules {
                nokia-submodules false
                nokia-combined-modules true
            }
            snmp {
                admin-state disable
            }
        }
        login-control {
            idle-timeout none
        }
        security {
            aaa {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                local-profiles {
                    profile "administrative" {
                        default-action permit-all
                        entry 10 {
                            match "configure system security"
                            action permit
                        }
                        entry 20 {
                            match "show system security"
                            action permit
                        }
                        entry 30 {
                            match "tools perform security"
                            action permit
                        }
                        entry 40 {
                            match "tools dump security"
                            action permit
                        }
                        entry 50 {
                            match "admin system security"
                            action permit
                        }
                        entry 100 {
                            match "configure li"
                            action deny
                        }
                        entry 110 {
                            match "show li"
                            action deny
                        }
                        entry 111 {
                            match "clear li"
                            action deny
                        }
                        entry 112 {
                            match "tools dump li"
                            action deny
                        }
                        netconf {
                            base-op-authorization {
                                action true
                                cancel-commit true
                                close-session true
                                commit true
                                copy-config true
                                create-subscription true
                                delete-config true
                                discard-changes true
                                edit-config true
                                get true
                                get-config true
                                get-data true
                                get-schema true
                                kill-session true
                                lock true
                                validate true
                            }
                        }
                    }
                    profile "default" {
                        entry 10 {
                            match "exec"
                            action permit
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    }
                    entry 20 {
                        match "exit"
                        action permit
                    }
                    entry 30 {
                        match "help"
                        action permit
                    }
                    entry 40 {
                        match "logout"
                        action permit
                    }
                    entry 50 {
                        match "password"
                        action permit
                    }
                    entry 60 {
                        match "show config"
                        action deny
                    }
                    entry 65 {
                        match "show li"
                        action deny
                    }
                    entry 66 {
                        match "clear li"
                        action deny
                    }
                    entry 67 {
                        match "tools dump li"
                        action deny
                    }
                    entry 68 {
                        match "state li"
                        action deny
                    }
                    entry 70 {
                        match "show"
                        action permit
                    }
                    entry 75 {
                        match "state"
                        action permit
                    }
                    entry 80 {
                        match "enable-admin"
                        action permit
                    }
                    entry 90 {
                        match "enable"
                        action permit
                    }
                    entry 100 {
                        match "configure li"
                        action deny
                    }
                }
            }
        }
        ssh {
            server-admin-state enable
            server-cipher-list-v2 {
                cipher 190 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        name aes256-ctr
                    }
                    cipher 192 {
                        name aes192-ctr
                    }
                    cipher 194 {
                        name aes128-ctr
                    }
                    cipher 200 {
                        name aes128-cbc
                    }
                    cipher 205 {
                        name 3des-cbc
                    }
                    cipher 225 {
                        name aes192-cbc
                    }
                    cipher 230 {
                        name aes256-cbc
                    }
                }
                client-cipher-list-v2 {
                    cipher 190 {
                        name aes256-ctr
                    }
                    cipher 192 {
                        name aes192-ctr
                    }
                    cipher 194 {
                        name aes128-ctr
                    }
                    cipher 200 {
                        name aes128-cbc
                    }
                    cipher 205 {
                        name 3des-cbc
                    }
                    cipher 225 {
                        name aes192-cbc
                    }
                    cipher 230 {
                        name aes256-cbc
                    }
                }
                server-mac-list-v2 {
                    mac 200 {
                        name hmac-sha2-512
                    }
                    mac 210 {
                        name hmac-sha2-256
                    }
                    mac 215 {
                        name hmac-sha1
                    }
                    mac 220 {
                        name hmac-sha1-96
                    }
                    mac 225 {
                        name hmac-md5
                    }
                    mac 240 {
                        name hmac-md5-96
                    }
                }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                client-mac-list-v2 {
                    mac 200 {
                        name hmac-sha2-512
                    }
                    mac 210 {
                        name hmac-sha2-256
                    }
                    mac 215 {
                        name hmac-sha1
                    }
                    mac 220 {
                        name hmac-sha1-96
                    }
                    mac 225 {
                        name hmac-md5
                    }
                    mac 240 {
                        name hmac-md5-96
                    }
                }
            }
            user-params {
                local-user {
                    user "admin" {
                        password "$2y$10$TQrZlpBDra86.qoexZUzQeBXDY1FcdDhGWdD9lLxMuFyPVSm0OGy6"
                        access {
                            console true
                            netconf true
                        }
                        console {
                            member ["administrative"]
                        }
                    }
                }
            }
        }
        time {
            prefer-local-time true
            zone {
                standard {
                    name cst
                }
            }
            dst-zone "CDT" {
                end {
                    day sunday
                    month november
                    hours-minutes "02:00"
                }
                start {
                    day sunday
                    month march
                    hours-minutes "02:00"
                }
            }
            sntp {
                admin-state enable
                server 135.227.160.253 {
                }
            }
        }
    }
}
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
# Finished 2022-12-14T13:38:31.1-06:00
```

bng-vport-per-subscriber-balancing

```
# TiMOS-C-22.10.R1 cpm/x86_64 Nokia 7750 SR Copyright (c) 2000-2022 Nokia.
# All rights reserved. All use subject to applicable license agreements.
# Built on Sun Oct 30 14:49:55 PDT 2022 by builder in /builds/c/2210B/R1/panos/main/sros
# Configuration format version 22.10 revision 0

# Generated 2022-12-14T13:42:02.4-06:00 by admin from 135.231.208.32
# Commit ID 5
#   Committed 2022-12-14T13:41:47.9-06:00 by admin (MD-CLI) from 135.231.208.32
# Commit ID 4
#   Committed 2022-12-14T13:41:11.6-06:00 by admin (MD-CLI) from 135.231.208.32
# Commit ID 3
#   Committed 2022-12-14T13:33:11.3-06:00 by admin (MD-CLI) from 135.231.208.32

configure {
    aaa {
        radius {
            server-policy "radius-server-1" {
                servers {
                    router-instance "Base"
                    server 1 {
                        server-name "free-radius-1"
                    }
                }
            }
        }
    }
    card 1 {
        card-type iom5-e
        mda 1 {
            mda-type me6-100gb-qsfp28
        }
        mda 2 {
            mda-type me6-100gb-qsfp28
        }
    }
    fwd-path-ext {
        sdp-id-range {
            start 17500
            end 17600
        }
        fpe 1 {
            description "pw-port on a lag"
            path {
                xc-lag-a "lag-2"
                xc-lag-b "lag-3"
            }
            application {
                pw-port-extension {
                }
            }
        }
    }
    lag "lag-1" {
        admin-state enable
        description "lag to p-1"
        mode hybrid
        max-ports 64
        port 1/1/c1/1 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
        port 1/1/c1/2 {
        }
    }
    lag "lag-2" {
        admin-state enable
        description "fpe pw-port pxc lag - transit side"
        mode hybrid
        max-ports 64
        port pxc-1.a {
        }
        port pxc-2.a {
        }
    }
    lag "lag-3" {
        admin-state enable
        description "fpe pw-port pxc lag - termination side"
        mode hybrid
        max-ports 64
        access {
            per-fp-ing-queuing false
            per-fp-egr-queuing false
            per-fp-sap-instance false
            adapt-qos {
                mode port-fair
            }
        }
        port pxc-1.b {
        }
        port pxc-2.b {
        }
    }
    log {
        filter "1001" {
            named-entry "10" {
                description "Collect only events of major severity or higher"
                action forward
                match {
                    severity {
                        gte major
                    }
                }
            }
        }
        log-id "100" {
            description "Default Serious Errors Log"
            filter "1001"
            source {
                main true
            }
            destination {
                memory {
                    max-entries 500
                }
            }
        }
        log-id "99" {
            description "Default System Log"
            source {
                main true
            }
            destination {
                memory {
                    max-entries 500
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            }
        }
    }
}
port 1/1/c1 {
    admin-state enable
    connector {
        breakout c10-10g
    }
}
port 1/1/c1/1 {
    admin-state enable
    description "lag-1 to p-1 access"
    ethernet {
        mode hybrid
    }
}
port 1/1/c1/2 {
    admin-state enable
    description "lag-1 to p-1 access"
    ethernet {
        mode hybrid
    }
}
port 1/1/c1/3 {
    admin-state enable
    ethernet {
        mode hybrid
    }
}
port 1/1/c1/4 {
    admin-state enable
    ethernet {
        mode hybrid
    }
}
port 1/1/c1/5 {
    admin-state enable
    ethernet {
        mode hybrid
    }
}
port 1/1/c1/6 {
    admin-state enable
    ethernet {
        mode hybrid
    }
}
port 1/1/c1/7 {
    admin-state enable
    description "to pe-2 network"
    ethernet {
        mode hybrid
    }
}
port 1/1/c1/8 {
    admin-state enable
    ethernet {
        mode hybrid
    }
}
port 1/1/c1/9 {
    admin-state enable
    ethernet {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            mode hybrid
            dot1x {
                tunneling true
            }
        }
    }
    port 1/1/c1/10 {
        admin-state enable
        description "RADIUS and mirroring"
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c2 {
        admin-state enable
        connector {
            breakout c1-100g
        }
    }
    port 1/1/c2/1 {
        admin-state enable
        ethernet {
            mode hybrid
            encap-type dot1q
            dot1x {
                tunneling true
            }
        }
    }
    port 1/1/c4 {
        admin-state enable
        connector {
            breakout c4-10g
        }
    }
    port 1/1/c4/1 {
        admin-state enable
        ethernet {
            mode hybrid
        }
    }
    port 1/1/c4/2 {
        admin-state enable
        description "PXC 1; pw-port lag 2,3"
        ethernet {
            mode hybrid
            dot1x {
                tunneling true
            }
        }
    }
    port 1/1/c4/3 {
        admin-state enable
        description "PXC 2; pw-port lag 2,3"
        ethernet {
            mode hybrid
            dot1x {
                tunneling true
            }
        }
    }
    port 1/1/c4/4 {
        admin-state enable
        ethernet {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            mode hybrid
        }
    }
    port pxc-1.a {
        admin-state enable
    }
    port pxc-1.b {
        admin-state enable
        description "PXC lag-3, pw-port termination"
        ethernet {
            access {
                egress {
                    virtual-port "vport-1" {
                        aggregate-rate {
                            rate 1500000
                        }
                        host-match {
                            int-dest-id "vport-1" { }
                        }
                    }
                    virtual-port "vport-10" {
                        aggregate-rate {
                            rate 9000000
                        }
                        host-match {
                            int-dest-id "vport-10" { }
                        }
                    }
                    virtual-port "vport-11" {
                        aggregate-rate {
                            rate 3000000
                        }
                        host-match {
                            int-dest-id "vport-11" { }
                        }
                    }
                    virtual-port "vport-12" {
                        aggregate-rate {
                            rate 3000000
                        }
                        host-match {
                            int-dest-id "vport-12" { }
                        }
                    }
                    virtual-port "vport-13" {
                        aggregate-rate {
                            rate 3000000
                        }
                        host-match {
                            int-dest-id "vport-13" { }
                        }
                    }
                    virtual-port "vport-14" {
                        aggregate-rate {
                            rate 3000000
                        }
                        host-match {
                            int-dest-id "vport-14" { }
                        }
                    }
                    virtual-port "vport-15" {
                        aggregate-rate {
                            rate 3000000
                        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        host-match {
                            int-dest-id "vport-15" { }
                        }
                    }
                    virtual-port "vport-16" {
                        aggregate-rate {
                            rate 3000000
                        }
                        host-match {
                            int-dest-id "vport-16" { }
                        }
                    }
                    virtual-port "vport-2" {
                        aggregate-rate {
                            rate 500000
                        }
                        host-match {
                            int-dest-id "vport-2" { }
                        }
                    }
                    virtual-port "vport-3" {
                        aggregate-rate {
                            rate 500000
                        }
                        host-match {
                            int-dest-id "vport-3" { }
                        }
                    }
                    virtual-port "vport-4" {
                        aggregate-rate {
                            rate 500000
                        }
                        host-match {
                            int-dest-id "vport-4" { }
                        }
                    }
                    virtual-port "vport-5" {
                        aggregate-rate {
                            rate 4500000
                        }
                        host-match {
                            int-dest-id "vport-5" { }
                        }
                    }
                    virtual-port "vport-6" {
                        aggregate-rate {
                            rate 1500000
                        }
                        host-match {
                            int-dest-id "vport-6" { }
                        }
                    }
                    virtual-port "vport-7" {
                        aggregate-rate {
                            rate 1500000
                        }
                        host-match {
                            int-dest-id "vport-7" { }
                        }
                    }
                    virtual-port "vport-8" {
                        aggregate-rate {
                            rate 1500000
                        }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    host-match {
                        int-dest-id "vport-8" { }
                    }
                }
                virtual-port "vport-9" {
                    aggregate-rate {
                        rate 9000000
                    }
                    host-match {
                        int-dest-id "vport-9" { }
                    }
                }
            }
        }
        egress {
            port-scheduler-policy {
                policy-name "lag-3-pxc"
            }
        }
    }
}
port pxc-2.a {
    admin-state enable
}
port pxc-2.b {
    admin-state enable
    description "PXC lag-3, pw-port termination"
    ethernet {
        access {
            egress {
                virtual-port "vport-1" {
                    aggregate-rate {
                        rate 1500000
                    }
                    host-match {
                        int-dest-id "vport-1" { }
                    }
                }
                virtual-port "vport-10" {
                    aggregate-rate {
                        rate 9000000
                    }
                    host-match {
                        int-dest-id "vport-10" { }
                    }
                }
                virtual-port "vport-11" {
                    aggregate-rate {
                        rate 3000000
                    }
                    host-match {
                        int-dest-id "vport-11" { }
                    }
                }
                virtual-port "vport-12" {
                    aggregate-rate {
                        rate 3000000
                    }
                    host-match {
                        int-dest-id "vport-12" { }
                    }
                }
                virtual-port "vport-13" {
                    aggregate-rate {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        rate 3000000
                    }
                    host-match {
                        int-dest-id "vport-13" { }
                    }
                }
                virtual-port "vport-14" {
                    aggregate-rate {
                        rate 3000000
                    }
                    host-match {
                        int-dest-id "vport-14" { }
                    }
                }
                virtual-port "vport-15" {
                    aggregate-rate {
                        rate 3000000
                    }
                    host-match {
                        int-dest-id "vport-15" { }
                    }
                }
                virtual-port "vport-16" {
                    aggregate-rate {
                        rate 3000000
                    }
                    host-match {
                        int-dest-id "vport-16" { }
                    }
                }
                virtual-port "vport-2" {
                    aggregate-rate {
                        rate 500000
                    }
                    host-match {
                        int-dest-id "vport-2" { }
                    }
                }
                virtual-port "vport-3" {
                    aggregate-rate {
                        rate 500000
                    }
                    host-match {
                        int-dest-id "vport-3" { }
                    }
                }
                virtual-port "vport-4" {
                    aggregate-rate {
                        rate 500000
                    }
                    host-match {
                        int-dest-id "vport-4" { }
                    }
                }
                virtual-port "vport-5" {
                    aggregate-rate {
                        rate 4500000
                    }
                    host-match {
                        int-dest-id "vport-5" { }
                    }
                }
                virtual-port "vport-6" {
                    aggregate-rate {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                            rate 1500000
                        }
                        host-match {
                            int-dest-id "vport-6" { }
                        }
                    }
                    virtual-port "vport-7" {
                        aggregate-rate {
                            rate 1500000
                        }
                        host-match {
                            int-dest-id "vport-7" { }
                        }
                    }
                    virtual-port "vport-8" {
                        aggregate-rate {
                            rate 1500000
                        }
                        host-match {
                            int-dest-id "vport-8" { }
                        }
                    }
                    virtual-port "vport-9" {
                        aggregate-rate {
                            rate 9000000
                        }
                        host-match {
                            int-dest-id "vport-9" { }
                        }
                    }
                }
            }
            egress {
                port-scheduler-policy {
                    policy-name "lag-3-pxc"
                }
            }
        }
    }
    port-xc {
        pxc 1 {
            admin-state enable
            description "pw-port lag 2,3"
            port-id 1/1/c4/2
        }
        pxc 2 {
            admin-state enable
            description "pw-port lag 2,3"
            port-id 1/1/c4/3
        }
    }
    pw-port 1 {
        encap-type qinq
        epipe "access" {
            admin-state enable
            fpe-id 1
        }
    }
    python {
        python-script "monitor" {
            admin-state enable
            urls ["ftp://135.231.216.68/pub/configs/alu/SIMS/ACGs/hashing-per-vport-on-pxc/
monitor_lag_sr.py"]
            version python3
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
        python-script "monitor-sched" {
            admin-state enable
            urls ["ftp://135.231.216.68/pub/configs/alu/SIMS/ACGs/hashing-per-vport-on-pxc/
monitor_lag_port_scheduler_sr.py"]
            version python3
        }
        python-script "monitor-sched-clear" {
            admin-state enable
            urls ["ftp://135.231.216.68/pub/configs/alu/SIMS/ACGs/hashing-per-vport-on-pxc/
monitor_lag_port_scheduler_clear_sr.py"]
            version python3
        }
    }
    qos {
        sap-egress "sap-egess-1" {
            policy-id 2
            queue 1 {
                port-parent {
                }
            }
        }
        port-scheduler-policy "lag-3-pxc" {
        }
    }
    router "Base" {
        autonomous-system 64500
        router-id 192.0.2.20
        interface "int-1-bng-1-p-1" {
            port lag-1:1
            ipv6 {
                address 2001:db8::501 {
                    prefix-length 120
                }
            }
        }
        interface "system" {
            ipv4 {
                primary {
                    address 192.0.2.20
                    prefix-length 32
                }
            }
            ipv6 {
                address 2001:db8::14 {
                    prefix-length 128
                }
            }
        }
        interface "to-radius" {
            port 1/1/c1/10:114
            ipv4 {
                primary {
                    address 192.168.114.20
                    prefix-length 24
                }
            }
        }
        bgp {
            admin-state enable
            vpn-apply-export true
            vpn-apply-import true
            rapid-withdrawal true
            rapid-update {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            evpn true
        }
        group "evpn" {
            peer-as 64500
            local-address 2001:db8::14
            family {
                evpn true
            }
        }
        neighbor "2001:db8::5" {
            group "evpn"
        }
    }
    isis 0 {
        admin-state enable
        advertise-passive-only false
        advertise-router-capability area
        ipv6-routing native
        level-capability 2
        traffic-engineering true
        area-address [49.0001]
        interface "int-1-bng-1-p-1" {
        }
        interface "system" {
        }
    }
    ldp {
        interface-parameters {
            interface "int-1-bng-1-p-1" {
                ipv6 {
                    admin-state enable
                }
            }
        }
    }
    radius {
        server "free-radius-1" {
            address 192.168.114.2
            secret "HDqTwZYSvEu934VNhUQy/pubZxTKpSDzvHg=" hash2
            accept-coa true
        }
    }
}
service {
    epipe "access" {
        admin-state enable
        service-id 1
        customer "1"
        bgp 1 {
        }
        bgp-evpn {
            evi 10
            local-attachment-circuit "bng" {
                eth-tag 2
            }
            remote-attachment-circuit "access" {
                eth-tag 1
            }
            mpls 1 {
                admin-state enable
                send-tunnel-encap {
                    mpls-over-udp false
                }
                auto-bind-tunnel {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    resolution any
                }
                route-next-hop {
                    system-ipv6
                }
            }
        }
    }
    vpls "capture-sap" {
        admin-state enable
        service-id 5
        customer "1"
        load-balancing {
            per-service-hashing true
        }
        capture-sap pw-1:*.* {
            radius-auth-policy "radius-1"
            trigger-packet {
                dhcp true
            }
            ipoe-session {
                admin-state enable
                ipoe-session-policy "ipoe-session-policy-1"
            }
        }
    }
    vprn "vport-hashing" {
        admin-state enable
        service-id 10
        customer "1"
        interface "loopback-1" {
            admin-state enable
            loopback true
            ipv4 {
                local-dhcp-server "dhcpv4"
                primary {
                    address 192.168.0.1
                    prefix-length 32
                }
                neighbor-discovery {
                    local-proxy-arp false
                    remote-proxy-arp false
                }
                dhcp {
                    admin-state enable
                }
            }
        }
        interface "to-ixia" {
            ipv4 {
                primary {
                    address 172.16.102.1
                    prefix-length 24
                }
            }
            sap 1/1/c1/7:2 {
            }
            ipv6 {
                address 2001:db8::6601 {
                    prefix-length 120
                }
            }
        }
        ipv6 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                router-advertisement {
                    interface "to-ixia" {
                        admin-state enable
                        max-advertisement-interval 15
                        min-advertisement-interval 10
                    }
                }
            }
            dhcp-server {
                dhcpv4 "dhcpv4" {
                    admin-state enable
                    pool-selection {
                        use-gi-address {
                            scope pool
                        }
                        use-pool-from-client {
                        }
                    }
                    pool "dhcpv4-1" {
                        max-lease-time 1200
                        subnet 10.10.0.0/24 {
                            address-range 10.10.0.10 end 10.10.0.100 {
                                failover-control-type access-driven
                            }
                        }
                    }
                }
            }
            subscriber-interface "sub-int-1" {
                admin-state enable
                ipv4 {
                    address 10.10.0.254 {
                        prefix-length 24
                    }
                }
                group-interface "group-int-1" {
                    admin-state enable
                    radius-auth-policy "radius-1"
                    oper-up-while-empty true
                    ipv4 {
                        neighbor-discovery {
                            remote-proxy-arp true
                            populate false
                        }
                        dhcp {
                            admin-state enable
                            server [192.168.0.1]
                            trusted true
                            gi-address 10.10.0.254
                            match-circuit-id true
                            option-82 {
                                action keep
                                vendor-specific-option {
                                    pool-name true
                                }
                            }
                            lease-populate {
                                max-leases 100
                            }
                            client-applications {
                                dhcp true
                            }
                        }
                    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    ipoe-session {
                        admin-state enable
                        ipoe-session-policy "ipoe-session-policy-1"
                        sap-session-limit 100
                        force-auth {
                            cid-change false
                            rid-change false
                        }
                    }
                }
            }
        }
    }
    sfm 1 {
        sfm-type m-sfm6-7/12
    }
    subscriber-mgmt {
        ipoe-session-policy "ipoe-session-policy-1" {
        }
        sub-profile "sub-profile-1" {
            egress {
                qos {
                    agg-rate {
                        rate 1000
                    }
                }
            }
        }
        sla-profile "sla-profile-1" {
            egress {
                qos {
                    sap-egress {
                        policy-name "sap-egess-1"
                        port-parent-location vport
                        overrides {
                            queue 1 {
                                stat-mode v4-v6
                            }
                        }
                    }
                }
            }
        }
        sub-ident-policy "sub-ident-policy-1" {
            sla-profile-map {
                use-direct-map-as-default true
            }
            sub-profile-map {
                use-direct-map-as-default true
            }
        }
        radius-authentication-policy "radius-1" {
            password "ncd8qyrNUMhYfa2SfrUqHMDZ9IXn3sVSmYBzbw==" hash2
            pppoe-access-method pap-chap
            radius-server-policy "radius-server-1"
            user-name {
                format circuit-id
            }
            include-radius-attribute {
                circuit-id true
                nas-identifier true
            }
        }
        msap-policy "msap-policy-1" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
            sub-sla-mgmt {
                subscriber-limit 1
                sub-ident-policy "sub-ident-policy-1"
                defaults {
                    subscriber-id {
                        sap-id
                    }
                }
                single-sub-parameters {
                    profiled-traffic-only true
                }
            }
            ies-vprn-only-sap-parameters {
                anti-spoof next-hop-ip-and-mac-addr
                ingress {
                    qos {
                        queuing-type service
                    }
                }
            }
        }
    }
    system {
        name "bng-1"
        management-interface {
            configuration-mode model-driven
            cli {
                cli-engine [md-cli classic-cli]
                md-cli {
                    environment {
                        more false
                        command-alias {
                            alias "monitor-lag" {
                                admin-state enable
                                python-script "monitor"
                                mount-point global { }
                            }
                            alias "monitor-schd-clear" {
                                admin-state enable
                                python-script "monitor-sched-clear"
                                mount-point global { }
                            }
                            alias "monitor-scheduler" {
                                admin-state enable
                                python-script "monitor-sched"
                                mount-point global { }
                            }
                        }
                    }
                }
            }
            netconf {
                admin-state enable
            }
            yang-modules {
                nokia-submodules false
                nokia-combined-modules true
            }
            snmp {
                admin-state disable
            }
        }
        login-control {
            idle-timeout none
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
        }
        security {
            aaa {
                local-profiles {
                    profile "administrative" {
                        default-action permit-all
                        entry 10 {
                            match "configure system security"
                            action permit
                        }
                        entry 20 {
                            match "show system security"
                            action permit
                        }
                        entry 30 {
                            match "tools perform security"
                            action permit
                        }
                        entry 40 {
                            match "tools dump security"
                            action permit
                        }
                        entry 50 {
                            match "admin system security"
                            action permit
                        }
                        entry 100 {
                            match "configure li"
                            action deny
                        }
                        entry 110 {
                            match "show li"
                            action deny
                        }
                        entry 111 {
                            match "clear li"
                            action deny
                        }
                        entry 112 {
                            match "tools dump li"
                            action deny
                        }
                        netconf {
                            base-op-authorization {
                                action true
                                cancel-commit true
                                close-session true
                                commit true
                                copy-config true
                                create-subscription true
                                delete-config true
                                discard-changes true
                                edit-config true
                                get true
                                get-config true
                                get-data true
                                get-schema true
                                kill-session true
                                lock true
                                validate true
                            }
                        }
                    }
                    profile "default" {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        entry 10 {
                            match "exec"
                            action permit
                        }
                        entry 20 {
                            match "exit"
                            action permit
                        }
                        entry 30 {
                            match "help"
                            action permit
                        }
                        entry 40 {
                            match "logout"
                            action permit
                        }
                        entry 50 {
                            match "password"
                            action permit
                        }
                        entry 60 {
                            match "show config"
                            action deny
                        }
                        entry 65 {
                            match "show li"
                            action deny
                        }
                        entry 66 {
                            match "clear li"
                            action deny
                        }
                        entry 67 {
                            match "tools dump li"
                            action deny
                        }
                        entry 68 {
                            match "state li"
                            action deny
                        }
                        entry 70 {
                            match "show"
                            action permit
                        }
                        entry 75 {
                            match "state"
                            action permit
                        }
                        entry 80 {
                            match "enable-admin"
                            action permit
                        }
                        entry 90 {
                            match "enable"
                            action permit
                        }
                        entry 100 {
                            match "configure li"
                            action deny
                        }
                    }
                }
            }
            ssh {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                server-admin-state enable
                server-cipher-list-v2 {
                    cipher 190 {
                        name aes256-ctr
                    }
                    cipher 192 {
                        name aes192-ctr
                    }
                    cipher 194 {
                        name aes128-ctr
                    }
                    cipher 200 {
                        name aes128-cbc
                    }
                    cipher 205 {
                        name 3des-cbc
                    }
                    cipher 225 {
                        name aes192-cbc
                    }
                    cipher 230 {
                        name aes256-cbc
                    }
                }
                client-cipher-list-v2 {
                    cipher 190 {
                        name aes256-ctr
                    }
                    cipher 192 {
                        name aes192-ctr
                    }
                    cipher 194 {
                        name aes128-ctr
                    }
                    cipher 200 {
                        name aes128-cbc
                    }
                    cipher 205 {
                        name 3des-cbc
                    }
                    cipher 225 {
                        name aes192-cbc
                    }
                    cipher 230 {
                        name aes256-cbc
                    }
                }
                server-mac-list-v2 {
                    mac 200 {
                        name hmac-sha2-512
                    }
                    mac 210 {
                        name hmac-sha2-256
                    }
                    mac 215 {
                        name hmac-sha1
                    }
                    mac 220 {
                        name hmac-sha1-96
                    }
                    mac 225 {
                        name hmac-md5
                    }
                    mac 240 {
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                    name hmac-md5-96
                }
            }
            client-mac-list-v2 {
                mac 200 {
                    name hmac-sha2-512
                }
                mac 210 {
                    name hmac-sha2-256
                }
                mac 215 {
                    name hmac-sha1
                }
                mac 220 {
                    name hmac-sha1-96
                }
                mac 225 {
                    name hmac-md5
                }
                mac 240 {
                    name hmac-md5-96
                }
            }
        }
        user-params {
            local-user {
                user "admin" {
                    password "$2y$10$TQrZlpBDra86.qoexZUzQeBXDY1FcdDhGWdD9lLxMuFyPVSm0OGy6"
                    access {
                        console true
                        netconf true
                    }
                    console {
                        member ["administrative"]
                    }
                }
            }
        }
    }
    time {
        prefer-local-time true
        zone {
            standard {
                name cst
            }
        }
        dst-zone "CDT" {
            end {
                day sunday
                month november
                hours-minutes "02:00"
            }
            start {
                day sunday
                month march
                hours-minutes "02:00"
            }
        }
        sntp {
            admin-state enable
            server 135.227.160.253 {
            }
        }
    }
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
    }
}


# Finished 2022-12-14T13:42:02.4-06:00
```

radius users

```
cid-1                   Cleartext-Password := "cse-password"
                        Alc-Subsc-ID-Str = "ipoe-ds-1",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-1",
                        Fall-Through = No

cid-2                   Cleartext-Password := "cse-password"
                        Alc-Subsc-ID-Str = "ipoe-ds-2",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-1",
                        Fall-Through = No

cid-3                   Cleartext-Password := "cse-password"
                        Alc-Subsc-ID-Str = "ipoe-ds-3",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-2",
                        Fall-Through = No

cid-4                   Cleartext-Password := "cse-password"
                        Alc-Subsc-ID-Str = "ipoe-ds-4",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-2",
                        Fall-Through = No

cid-5                   Cleartext-Password := "cse-password"
                        Alc-Subsc-ID-Str = "ipoe-ds-5",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-3",
                        Fall-Through = No

cid-6                   Cleartext-Password := "cse-password"
                        Alc-Subsc-ID-Str = "ipoe-ds-6",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-3",
                        Fall-Through = No

cid-7                   Cleartext-Password := "cse-password"
                        Alc-Subsc-ID-Str = "ipoe-ds-7",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-4",
                        Fall-Through = No

cid-8                   Cleartext-Password := "cse-password"
                        Alc-Subsc-ID-Str = "ipoe-ds-8",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-4",
                        Fall-Through = No

cid-9                   Cleartext-Password := "cse-password"
                        Alc-Subsc-ID-Str = "ipoe-ds-9",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-5",
                        Fall-Through = No

cid-10                  Cleartext-Password := "cse-password"
                        Alc-Subsc-ID-Str = "ipoe-ds-10",
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-5",
                        Fall-Through = No

cid-11                  Cleartext-Password := "cse-password"
                        Alc-Subsc-ID-Str = "ipoe-ds-11",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-6",
                        Fall-Through = No

cid-12                  Cleartext-Password := "cse-password"
                        Alc-Subsc-ID-Str = "ipoe-ds-12",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-6",
                        Fall-Through = No

cid-13                  Cleartext-Password := "cse-password"
                        Alc-Subsc-ID-Str = "ipoe-ds-13",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-7",
                        Fall-Through = No

cid-14                  Cleartext-Password := "cse-password"
                        Alc-Subsc-ID-Str = "ipoe-ds-14",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-7",
                        Fall-Through = No

cid-15                  Cleartext-Password := "cse-password"
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        Alc-Subsc-ID-Str = "ipoe-ds-15",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-8",
                        Fall-Through = No

cid-16                  Cleartext-Password := "cse-password"
                        Alc-Subsc-ID-Str = "ipoe-ds-16",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-8",
                        Fall-Through = No

cid-17                  Cleartext-Password := "cse-password"
                        Alc-Subsc-ID-Str = "ipoe-ds-17",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-9",
                        Fall-Through = No

cid-18                  Cleartext-Password := "cse-password"
                        Alc-Subsc-ID-Str = "ipoe-ds-18",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-9",
                        Fall-Through = No

cid-19                  Cleartext-Password := "cse-password"
                        Alc-Subsc-ID-Str = "ipoe-ds-19",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-10",
                        Fall-Through = No
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
cid-20                  Cleartext-Password := "cse-password"
                        Alc-Subsc-ID-Str = "ipoe-ds-20",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-10",
                        Fall-Through = No

cid-21                  Cleartext-Password := "cse-password"
                        Alc-Subsc-ID-Str = "ipoe-ds-21",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-11",
                        Fall-Through = No

cid-22                  Cleartext-Password := "cse-password"
                        Alc-Subsc-ID-Str = "ipoe-ds-22",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-11",
                        Fall-Through = No

cid-23                  Cleartext-Password := "cse-password"
                        Alc-Subsc-ID-Str = "ipoe-ds-23",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-12",
                        Fall-Through = No

cid-24                   Cleartext-Password := "cse-password"
                        Alc-Subsc-ID-Str = "ipoe-ds-24",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-12",
                        Fall-Through = No
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
cid-25                  Cleartext-Password := "cse-password"
                        Alc-Subsc-ID-Str = "ipoe-ds-25",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-13",
                        Fall-Through = No

cid-26                  Cleartext-Password := "cse-password"
                        Alc-Subsc-ID-Str = "ipoe-ds-26",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-13",
                        Fall-Through = No

cid-27                  Cleartext-Password := "cse-password"
                        Alc-Subsc-ID-Str = "ipoe-ds-27",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-14",
                        Fall-Through = No

cid-28                  Cleartext-Password := "cse-password"
                        Alc-Subsc-ID-Str = "ipoe-ds-28",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-14",
                        Fall-Through = No

cid-29                  Cleartext-Password := "cse-password"
                        Alc-Subsc-ID-Str = "ipoe-ds-29",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-15",
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

```
                        Fall-Through = No

cid-30                  Cleartext-Password := "cse-password"
                        Alc-Subsc-ID-Str = "ipoe-ds-30",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-15",
                        Fall-Through = No

cid-31                  Cleartext-Password := "cse-password"
                        Alc-Subsc-ID-Str = "ipoe-ds-31",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-16",
                        Fall-Through = No

cid-32                  Cleartext-Password := "cse-password"
                        Alc-Subsc-ID-Str = "ipoe-ds-32",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-MSAP-Interface = "group-int-1",
                        Alc-MSAP-Policy = "msap-policy-1",
                        Alc-MSAP-Serv-Id = "10",
                        Framed-Pool = "dhcpv4-1",
                        Framed-IPv6-Pool = "dhcpv6-1",
                        Alc-Delegated-IPv6-Pool = "dhcpv6-1",
                        Alc-Int-Dest-Id-Str = "vport-16",
                        Fall-Through = No
```


© 2023 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

Customer document and product support

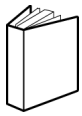

Customer documentation

Customer documentation welcome page

Technical support

Product support portal

Documentation feedback

Customer documentation feedback