



7450 Ethernet Service Switch
7750 Service Router
7950 Extensible Routing System
Virtualized Service Router
Releases up to 26.3.R3

Services Overview Advanced Configuration Guide for Classic CLI

3HE 20806 AAAG TQZZA
Edition: 01
July 2026

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.

© 2026 Nokia.

Table of contents

List of tables..... 4

List of figures.....5

Preface..... 7

BGP Selective Label-IPv4 Route Installation..... 8

G.8032 Ethernet Ring Protection Multiple Ring Topology..... 23

G.8032 Ethernet Ring Protection Single Ring Topology..... 60

GRE Tunnel Origination and Termination Using Non-system IP Addresses.....78

Inter-AS Option B Label Security for IP-VPN and EVPN Routes.....93

Network Group Encryption Helper..... 110

Seamless BFD Application — Auto-bind tunnel..... 142

List of tables

Table 1: Selective BGP-LU installation logic by service type..... 9

Table 2: Terminology comparison.....25

Table 3: Untrusted interfaces with default-forwarding forward option allow all IP-VPN and EVPN routes....94

Table 4: BGP neighbor-trust defines what traffic is allowed on untrusted interfaces with default-forwarding
drop option.....94

List of figures

Figure 1: Example topology.....	10
Figure 2: VPRN 1 uses a BGP transport tunnel with endpoint 192.0.1.21 on PE-2.....	13
Figure 3: VPRN 2, VPLS 3, and Epipe 4 use user-provisioned SDP 1 with BGP tunnel.....	17
Figure 4: PE-1 receives BGP-VPLS and BGP-AD routes with next-hop 192.0.1.23.....	19
Figure 5: G.8032 major ring and subring.....	26
Figure 6: G.8032 ring components.....	27
Figure 7: G.8032 subring interconnection components.....	28
Figure 8: Ethernet example topology.....	31
Figure 9: ETH-CFM MEP associations.....	33
Figure 10: Subring to VPLS topology.....	53
Figure 11: G.8032 operation and topologies.....	62
Figure 12: Example topology.....	63
Figure 13: Ethernet CFM configuration.....	67
Figure 14: Example topology.....	81
Figure 15: Mismatched T-LDP transport addresses.....	83
Figure 16: Matching T-LDP transport addresses.....	84
Figure 17: L2 services on PE-1 and PE-2.....	86
Figure 18: L3 services on PE-1 and PE-2.....	90
Figure 19: Example topology with services on PEs.....	95
Figure 20: Example topology with services on PEs and on ASBR-2.....	107
Figure 21: General architecture using an NGE helper.....	111

Figure 22: BGP topology for learning BGP label routes..... 114

Figure 23: Operation of NGE helper for MP-BGP auto-bind VPRN or NG-MVPN multicast..... 117

Figure 24: NGE helper for T-LDP signaled Epipe or VPLS services..... 120

Figure 25: NGE helper for BGP VPLS or BGP VPWS using GRE SDPs with auto-GRE SDP..... 123

Figure 26: S-BFD session establishment – continuity check..... 143

Figure 27: Example topology..... 144

Figure 28: Primary path of SR-TE LSP via PE-4..... 150

Figure 29: Remote failure in the primary path of the SR-TE LSP..... 151

Figure 30: SR-TE LSP reconnects after retry timer expires..... 153

Preface

About This Guide

Each Advanced Configuration Guide is organized alphabetically and provides feature and configuration explanations, CLI descriptions, and overall solutions. The Advanced Configuration Guide chapters are written for and based on several Releases, up to 26.3.R3. The Applicability section in each chapter specifies on which release the configuration is based.

The Advanced Configuration Guides supplement 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.

BGP Selective Label-IPv4 Route Installation

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

Topics in this chapter include:

- [Applicability](#)
- [Overview](#)
- [Configuration](#)
- [Conclusion](#)

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:PE-1# tree flat detail | match selective-label-ipv4-install
configure router bgp group neighbor selective-label-ipv4-install
configure router bgp group neighbor no selective-label-ipv4-install
configure router bgp group no selective-label-ipv4-install
configure router bgp group selective-label-ipv4-install
configure router bgp no selective-label-ipv4-install
configure router bgp selective-label-ipv4-install
```

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:

- 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.
VPN 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.
VPN with auto-bind-tunnel	If the auto-bind VPN 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.

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: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
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
  router Base
    policy-options
      begin
        prefix-list "192.0.1.0/24"
        prefix 192.0.1.0/24 prefix-length-range 32-32
      exit
    policy-statement "export-svc-lu-bgp"
      entry 10
        from
          prefix-list "192.0.1.0/24"
        exit
        action accept
      exit
    exit
  commit
exit
bgp
  split-horizon
  next-hop-resolution
    labeled-routes
      transport-tunnel
        family label-ipv4
          resolution-filter
            no ldp
            sr-isis
          exit
          resolution filter
        exit
      exit
    exit
  exit
exit
group "iBGPv4"
  family vpn-ipv4 label-ipv4
  peer-as 64500
  neighbor 192.0.2.1
    export "export-svc-lu-bgp"
  exit
exit
exit
```

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

```
*A: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                          524285
u*>i  192.0.1.22/32                        100        None
      192.0.2.2                          None       10
      No As-Path                          524285
u*>i  192.0.1.23/32                        100        None
      192.0.2.2                          None       10
      No As-Path                          524285
u*>i  192.0.1.24/32                        100        None
      192.0.2.2                          None       10
      No As-Path                          524285
-----
Routes : 4
=====

```

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

```

*A: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
=====

```

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: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 00h02m54s 170
      192.0.2.2 (tunneled:SR-ISIS:524290) 10
192.0.1.22/32           Remote BGP_LABEL 00h02m54s 170
      192.0.2.2 (tunneled:SR-ISIS:524290) 10
192.0.1.23/32           Remote BGP_LABEL 00h02m54s 170
      192.0.2.2 (tunneled:SR-ISIS:524290) 10
192.0.1.24/32           Remote BGP_LABEL 00h02m54s 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: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 .

VPNR 1 with auto-bind-tunnel

VPNR 1 in [Figure 2: VPNR 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.

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



35966

VPNR 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-VPNR1" sets the next-hop to 192.0.1.21 and adds the community "target:64500:1", which matches the vrf-target of VPNR 1.

```
# on PE-2:
configure
router Base
policy-options
begin
community "target:64500:1"
members "target:64500:1"
```

```

    exit
    policy-statement "export-VPN1"
    entry 10
    action accept
    next-hop 192.0.1.21
    community add "target:64500:1"
    exit
  exit
exit
commit
exit
exit
service
  vprn 1 name "VPRN 1" customer 1 create
  interface "lo1" create
  address 172.31.1.2/32
  loopback
  exit
  bgp-ipvpn
  mpls
  auto-bind-tunnel
  resolution-filter
  exit
  resolution filter
  exit
  route-distinguisher 64500:1
  vrf-export "export-VPN1"
  vrf-target target:64500:1
  no shutdown
  exit
exit
no shutdown

```

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: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                               524287
-----
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
  exit
```

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: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                             524285
i     192.0.1.22/32                          100         None
      192.0.2.2                             None        10
      No As-Path                             524285
i     192.0.1.23/32                          100         None
      192.0.2.2                             None        10
      No As-Path                             524285
i     192.0.1.24/32                          100         None
      192.0.2.2                             None        10
      No As-Path                             524285
-----
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: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: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 00h04m01s 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
=====
```

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 "lo2"
    exit
  exit
  no shutdown
exit
exit
service
  sdp 1 mpls create
    signaling tldp      # default
    far-end 192.0.1.12
    bgp-tunnel
    no shutdown
  exit
exit
```

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

```
*A: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: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             524285
u*>i  192.0.1.22/32          100        None
      192.0.2.2             None        10
      No As-Path             524285
i     192.0.1.23/32          100        None
      192.0.2.2             None        10
      No As-Path             524285
i     192.0.1.24/32          100        None
      192.0.2.2             None        10
      No As-Path             524285
-----
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
```

```
router Base
  policy-options
    begin
      community "target:64500:2"
        members "target:64500:2"
      exit
    policy-statement "export-VRPN2"
      entry 10
        action accept
        next-hop 192.0.1.12
        community add "target:64500:2"
      exit
    exit
  exit
  commit
exit
exit
service
  vprn 2 name "VRPN 2" customer 1 create
    interface "lo1" create
      address 172.31.2.1/32
      loopback
    exit
    bgp-ipvpn
      mpls
        route-distinguisher 64500:2
        vrf-export "export-VRPN2"
        vrf-target target:64500:2
        no shutdown
      exit
    exit
    spoke-sdp 1:2 create
    exit
    no shutdown
  exit
exit
```


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

```
# on PE-1:
configure
  service
    vpls 3 name "VPLS 3" customer 1 create
      sap 1/1/c2/1:3 create
      exit
      spoke-sdp 1:3 create
      exit
      no shutdown
    exit
    epipe 4 name "Epipe 4" customer 1 create
      sap 1/1/c2/1:4 create
      exit
      spoke-sdp 1:4 create
      exit
      no shutdown
    exit
```

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"
      family l2-vpn
      type internal
      local-address 192.0.1.23
      neighbor 192.0.1.13
    exit
  exit
exit
```

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 "lo3"
    exit
  exit
  no shutdown
exit
```

The service configuration is as follows:

```
# on PE-2:
configure
service
  pw-template 1 name "PW1" create
```

```

exit
vpls 5 name "BGP-VPLS 5" customer 1 create
  bgp
    route-distinguisher 64500:5
    route-target export target:64500:5 import target:64500:5
    pw-template-binding 1 import-rt "target:64500:5"
  exit
  exit
  bgp-vpls
    max-ve-id 100
    ve-name "PE-2"
    ve-id 2
  exit
  no shutdown
exit
sap 1/1/c2/1:5 create
exit
no shutdown
exit
vpls 6 name "BGP-AD VPLS 6" customer 1 create
  bgp
    route-distinguisher 64500:6
    route-target export target:64500:6 import target:64500:6
    pw-template-binding 1
  exit
  exit
  bgp-ad
    vpls-id 64500:6
    vsi-id
    prefix 192.0.1.23
  exit
  no shutdown
exit
sap 1/1/c2/1:6 create
exit
no shutdown
exit

```

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

```

*A: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
NextHop VeId BlockSize LocalPref
As-Path BaseOffset vplsLabelBase
-----
u*>i VPLS - - 0
64500:5 - - -
192.0.1.23 2 8 100
No As-Path 1 524273
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:PE-1# show service id 5 sdp
```

```

=====
Services: Service Destination Points
=====
SdpId           Type      Far End addr  Adm   Opr      I.Lbl  E.Lbl
-----
32767:4294967295 BgpVpls  192.0.1.23   Up    Up        524274 524273
-----
Number of SDPs : 1
=====

```

```
*A:PE-1# show service id 6 sdp
```

```

=====
Services: Service Destination Points
=====
SdpId           Type      Far End addr  Adm   Opr      I.Lbl  E.Lbl
-----
32766:4294967294 BgpAd    192.0.1.23   Up    Up        524268 524268
-----
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: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             524285
u*>i  192.0.1.22/32             100        None
      192.0.2.2             None        10
      No As-Path             524285

```

```
u*>i 192.0.1.23/32      100      None
      192.0.2.2        None      10
      No As-Path       524285
i    192.0.1.24/32      100      None
      192.0.2.2        None      10
      No As-Path       524285
-----
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:PE-1# show router bgp routes label-ipv4 hunt | match "Invalid" pre-lines 16

Network      : 192.0.1.24/32
Nextthop     : 192.0.2.2
Path Id      : None
From         : 192.0.2.2
Res. Nextthop : 192.0.2.2 (ISIS Tunnel)
Local Pref.  : 100
Aggregator AS : None
Atomic Aggr. : Not Atomic
AIGP Metric  : None
Connector    : None
Community    : No Community Members
Cluster      : No Cluster Members
Originator Id : None
Fwd Class    : None
IPv4 Label   : 524285
Flags        : Invalid IGP Label-Unicast-No-Svc

Interface Name : NotAvailable
Aggregator     : None
MED            : None
IGP Cost       : 10
Peer Router Id : 192.0.2.2
Priority       : None
```

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.

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](#)
- [Overview](#)
- [Configuration](#)
- [Conclusion](#)

Applicability

Initially, this chapter was written for SR OS Release 12.0.R5, but the 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.

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 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.

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.

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.

- 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



26167

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 **new-qinq-untagged-sap** parameter (**configure system ethernet new-qinq-untagged-sap**), 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.

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 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 <down-timeout>] [up <up-timeout>] }
  compatible-version <version> # [1..2] - Default: 2
  description <description-string>
  guard-time <time> # [1..20] in deciseconds - Default: 5
  node-id <xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx>
  path {a|b} [ { <port-id>|<lag-id> } raps-tag <qtag1>[.<qtag2>] ]
    description <description-string>
    eth-cfm
      mep <mep-id> domain <md-index> association <ma-index>
      <...>
    rpl-end
    shutdown
  revert-time <time> # [60..720] in seconds - Default: 300
  rpl-node {owner|nbr}
  shutdown
  sub-ring {virtual-link|non-virtual-link}
    interconnect { ring-id <ring-index> | vpls }
    propagate-topology-change
```

Parameters:

- **<ring-index>** — The ring index is the number by which the ring is referenced; values: 1 to 128.
- **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 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:

```
<down-timeout> : [0..5000] in centiseconds - default: 0; 1 centisecond = 10 ms
<up-timeout> : [0..5000] in deciseconds - default: 20; 1 decisecond = 100 ms
```

- 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 that can be used 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:

```
[1..20] in deciseconds - default: 5; 1 decisecond = 100ms
```


- The **node-id** (*<xx:xx:xx:xx:xx:xx>* or *<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.
- **path** {*a* | *b*} [*{<port-id> | <lag-id>}*] **raps-tag** *<qtag1>*[*<qtag2>*] — 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 **description** includes a text string of maximum 80 characters to describe the use of the path.
 - The **eth-cfm** context contains the associated Ethernet CFM parameters.
 - **mep** *<mep-id>* **domain** *<md-index>* **association** *<ma-index>* — 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.
 - When the **rpl-end** parameter is configured, the path is expected to be one end of the RPL. The **rpl-end** parameter must be configured in conjunction with the **rpl-node** parameter.
 - The **shutdown** command disables the path.
- 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. Configuring **no revert-time** disables reversion, effectively setting the revert time to zero. Values: [60..720] in seconds - Default: 300.
- **rpl-node** {*owner* | *nbr*} — A node can be designated as either the **owner** of the RPL, in which case this node is responsible for the RPL, or the **nbr** (neighbor), in which case this node is expected to be

the neighbor to the RPL owner across the RPL. The **nbr** is optional and is included to be compliant with the specification. This parameter must be configured in conjunction with the **rpl-end** command. On a subring without virtual channel it is mandatory to configure **sub-ring non-virtual-link** on all nodes on the subring to ensure propagation of the R-APS messages around the subring.

- **shutdown** — This command disables the ring.
- **sub-ring {virtual-link | non-virtual-link}** — This command is configured on the interconnection node between the subring and its major ring or subring to indicate that this ring is a subring. The parameter specifies whether it uses a virtual link through the major ring or subring for the R-APS messages or not. A ring configured as a subring can only be configured with a single path.
 - **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 **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](#).

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
- ETH-CFM
- Ethernet ring for major ring 1
- Ethernet ring for subring 2
- control channel service and Ethernet ring SAPs
- user data channels

Configure the encapsulation for the ring ports.

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 **new-qinq-untagged-sap** parameter (**configure system ethernet new-qinq-untagged-sap**), 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
    ethernet
      mode access
      encap-type qinq
    exit
    no shutdown
  exit
  port 1/1/c2/1
    ethernet
      mode access
      encap-type qinq
    exit
    no shutdown
  exit
  port 1/1/c3/1
    ethernet
      mode access
      encap-type qinq
    exit
```

```
no shutdown
exit
```

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-enable** from under the path MEPs and to remove the **remote-mepid** on the corresponding ETH CFM configuration.

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

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



al_0533

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 12 and 13 are used for the major ring and association 14 is used for the subring.

```
# on PE-1:
configure
eth-cfm
  domain 1 format none level 2 admin-name "domain-1"
    association 12 format icc-based name "Association12" admin-name "association-12"
      ccm-interval 1
      remote-mepid 122
    exit
    association 13 format icc-based name "Association13" admin-name "association-13"
      ccm-interval 1
      remote-mepid 133
    exit
    association 14 format icc-based name "Association14" admin-name "association-14"
      ccm-interval 1
      remote-mepid 144
    exit
  exit
exit
```

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

```
# on PE-2:
configure
eth-cfm
  domain 1 format none level 2 admin-name "domain-1"
    association 12 format icc-based name "Association12" admin-name "association-12"
      ccm-interval 1
      remote-mepid 121
    exit
    association 23 format icc-based name "Association23" admin-name "association-23"
      ccm-interval 1
      remote-mepid 233
    exit
  exit
exit
```

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

```
# on PE-3:
configure
eth-cfm
  domain 1 format none level 2 admin-name "domain-1"
    association 13 format icc-based name "Association13" admin-name "association-13"
      ccm-interval 1
      remote-mepid 131
    exit
    association 23 format icc-based name "Association23" admin-name "association-23"
      ccm-interval 1
      remote-mepid 232
    exit
    association 34 format icc-based name "Association34" admin-name "association-34"
      ccm-interval 1
      remote-mepid 344
    exit
  exit
exit
```

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

```
# on PE-4
configure
  eth-cfm
    domain 1 format none level 2 admin-name "domain-1"
      association 14 format icc-based name "Association14" admin-name "association-14"
        ccm-interval 1
        remote-mepid 141
      exit
    association 34 format icc-based name "Association34" admin-name "association-34"
      ccm-interval 1
      remote-mepid 343
    exit
  exit
```

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 its minimum value of 60 seconds and CCM messages are enabled on the MEP. The **control-mep** parameter is required to indicate 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
    description "Ethernet ring 1"
    revert-time 60
    path a 1/1/c1/1 raps-tag 1.1
      description "Ethernet ring 1 - pathA"
      eth-cfm
        mep 121 domain 1 association 12
        ccm-enable
        control-mep
        no shutdown
      exit
    exit
    no shutdown
  exit
  path b 1/1/c3/1 raps-tag 1.1
    description "Ethernet Ring 1 - PathB"
    eth-cfm
      mep 131 domain 1 association 13
      ccm-enable
      control-mep
      no shutdown
    exit
  exit
  no shutdown
exit
no shutdown
exit
```

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

```
*A:PE-1>config>eth-ring# path a 1/1/c1/1 raps-tag 1.1
```

```
*A:PE-1>config>eth-ring>path# no shutdown
INFO: ERMGR #1001 Not permitted - must configure eth-cfm MEP first
```

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
    description "Ethernet Ring 1"
    revert-time 60
    rpl-node owner
    path a 1/1/c1/1 raps-tag 1.1
      description "Ethernet ring 1 - PathA"
      rpl-end
      eth-cfm
        mep 232 domain 1 association 23
        ccm-enable
        control-mep
        no shutdown
      exit
    exit
    no shutdown
  exit
  path b 1/1/c2/1 raps-tag 1.1
    description "Ethernet ring 1 - PathB"
    eth-cfm
      mep 122 domain 1 association 12
      ccm-enable
      control-mep
      no shutdown
    exit
  exit
  no shutdown
exit
no shutdown
exit
```

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 **nbr** otherwise the following error message is reported.

```
*A:PE-2>config>eth-ring>path# rpl-end
INFO: ERMGR #1001 Not permitted - path-type rpl-end is not consistent with eth-ring
'rpl-node' type
```

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
    description "Ethernet ring 1"
    revert-time 60
    rpl-node nbr
    path a 1/1/c3/1 raps-tag 1.1
      description "Ethernet ring 1 - PathA"
      eth-cfm
        mep 133 domain 1 association 13
```

```

        ccm-enable
        control-mep
        no shutdown
    exit
    exit
    no shutdown
exit
path b 1/1/c2/1 raps-tag 1.1
    description "Ethernet ring 1 - PathB"
    rpl-end
    eth-cfm
        mep 233 domain 1 association 23
        ccm-enable
        control-mep
        no shutdown
    exit
    exit
    no shutdown
exit
    no shutdown
exit

```

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 (ring 1). The subring in this case uses a virtual link. The interconnection ring instance identifier (**ring-id**) is specified and **propagate-topology-change** 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 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
        description "Ethernet subring 2 on major ring 1"
        revert-time 60
        sub-ring virtual-link
            interconnect ring-id 1
            propagate-topology-change
        exit
    exit
    path a 1/1/c2/1 raps-tag 2.1
        description "Ethernet ring 2 - PathA"
        eth-cfm
            mep 141 domain 1 association 14
            ccm-enable
            control-mep
            no shutdown
        exit

```

```

        exit
        no shutdown
    exit
    no shutdown
exit

```

The configuration of 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
    description "Ethernet subring 2 on major ring 1"
    revert-time 60
    rpl-node nbr
    sub-ring virtual-link
      interconnect ring-id 1
      propagate-topology-change
    exit
  exit
  path a 1/1/c1/1 raps-tag 2.1
    description "Ethernet ring 2 - PathA"
    rpl-end
    eth-cfm
      mep 343 domain 1 association 34
      ccm-enable
      control-mep
      no shutdown
    exit
  exit
  no shutdown
exit
no shutdown
exit

```

Ring node PE-4 only has configuration for the subring 2. 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
    description "Ethernet subring 2"
    revert-time 60
    rpl-node owner
    path a 1/1/c1/1 raps-tag 2.1
      description "Ethernet ring 2 - PathA"
      eth-cfm
        mep 144 domain 1 association 14
        ccm-enable
        control-mep
        no shutdown
      exit
    exit
    no shutdown
  exit
  path b 1/1/c2/1 raps-tag 2.1
    description "Ethernet ring 2 - PathB"
    rpl-end
    eth-cfm
      mep 344 domain 1 association 34
      ccm-enable
      control-mep

```

```

        no shutdown
    exit
    exit
    no shutdown
exit
no shutdown
exit

```

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:PE-1# show eth-ring 1

=====
Ethernet Ring 1 Information
=====
Description       : Ethernet ring 1
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
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 1 name "control-VPLS-1" customer 1 create
  description "Control VID 1.1 for ring 1 - major ring"
  sap 1/1/c1/1:1.1 eth-ring 1 create
  exit
  sap 1/1/c3/1:1.1 eth-ring 1 create
  exit
  no shutdown
exit

```

```

# on PE-2:
configure
service
  vpls 1 name "control-VPLS-1" customer 1 create
  description "Control VID 1.1 for ring 1 - major ring"
  sap 1/1/c1/1:1.1 eth-ring 1 create
  exit
  sap 1/1/c2/1:1.1 eth-ring 1 create
  exit
  no shutdown
exit

```

```

# on PE-3:
configure
service
  vpls 1 name "control-VPLS-1" customer 1 create
  description "Control VID 1.1 for ring 1 - major ring"
  sap 1/1/c2/1:1.1 eth-ring 1 create
  exit
  sap 1/1/c3/1:1.1 eth-ring 1 create
  exit
  no shutdown
exit

```

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 this parameter to a control channel VPLS results in the following message being displayed.

```

*A:PE-1>config>service>vpls# sap 1/1/c4/1:1 create
MINOR: SVCMgr #1321 Service contains an Ethernet ring control SAP

```

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 2 name "control-VPLS-2" customer 1 create
  description "control/virtual channel VID 2.1 for ring 2"
  split-horizon-group "shg-ring2" create
  exit
  sap 1/1/c1/1:2.1 split-horizon-group "shg-ring2" eth-ring 1 create
  description "ring 2 interconnection using ring 1"
  exit
  sap 1/1/c2/1:2.1 eth-ring 2 create

```

```

exit
sap 1/1/c3/1:2.1 split-horizon-group "shg-ring2" eth-ring 1 create
description "ring 2 interconnection using ring 1"
exit
no shutdown
exit

```

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 2 name "control-VPLS-2" customer 1 create
description "virtual channel VID 2.1 for ring 2"
sap 1/1/c1/1:2.1 eth-ring 1 create
exit
sap 1/1/c2/1:2.1 eth-ring 1 create
exit
no shutdown
exit

```

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 2 name "control-VPLS-2" customer 1 create
description "control/virtual channel VID 2.1 for ring 2"
split-horizon-group "shg-ring2" create
exit
sap 1/1/c1/1:2.1 eth-ring 2 create
exit
sap 1/1/c2/1:2.1 split-horizon-group "shg-ring2" eth-ring 1 create
description "ring 2 interconnection using ring 1"
exit
sap 1/1/c3/1:2.1 split-horizon-group "shg-ring2" eth-ring 1 create
description "ring 2 interconnection using ring 1"
exit
no shutdown
exit

```

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 2 name "control-VPLS-2" customer 1 create
description "Control VID 2.1 for ring 2 Sub-ring"

```

```
sap 1/1/c1/1:2.1 eth-ring 2 create
exit
sap 1/1/c2/1:2.1 eth-ring 2 create
exit
no shutdown
exit
```

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:PE-1# show eth-ring status
```

```
=====
Ethernet Ring (Status information)
=====
```

Ring ID	Admin State	Oper State	Path Information		State	MEP Information		
			Path	Tag		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:PE-1# show eth-ring
```

```
=====
Ethernet Rings (summary)
=====
```

Ring ID	Int ID	Admin State	Oper State	Paths Summary		Path States			
				a	b	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:PE-1# show eth-ring 1
```

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

Description	: Ethernet 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   : 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:PE-2# show eth-ring 1
```

Ethernet Ring 1 Information

```
Description      : Ethernet Ring 1
Admin State      : Up
Node ID          : 02:0b:ff:00:00:00
Guard Time       : 5 deciseconds
Max Revert Time  : 60 seconds
CCM Hold Down Time : 0 centiseconds
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
```

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:PE-3# show eth-ring 1
```

Ethernet Ring 1 Information

```

=====
Description      : Ethernet 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     : 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: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     : 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: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:PE-4# show eth-ring 2

=====
Ethernet Ring 2 Information
=====
Description      : Ethernet subring 2
Admin State      : Up                               Oper State      : Up
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:PE-1# show eth-ring 1 path a

=====
Ethernet Ring 1 Path Information
=====
Description      : Ethernet ring 1 - pathA
Port             : 1/1/c1/1          Raps-Tag       : 1.1
Admin State      : Up                Oper State      : Up
Path Type        : normal            Fwd State       : unblocked
                                      Fwd State Change : 05/10/2023 07:35:33
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: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
```

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 11 name "VPLS-11" customer 1 create
        description "data VPLS"
        sap 1/1/c1/1:1.11 eth-ring 1 create
        exit
        sap 1/1/c2/1:1.11 eth-ring 2 create
        exit
        sap 1/1/c3/1:1.11 eth-ring 1 create
        exit
        sap 1/1/c4/1:11 create
        description "sample customer service SAP"
```

```

        exit
    no shutdown
exit

```

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 11 name "VPLS-11" customer 1 create
        description "data VPLS"
        sap 1/1/c1/1:1.11 eth-ring 1 create
        exit
        sap 1/1/c2/1:1.11 eth-ring 1 create
        exit
        sap 1/1/c4/1:1.11 create
            description "sample customer service SAP"
        exit
    no shutdown
exit

```

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 11 name "VPLS-11" customer 1 create
        description "data VPLS"
        sap 1/1/c1/1:1.11 eth-ring 2 create
        exit
        sap 1/1/c2/1:1.11 eth-ring 2 create
        exit
        sap 1/1/c4/1:1.11 create
            description "sample customer service SAP"
        exit
    no shutdown
exit

```

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: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:PE-1# show eth-cfm statistics

=====
ETH-CFM System Statistics
=====
Rx Count      : 3458          Tx Count      : 3168
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           3008      3099
lbr            0         0
lbrm          0         0
ltr            0         0
ltm            0         0
ais            0         0
lck            0         0
tst            0         0
laps          0         0
raps          450        69
mcc            0         0
lmr            0         0
lmm            0         0
ldm            0         0
dmr            0         0
dmm            0         0
exr            0         0
exm            0         0
csf            0         0
vsr            0         0
vsm            0         0
lsl            0         0
slr            0         0
slm            0         0
gnm            0         0
other          0         0
-----
Total          3458      3168
=====
```

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
    shutdown
```

```
84 2023/05/10 07:54:04.850 UTC MINOR: ETH_CFM #2001 Base
"MEP 1/12/122 highest defect is now defRemoteCCM"

83 2023/05/10 07:54:01.310 UTC MAJOR: SVCMMGR #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."

82 2023/05/10 07:54:01.301 UTC MINOR: ERING #2001 Base eth-ring-1
"Eth-Ring 1 path a changed fwd state to unblocked"

81 2023/05/10 07:54:01.301 UTC MINOR: ERING #2001 Base eth-ring-1
"Eth-Ring 1 path b changed fwd state to blocked"

80 2023/05/10 07:54:01.300 UTC 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: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 2/2 tm 000 01:56:04.290/000 02:01:31.070
    state: Cnt 5 B/F 000 02:22:01.000/000 02:19:55.750, flag: 0x0
  path-b, port 1/1/c2/1 (Up), tag 1.1(Up) status (Up/Up/Fwd)
    cc (Dn/Up): Cnt 3/3 tm 000 02:19:59.300/000 02:20:44.520
    state: Cnt 6 B/F 000 02:19:55.750/000 02:22:01.000, 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= v1-0x0000 Nr, cmd= None
path-b Normal, RxId(I)= 02:09:ff:00:00:00, rx= v1-0x0000 Nr, cmd= None
DebugInfo: aPathSts 3, bPathSts 5, pm (set/clr) 0/0, txFlush 0
RxRaps: ok 20 nok 0 self 30, TmrExp - wtr 2(1), grd 3, wtb 0
Flush: cnt 8 (5/3/0) tm 000 02:22:01.000-000 02:22:01.000 Out/Ack 0/1
RxRawRaps: aPath 49 bPath 43 vPath 0
Now: 000 02:24:13.310 , softReset: No - noTx 0

Seq Event  RxInfo(Path: NodeId-Bytes)
          state:TxInfo (Bytes)          Dir  pA  pB          Time
=== =====
013  pdu A: 02:0d:ff:00:00:00-0xb060 Sf(DNF)
      PEND-G: 0x0000 Nr                  Rx<-- Blk Fwd 000 02:01:33.630
014  pdu B: 02:0d:ff:00:00:00-0xb060 Sf(DNF)
      PEND-G: 0x0000 Nr                  Rx<-- Blk Fwd 000 02:01:33.630
```

```

015 pdu A: 02:0d:ff:00:00:00-0xb060 Sf(DNF)
      PEND-G: 0x0000 Nr Rx<-- Blk Fwd 000 02:01:33.730
016 pdu B: 02:0d:ff:00:00:00-0xb060 Sf(DNF)
      PEND-G: 0x0000 Nr Rx<-- Blk Fwd 000 02:01:33.730
017 pdu A: 02:0d:ff:00:00:00-0x0020 Nr
      PEND-G: 0x0000 Nr Rx<-- Blk Fwd 000 02:01:33.830
018 pdu B: 02:0d:ff:00:00:00-0x0020 Nr
      PEND-G: 0x0000 Nr Rx<-- Blk Fwd 000 02:01:33.830
019 pdu A: 02:0d:ff:00:00:00-0x0020 Nr
      PEND-G: 0x0000 Nr Rx<-- Blk Fwd 000 02:01:33.930
000 pdu B: 02:0d:ff:00:00:00-0x0020 Nr
      PEND-G: 0x0000 Nr Rx<-- Blk Fwd 000 02:01:33.930
001 pdu A: 02:0d:ff:00:00:00-0x0020 Nr
      PEND-G: 0x0000 Nr Rx<-- Blk Fwd 000 02:01:34.030
002 pdu B: 02:0d:ff:00:00:00-0x0020 Nr
      PEND-G: 0x0000 Nr Rx<-- Blk Fwd 000 02:01:34.030
003 pdu A: 02:0d:ff:00:00:00-0x0020 Nr
      PEND : 0x0000 Nr Rx<-- Blk Fwd 000 02:01:38.030
004 pdu
      PEND : ----- Fwd Fwd 000 02:01:38.030
005 pdu B: 02:0d:ff:00:00:00-0x0020 Nr
      PEND : Rx<-- Fwd Fwd 000 02:01:38.030
006 xWtr
      IDLE : 0x0080 Nr(RB ) TxF-> Blk Fwd 000 02:02:38.000
007 bDn
      PROT : 0xb020 Sf TxF-> Fwd Blk 000 02:19:55.750
008 pdu A: 02:09:ff:00:00:00-0xb000 Sf
      PROT : 0xb020 Sf RxF<- Fwd Blk 000 02:19:59.520
009 bUp
      PEND-G: 0x0020 Nr Tx-> Fwd Blk 000 02:20:46.500
010 pdu B: 02:09:ff:00:00:00-0x0000 Nr
      PEND : 0x0020 Nr Rx<-- Fwd Blk 000 02:20:47.360
011 pdu A: 02:09:ff:00:00:00-0x0000 Nr
      PEND : 0x0020 Nr Rx<-- Fwd Blk 000 02:20:47.360
012 xWtr
      IDLE : 0x0080 Nr(RB ) TxF-> Blk Fwd 000 02:22:01.000

```

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 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
    description "Ethernet subring 2 on major ring 1"

```

```

    revert-time 60
    sub-ring non-virtual-link
        interconnect ring-id 1
        propagate-topology-change
    exit
exit
path a 1/1/c2/1 raps-tag 2.1
    description "Ethernet ring 2 - PathA"
    eth-cfm
        mep 141 domain 1 association 14
        ccm-enable
        control-mep
        no shutdown
    exit
    exit
    no shutdown
exit
    no shutdown
exit

```

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

```

# on PE-4:
configure
    eth-ring 2
        description "Ethernet subring 2"
        revert-time 60
        rpl-node owner
        sub-ring non-virtual-link
    exit
    path a 1/1/c1/1 raps-tag 2.1
        description "Ethernet ring 2 - PathA"
        eth-cfm
            mep 144 domain 1 association 14
            ccm-enable
            control-mep
            no shutdown
        exit
    exit
    no shutdown
exit
    path b 1/1/c2/1 raps-tag 2.1
        description "Ethernet ring 2 - PathB"
        rpl-end
        eth-cfm
            mep 344 domain 1 association 34
            ccm-enable
            control-mep
            no shutdown
        exit
    exit
    no shutdown
exit
    no shutdown
exit

```

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

```

*A:PE-1# show service sap-using eth-ring

```

```

=====

```

Service Access Points (Ethernet Ring)

SapId	SvcId	Eth-Ring	Path	Admin State	Oper State	Blocked	Control/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							

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:PE-1# show eth-ring 2
```

Ethernet Ring 2 Information

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

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

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 messages (flush-all-from-me) 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.

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
    description "Ethernet subring 2 on VPLS"
    revert-time 60
```

```

    sub-ring non-virtual-link
    interconnect vpls
    propagate-topology-change
    exit
exit
path a 1/1/c2/1 raps-tag 2.1
  description "Ethernet ring 2 - PathA"
  eth-cfm
    mep 141 domain 1 association 14
    ccm-enable
    control-mep
    no shutdown
  exit
exit
no shutdown
exit
no shutdown
exit

```

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

```

# on PE-4:
configure
  eth-ring 2
    description "Ethernet subring 2"
    revert-time 60
    rpl-node owner
    sub-ring non-virtual-link
    exit
    path a 1/1/c1/1 raps-tag 2.1
      description "Ethernet ring 2 - PathA"
      eth-cfm
        mep 144 domain 1 association 14
        ccm-enable
        control-mep
        no shutdown
      exit
    exit
    no shutdown
  exit
  path b 1/1/c2/1 raps-tag 2.1
    description "Ethernet ring 2 - PathB"
    rpl-end
    eth-cfm
      mep 344 domain 1 association 34
      ccm-enable
      control-mep
      no shutdown
    exit
  exit
  no shutdown
exit
no shutdown
exit

```

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

```

# on PE-1:
configure
  service
    vpls 11 name "VPLS-11" customer 1 create

```

```

description "data VPLS"
sap 1/1/c2/1:1.11 eth-ring 2 create
no shutdown
exit
sap 1/1/c4/1:1.11 create
description "sample customer service SAP"
no shutdown
exit
mesh-sdp 12:11 create
no shutdown
exit
mesh-sdp 13:11 create
no shutdown
exit
no shutdown
exit

```

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:PE-1# show eth-ring 2
```

```
=====
Ethernet Ring 2 Information
=====
```

```

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

```

```

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

```

```
-----
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, as follows:

```

*A:PE-1# debug router ldp peer 192.0.2.2 packet init
*A:PE-1# debug router ldp peer 192.0.2.3 packet init

```

```

# on PE-1:
debug
router "Base"
  ldp
    peer 192.0.2.2
      event

```

```

        exit
        packet
            init
        exit
    exit
peer 192.0.2.3
    event
    exit
    packet
        init
    exit
exit
exit
exit
exit
exit

```

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

```

# on PE-1:
configure
    port 1/1/c2/1
        shutdown

```

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 flush messages.

```

2 2023/05/10 09:37:40.672 UTC WARNING: SNMP #2004 Base 1/1/c2/1
"Interface 1/1/c2/1 is not operational"

3 2023/05/10 09:37:40.672 UTC MINOR: ERING #2001 Base eth-ring-2
"Eth-Ring 2 path a changed fwd state to blocked"

1 2023/05/10 09:37:40.673 UTC MINOR: DEBUG #2001 Base LDP
"LDP: LDP
Send Address Withdraw packet (msgId 10173) 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 2023/05/10 09:37:40.673 UTC MINOR: DEBUG #2001 Base LDP
"LDP: LDP
Send Address Withdraw packet (msgId 10164) to 192.0.2.3:0
MAC Flush (All MACs learned from me)
Service FEC PWE3: ENET(5)/11 Group ID = 0 cBit = 0
"

4 2023/05/10 09:37:40.691 UTC MAJOR: SVCNMR #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."

3 2023/05/10 09:37:44.028 UTC MINOR: DEBUG #2001 Base LDP
"LDP: LDP
Recv Address Withdraw packet (msgId 10160) from 192.0.2.3:0
"

5 2023/05/10 09:37:44.081 UTC 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:PE-1# show eth-ring 1

=====
Ethernet Ring 1 Information
=====
Description      : Ethernet ring 1
Admin State      : Up
Node ID          : 02:09:ff:00:00:00
Guard Time       : 5 deciseconds
Max Revert Time  : 60 seconds
CCM Hold Down Time : 0 centiseconds
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 on PE-1 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:PE-1# show eth-ring 1

=====
Ethernet Ring 1 Information
=====
Description      : Ethernet ring 1
Admin State      : Up
Node ID          : 02:09:ff:00:00:00
Guard Time       : 5 deciseconds
Max Revert Time  : 60 seconds
CCM Hold Down Time : 0 centiseconds
RPL Node         : rplNone
Time to Revert   : N/A
CCM Hold Up Time : 20 deciseconds

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
=====
```

```
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

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

```
*A:PE-1# tools perform eth-ring clear 1
```

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

```
*A:PE-1# show eth-ring 1
```

Ethernet Ring 1 Information

```
Description      : Ethernet ring 1
Admin State      : Up
Node ID         : 02:09:ff:00:00:00
Guard Time      : 5 deciseconds
Max Revert Time  : 60 seconds
CCM Hold Down Time : 0 centiseconds
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: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.

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](#)
- [Overview](#)
- [Configuration](#)
- [Conclusion](#)

Applicability

The chapter was initially written for SR OS Release 8.0.R7, but the 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) 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: ETH 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.

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](#).

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 <down-timeout>] [up <up-timeout>] }
  compatible-version <version> # [1..2] - Default: 2
  description <description-string>
  guard-time <time> # [1..20] in deciseconds - Default: 5
  node-id <xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx>
  path {a|b} [ { <port-id>|<lag-id> } raps-tag <qtag1>[.<qtag2>] ]
  description <description-string>
  eth-cfm
    mep <mep-id> domain <md-index> association <ma-index>
    <...>
  rpl-end
  shutdown
  revert-time <time> # [60..720] in seconds - Default: 300
  rpl-node {owner|nbr}
  shutdown
  sub-ring {virtual-link|non-virtual-link} # beyond the scope
```

Parameters:

- **ring-index** — This is the number by which the ring is referenced, values: 1 to 128.
- **ccm-hold-time** {[down <down-timeout>] [up <up-timeout>]}
 - **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.
- timeout values:

```
<down-timeout>      : [0..5000] in 100ths of seconds - Default: 0
<up-timeout>        : [0..5000] in 10ths of seconds - Default: 20
```

- **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 <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 <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.

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

```
[1..20] in deciseconds - Default: 5
```

- **node-id <xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx>** — 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 {a|b} [{<port-id>|<lag-id>} raps-tag <qtag1>[.<qtag2>]]** — 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 <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 <mep-id> domain <md-index> association <ma-index>** — 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**.
- **shutdown** — This command disables the path.
- **revert-time <time>** — 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. Configuring **no revert-time** disables reversion, effectively setting the revert-time to zero.

Values:

[60..720] in seconds - Default: 300

- **rpl-node {owner|nbr}** — A node can be designated as either the **owner** of the RPL, in which case this node is responsible for the RPL, or the **nbr**, 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.
- **shutdown** — This command disables the ring.
- **sub-ring {virtual-link|non-virtual-link}** — The **sub-ring** command is beyond the scope of this chapter because it is only required for multiple ring topologies.

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 1 name "log1"
    from main
    to console
exit
exit
```

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 **new-qinq-untagged-sap** parameter (**configure system ethernet new-qinq-untagged-sap**), 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 this example, 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
    ethernet
      mode access
      encap-type dot1q
    exit
  no shutdown
exit
port 1/1/c2/1
  ethernet
    mode access
    encap-type dot1q
  exit
  no shutdown
exit
```

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-enable** from under the path MEPs and remove the **remote-mepid** 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 1 format none level 3 admin-name "domain-1"
    association 1 format icc-based name "ring1_1_2" admin-name "association-1"
      ccm-interval 1
      remote-mepid 1122
    exit
    association 2 format icc-based name "ring1_1_3" admin-name "association-2"
      ccm-interval 1
      remote-mepid 1133
    exit
  exit
exit
```

```
# on PE-2:
configure
eth-cfm
  domain 1 format none level 3 admin-name "domain-1"
    association 1 format icc-based name "ring1_2_3" admin-name "association-1"
      ccm-interval 1
      remote-mepid 1233
    exit
    association 2 format icc-based name "ring1_1_2" admin-name "association-2"
      ccm-interval 1
      remote-mepid 1121
    exit
  exit
exit
```

```
# on PE-3:
configure
eth-cfm
```

```
domain 1 format none level 3 admin-name "domain-1"
  association 1 format icc-based name "ring1_1_3" admin-name "association-1"
    ccm-interval 1
    remote-mepid 1131
  exit
  association 2 format icc-based name "ring1_2_3" admin-name "association-2"
    ccm-interval 1
    remote-mepid 1232
  exit
exit
```

Configure Ethernet ring

Two paths need to be configured to form a ring. 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
    path a 1/1/c1/1 raps-tag 1
      eth-cfm
        mep 1121 domain 1 association 1
        ccm-enable
        control-mep
        no shutdown
      exit
    exit
    no shutdown
  exit
  path b 1/1/c2/1 raps-tag 1
    eth-cfm
      mep 1131 domain 1 association 2
      ccm-enable
      control-mep
      no shutdown
    exit
  exit
  no shutdown
exit
no shutdown
exit
```

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

```
*A:PE-1>config>eth-ring# path a 1/1/c1/1 raps-tag 1
*A:PE-1>config>eth-ring>path# no shutdown
INFO: ERMGR #1001 Not permitted - must configure eth-cfm MEP first
```

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
    revert-time 60
```

```
rpl-node owner
path a 1/1/c1/1 raps-tag 1
  eth-cfm
    mep 1232 domain 1 association 1
      ccm-enable
      control-mep
      no shutdown
    exit
  exit
  no shutdown
exit
no shutdown
path b 1/1/c2/1 raps-tag 1
  rpl-end
  eth-cfm
    mep 1122 domain 1 association 2
      ccm-enable
      control-mep
      no shutdown
    exit
  exit
  no shutdown
exit
no shutdown
exit
```

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 **nbr** otherwise the following error message is reported.

```
*A:PE-2>config>eth-ring>path# rpl-end
INFO: ERMGR #1001 Not permitted - path-type rpl-end is not consistent with eth-ring 'rpl-node'
type
```

```
# on PE-3:
configure
  eth-ring 1
    path a 1/1/c1/1 raps-tag 1
      eth-cfm
        mep 1133 domain 1 association 1
          ccm-enable
          control-mep
          no shutdown
        exit
      exit
      no shutdown
    exit
    path b 1/1/c2/1 raps-tag 1
      eth-cfm
        mep 1233 domain 1 association 2
          ccm-enable
          control-mep
          no shutdown
        exit
      exit
      no shutdown
    exit
  no shutdown
exit
```

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:PE-1# show eth-ring 1
```

```
=====
```

```
Ethernet Ring 1 Information
```

```
=====
```

```
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 **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 1 name "VPLS-1" customer 1 create
  description "control channel VPLS 1 tag 1"
  sap 1/1/c1/1:1 eth-ring 1 create
  exit
  sap 1/1/c2/1:1 eth-ring 1 create
  exit
  no shutdown
exit
```

```
# on PE-2:
configure
service
  vpls 1 name "VPLS-1" customer 1 create
  description "control channel VPLS 1 tag 1"
```

```

        sap 1/1/c1/1:1 eth-ring 1 create
    exit
        sap 1/1/c2/1:1 eth-ring 1 create
    exit
    no shutdown
exit

```

```

# on PE-3:
configure
    service
        vpls 1 name "VPLS-1" customer 1 create
            description "control channel VPLS 1 tag 1"
            sap 1/1/c1/1:1 eth-ring 1 create
        exit
        sap 1/1/c2/1:1 eth-ring 1 create
        exit
    no shutdown
exit

```

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 the following message being displayed. In the following example, SAP 1/1/c3/1:1 is added to control VPLS 1 without the **eth-ring** parameter.

```

*A:PE-1>config>service>vpls# sap 1/1/c3/1:1 create
MINOR: SVCNMR #1321 Service contains an Ethernet ring control SAP

```

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: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

```

The following command shows the ring and path forwarding states.

```

*A: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: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: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        : 0xA0 ( RB BPR )
                  Node ID       : 02:0b: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/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 **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: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   : 53 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.

```
72 2023/05/04 12:46:08.692 UTC 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: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
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: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 : 05/04/2023 12:45:09
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: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 : 05/04/2023 12:46:09
Last Switch Command: noCmd
APS Rx PDU       : Request State: 0x0
                  Sub-Code       : 0x0
                  Status         : 0x20 ( BPR )
                  Node ID        : 02:0d:ff:00:00:00

=====
```

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 and/or SDPs.

```
# on PE-1:
configure
service
    vpls 100 name "VPLS-100" customer 1 create
    description "data channel VPLS 100"
    sap 1/1/c1/1:100 eth-ring 1 create
    exit
    sap 1/1/c2/1:100 eth-ring 1 create
    exit
    sap 1/1/c3/1:100 create
    exit
    no shutdown
exit
```

```
# on PE-2:
```

```
configure
service
  vpls 100 name "VPLS-100" customer 1 create
    description "data channel VPLS 100"
    sap 1/1/c1/1:100 eth-ring 1 create
    exit
    sap 1/1/c2/1:100 eth-ring 1 create
    exit
    sap 1/1/c3/1:100 create
    exit
  no shutdown
exit
```

```
# on PE-3:
configure
service
  vpls 100 name "VPLS-100" customer 1 create
    description "data channel VPLS 100"
    sap 1/1/c1/1:100 eth-ring 1 create
    exit
    sap 1/1/c2/1:100 eth-ring 1 create
    exit
    sap 1/1/c3/1:100 create
    exit
  no shutdown
exit
```

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

```
*A:PE-1# show service sap-using eth-ring
```

```
=====
Service Access Points (Ethernet Ring)
=====
```

SapId	SvcId	Eth-Ring	Path	Admin State	Oper State	Blocked	Control/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
  shutdown
```

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

```
85 2023/05/04 12:49:46.602 UTC MINOR: ETH_CFM #2001 Base
"MEP 1/1/1232 highest defect is now defRemoteCCM"
```

```

84 2023/05/04 12:49:43.312 UTC MAJOR: SVCGR #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."

83 2023/05/04 12:49:43.308 UTC MINOR: ERING #2001 Base eth-ring-1
"Eth-Ring 1 path b changed fwd state to unblocked"

82 2023/05/04 12:49:43.308 UTC MINOR: ERING #2001 Base eth-ring-1
"Eth-Ring 1 path a changed fwd state to blocked"

81 2023/05/04 12:49:43.308 UTC 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: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)
  cc (Dn/Up): Cnt 4/3 tm 000 00:12:39.000/000 00:07:58.420
  state: Cnt 7 B/F 000 00:12:35.700/000 00:08:01.000, 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 497 02:27:20.970/000 00:03:59.980
  state: Cnt 8 B/F 000 00:09:01.090/000 00:12:35.700, 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 3, pm (set/clear) 0/0, txFlush 0
  RxRaps: ok 14 nok 0 self 144, TmrExp - wtr 2(0), grd 3, wtb 0
  Flush: cnt 9 (7/2/0) tm 000 00:12:39.430-000 00:12:39.430 Out/Ack 0/1
  RxRawRaps: aPath 106 bPath 127 vPath 0
  Now: 000 00:13:19.130 , softReset: No - noTx 0

Seq Event RxInfo(Path: NodeId-Bytes)
      state:TxInfo (Bytes) Dir pA pB Time
=== =====
009 pdu B: 02:09:ff:00:00:00-0x0020 Nr
    PROT : 0xb060 Sf(DNF) Rx<-- Fwd Blk 000 00:04:01.450
010 bUp
    PEND-G: 0x0020 Nr Tx--> Fwd Blk 000 00:04:01.990
011 pdu A: 02:0d:ff:00:00:00-0x0000 Nr
    PEND-G: 0x0020 Nr Rx<-- Fwd Blk 000 00:04:01.990
012 pdu A: 02:0d:ff:00:00:00-0x0000 Nr
    PEND-G: 0x0020 Nr Rx<-- Fwd Blk 000 00:04:02.090
013 pdu B: 02:0d:ff:00:00:00-0x0000 Nr
    PEND-G: 0x0020 Nr Rx<-- Fwd Blk 000 00:04:02.090
014 pdu A: 02:0d:ff:00:00:00-0x0000 Nr
    PEND-G: 0x0020 Nr Rx<-- Fwd Blk 000 00:04:02.190
015 pdu B: 02:0d:ff:00:00:00-0x0000 Nr
    PEND-G: 0x0020 Nr Rx<-- Fwd Blk 000 00:04:02.190
016 pdu A: 02:0d:ff:00:00:00-0x0000 Nr
    PEND : 0x0020 Nr Rx<-- Fwd Blk 000 00:04:06.390
017 pdu

```

```

018      pdu B: 02:0d:ff:00:00:00-0x0000 Nr
      PEND : ----- Fwd Fwd 000 00:04:06.390
019 xWtr      PEND : Rx<-- Fwd Fwd 000 00:04:06.390
      IDLE : 0x00a0 Nr(RB ) TxF-> Fwd Blk 000 00:05:06.090
000 aDn
      PROT : 0xb000 Sf TxF-> Blk Fwd 000 00:07:17.900
001 pdu B: 02:0d:ff:00:00:00-0xb020 Sf
      PROT : 0xb000 Sf RxF<- Blk Fwd 000 00:07:21.420
002 aUp
      PEND-G: 0x0000 Nr Tx--> Blk Fwd 000 00:08:00.390
003 pdu A: 02:0d:ff:00:00:00-0x0020 Nr
      PEND : 0x0000 Nr Rx<-- Blk Fwd 000 00:08:01.000
004 pdu
      PEND : ----- Fwd Fwd 000 00:08:01.000
005 pdu B: 02:0d:ff:00:00:00-0x0020 Nr
      PEND : Rx<-- Fwd Fwd 000 00:08:01.000
006 xWtr      IDLE : 0x00a0 Nr(RB ) TxF-> Fwd Blk 000 00:09:01.090
007 aDn
      PROT : 0xb000 Sf TxF-> Blk Fwd 000 00:12:35.700
008 pdu B: 02:0d:ff:00:00:00-0xb020 Sf
      PROT : 0xb000 Sf RxF<- Blk Fwd 000 00:12:39.430

```

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.

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](#)
- [Overview](#)
- [Configuration](#)
- [Conclusion](#)

Applicability

This chapter was initially written based on SR OS Release 16.0.R5, but the 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 (**use-provisioned-sdp** or **prefer-provisioned-sdp** CLI 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:
configure
  router Base
    interface "lo1"
      address 10.0.1.1/24
      loopback
      gre-termination
      no shutdown
    exit
```

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:

```
*A:PE-1>config>router>if$ gre-termination
MINOR: CLI Could not set gre-termination for interface "lo2".
MINOR: PIP #2078 Cannot config GRE termination - already set on interface "lo1"
```

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.

```
*A:PE-1>config>router>if$ gre-termination
MINOR: CLI Could not set gre-termination for interface "int-PE-1-PE-2".
MINOR: PIP #2078 Cannot config GRE termination - already set on interface "lo1"
```

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 create
```

```

    far-end 10.0.2.1
    local-end 10.0.1.1
    no shutdown
exit

```

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:

```

*A:PE-1>config>service# sdp 122 mpls create
*A:PE-1>config>service>sdp$ local-end 10.0.1.1
MINOR: SVCMMGR #7825 Invalid local-end address - local-end not supported for this sdp 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 create
    far-end 10.0.2.1
    local-end 192.168.12.1
    no shutdown
exit

```

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.

```

*A:PE-1>config>service# sdp 120 create
*A:PE-1>config>service>sdp$ local-end 10.99.1.1
MINOR: SVCMMGR #7827 Cannot configure local-end IP address - Local router interface with
address does not exist, or address is not primary

```

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
exit

```

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



28868

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

- 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"
      address 10.0.1.1/24
      loopback
      gre-termination
      no shutdown
    exit
```

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"
    exit
```

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
  service
    sdp 120 create
      signaling tldp          # default
      far-end 10.0.2.1
      local-end 10.0.1.1
      no shutdown
    exit
    sdp 121 create
      signaling bgp
```

```


far-end 10.0.2.1
local-end 10.0.1.1
no shutdown
exit

```

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



28869

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 "lo1"
    exit
# on PE-2:
configure
router Base
  ldp
    targeted-session
      peer 10.0.1.1
      local-lsr-id "lo1"

```

exit

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



28870

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
  bgp
    rapid-withdrawal
    split-horizon
    group "internal"
      family vpn-ipv4 l2-vpn
      type internal
      local-address 10.0.1.1
      neighbor 10.0.3.1
    exit
  exit
no shutdown
```

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

```
# on RR-3:
configure
router Base
  bgp
    rapid-withdrawal
    split-horizon
    group "internal"
```

```
family vpn-ipv4 l2-vpn
type internal
cluster 10.0.3.1
local-address 10.0.3.1
neighbor 10.0.1.1
exit
neighbor 10.0.2.1
exit
exit
no shutdown
exit
```

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)

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



28871

For a description of the BGP-VPLS parameters, see the "BGP VPLS" chapter in *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Advanced Configuration Guide for Classic CLI*; for BGP-AD, see the "LDP VPLS Using BGP Auto-Discovery" chapter in *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Advanced Configuration Guide for Classic CLI*; for BGP-VPWS, see the "BGP Virtual Private Wire Services" chapter in *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Advanced Configuration Guide for Classic CLI*. For BGP-VPLS, BGP-AD, and BGP-VPWS, PW template 1 is configured with the **use-provisioned-sdp** 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 create
      far-end 10.0.2.1
      local-end 10.0.1.1
      keep-alive
      shutdown
    exit
    no shutdown
  exit
  sdp 121 create
    signaling bgp
    far-end 10.0.2.1
    local-end 10.0.1.1
    keep-alive
    shutdown
  exit
  no shutdown
exit
pw-template 1 name "PW1-use-prov-SDP" use-provisioned-sdp create
exit
vpls 1 name "VPLS-1" customer 1 create
  description "VPLS 1 with manually configured spoke-SDP"

```

```
    stp
      shutdown
    exit
    sap pxc-10.a:1 create
      no shutdown
    exit
    spoke-sdp 120:1 create
      no shutdown
    exit
    no shutdown
  exit
  epipe 2 name "Epipe-2" customer 1 create
    description "Epipe 2 with manually configured spoke-SDP"
    sap pxc-10.a:2 create
      no shutdown
    exit
    spoke-sdp 120:2 create
      no shutdown
    exit
    no shutdown
  exit
  vpls 3 name "BGP-VPLS-3" customer 1 create
    description "BGP-VPLS with use provisioned SDP"
    bgp
      route-distinguisher 64500:3
      route-target export target:64500:3 import target:64500:3
      pw-template-binding 1
    exit
  exit
  bgp-vpls
    max-ve-id 100
    ve-name "PE-1"
    ve-id 1
  exit
  no shutdown
  exit
  stp
    shutdown
  exit
  sap pxc-10.a:3 create
    no shutdown
  exit
  no shutdown
  exit
  vpls 4 name "BGP-AD VPLS-4" customer 1 create
    description "BGP-AD for LDP VPLS with use provisioned SDP"
    bgp
      route-distinguisher 64500:4
      route-target export target:64500:4 import target:64500:4
      pw-template-binding 1
    exit
  exit
  bgp-ad
    vpls-id 64500:4
    no shutdown
  exit
  stp
    shutdown
  exit
  sap pxc-10.a:4 create
    no shutdown
  exit
  no shutdown
  exit
```

```

epipe 5 name "BGP-VPWS-5" customer 1 create
description "BGP-VPWS with use provisioned SDP"
bgp
    route-distinguisher 64500:5
    route-target export target:64500:5 import target:64500:5
    pw-template-binding 1
    exit
exit
bgp-vpws
    ve-name "PE-1"
    ve-id 1
    exit
    remote-ve-name "PE-2"
    ve-id 2
    exit
    no shutdown
exit
sap pxc-10.a:5 create
    no shutdown
exit
    no shutdown
exit

```

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

```

*A: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    13    0 00h02m48s 0/0/0 (VpnIPv4)
                   15     0      3/3/3 (L2VPN)
-----

```

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

```

*A: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 52        53        0d 00:03:39
-----
No. of IPv4 Sessions: 1
=====

```

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: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:PE-1# show service sdp-using
```

```
=====
```

SDP Using							
SvcId	SdpId	Type	Far End	Opr State	I.Label	E.Label	
1	120:1	Spok	10.0.2.1	Up	524286	524286	
2	120:2	Spok	10.0.2.1	Up	524285	524285	
3	121:4294967295	BgpVp*	10.0.2.1	Up	524278	524277	
4	120:4294967294	BgpAd	10.0.2.1	Up	524275	524275	
5	121:4294967293	BgpVp*	10.0.2.1	Up	524276	524276	

```
-----
```

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:PE-1# ping router 11 10.0.11.21 rapid
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 = 3.58ms, avg = 5.11ms, max = 10.3ms, stddev = 2.59ms
*A:PE-1# ping router 12 10.0.12.22 rapid
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 = 3.37ms, avg = 4.54ms, max = 8.83ms, stddev = 2.15ms
*A:PE-1# ping router 13 10.0.13.23 rapid
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 = 3.24ms, avg = 4.32ms, max = 8.02ms, stddev = 1.85ms
*A:PE-1# ping router 14 10.0.14.24 rapid
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 = 3.31ms, avg = 4.45ms, max = 8.72ms, stddev = 2.14ms
*A:PE-1# ping router 15 10.0.15.25 rapid
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 = 3.34ms, avg = 4.93ms, max = 8.62ms, stddev = 1.98ms
```

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



28872

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 10.0.1.1 comm-val 60000 to 65000
exit
vprn 6 name "VPRN-6 with GRE spoke-SDP" customer 1 create
    interface "lo6" create
        address 172.16.16.1/32
```

```

        loopback
    exit
    bgp-ipvpn
    mpls
        route-distinguisher auto-rd
        vrf-target target:64500:6
        no shutdown
    exit
exit
spoke-sdp 120:6 create
exit
no shutdown
exit

```

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

```
*A:PE-1# show router 6 fib 1
```

FIB Display

Prefix [Flags] NextHop	Protocol
172.16.16.1/32 172.16.16.1 (lo6)	LOCAL
172.16.26.2/32 10.0.2.1 (VPRN Label:524274 Transport:SDP:120)	BGP_VPN
Total Entries : 2	

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
        exit
    vprn 7 name "VPRN-7 with auto-bind GRE" customer 1 create
        interface "lo7" create
            address 172.16.17.1/24
            loopback
        exit
    bgp-ipvpn
    mpls
        auto-bind-tunnel
        resolution-filter
            gre
        exit
        resolution filter
    exit
    route-distinguisher auto-rd
    vrf-target target:64500:7
    no shutdown
exit
exit

```

```
no shutdown
exit
```

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

```
*A: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:524273 Transport:GRE)
-----
Total Entries : 2
=====
```

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

```
# on PE-1:
configure
service
  ies 8 name "IES-8" customer 1 create
  interface "lo8" create
    address 172.16.18.1/24
    loopback
  exit
  interface "int-IES8-PE-1-PE-2" create
    address 172.16.128.1/30
    spoke-sdp 120:8 create
    no shutdown
  exit
exit
no shutdown
exit
```

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

```
*A:PE-1# ping 172.16.128.2 rapid
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 = 2.44ms, avg = 2.54ms, max = 2.69ms, stddev = 0.081ms
```

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.

Inter-AS Option B Label Security for IP-VPN and EVPN Routes

This chapter provides information about inter-AS option B label security for IP-VPN and EVPN routes.

Topics in this chapter include:

- [Applicability](#)
- [Overview](#)
- [Configuration](#)
- [Conclusion](#)

Applicability

The information and the configuration in this chapter are based on SR OS Release 24.3.R1. Inter-AS option B label security for IP-VPN routes is supported in SR OS Release 16.0.R4, and later. Inter-AS option B label security for EVPN routes is supported in SR OS Release 23.3.R2, and later.

Overview

In inter-AS option B interconnects, the Autonomous System Border Routers (ASBRs) can filter BGP IP-VPN or BGP EVPN routes based on route target (RT). In addition, BGP neighbor trust prevents label spoofing in inter-AS option B for the VPN-IPv4, VPN-IPv6, and EVPN address families. In networks where ASBRs advertise routes to multiple peer ASBRs, an ASBR may drop packets on IP interfaces that are configured as **untrusted** with the **default-forwarding** argument set to the **drop** command option:

```
# on ASBR: configure router interface <..> untrusted default-forwarding drop
```

By default, all IP interfaces between ASBRs are trusted and the datapath allows all packets. It is possible to configure a number of maximum 15 interfaces as **untrusted**. The **default-forwarding** argument can be set to the **forward** option (default behavior) or to the **drop** option.



Note:

When an IP interface is configured as **untrusted** without the **default-forwarding drop** option or when the untrusted IP interface is configured with the (default) **default-forwarding forward** option, the datapath allows all packets and the behavior is the same as when the **untrusted** command is not configured.

Traffic is only dropped when the IP interface is configured with **untrusted default-forwarding drop**.

[Table 3: Untrusted interfaces with default-forwarding forward option allow all IP-VPN and EVPN routes](#) shows that the datapath allows all IP-VPN and EVPN traffic when the interface is configured as **untrusted**

with **default-forwarding** set to **forward**. There is no need to configure neighbor-trust for VPN-IPv4, VPN-IPv6, or EVPN.

Table 3: Untrusted interfaces with default-forwarding forward option allow all IP-VPN and EVPN routes

untrusted configuration	neighbor-trust configured			traffic allowed		
	VPN-IPv4	VPN-IPv6	EVPN	VPN-IPv4	VPN-IPv6	EVPN
untrusted forward	no	no	no	yes	yes	yes
untrusted forward	no	no	yes	yes	yes	yes
untrusted forward	no	yes	no	yes	yes	yes
untrusted forward	no	yes	yes	yes	yes	yes
untrusted forward	yes	no	no	yes	yes	yes
untrusted forward	yes	no	yes	yes	yes	yes
untrusted forward	yes	yes	no	yes	yes	yes
untrusted forward	yes	yes	yes	yes	yes	yes

In contrast, the datapath drops all labeled packets on untrusted IP interfaces configured with the **default-forwarding drop** option. To allow the datapath to provide an exception to the default forwarding handling for Ingress Label Maps (ILMs), BGP must flag those ILMs to the data path. The following **neighbor-trust** command is used to enable the exceptional ILM forwarding behavior for multiple VPN address families: VPN-IPv4, VPN-IPv6, and EVPN:

```
# on ASBR: configure router bgp neighbor-trust { vpn-ipv4 | vpn-ipv6 | evpn }
```

Table 4: BGP neighbor-trust defines what traffic is allowed on untrusted interfaces with default-forwarding drop option shows what traffic is allowed on an untrusted interface configured with the **default-forwarding drop** option when BGP **neighbor-trust** is configured for VPN-IP or EVPN address families.

Table 4: BGP neighbor-trust defines what traffic is allowed on untrusted interfaces with default-forwarding drop option

untrusted configuration	neighbor-trust configured			traffic allowed		
	VPN-IPv4	VPN-IPv6	EVPN	VPN-IPv4	VPN-IPv6	EVPN
untrusted drop	no	no	no	no	no	no
untrusted drop	no	no	yes	no	no	yes

untrusted configuration	neighbor-trust configured			traffic allowed		
	VPN-IPv4	VPN-IPv6	EVPN	VPN-IPv4	VPN-IPv6	EVPN
untrusted drop	no	yes	no	no	yes	no
untrusted drop	no	yes	yes	no	yes	yes
untrusted drop	yes	no	no	yes	no	no
untrusted drop	yes	no	yes	yes	no	yes
untrusted drop	yes	yes	no	yes	yes	no
untrusted drop	yes	yes	yes	yes	yes	yes

Configuration

The following scenarios are described in this chapter:


- [Inter-AS option B label security with services configured on PEs only](#)
- [Inter-AS option B label security with services configured on PEs and on ASBR](#)

Inter-AS option B label security with services configured on PEs only

[Figure 19: Example topology with services on PEs](#) shows the example topology with the following services configured on PE-1 and PE-4:

- BGP-IPVPN "VPRN-1"
- BGP-EVPN "VPRN-2"
- EVPN VPLS "VPLS-3"
- EVPN VPWS "Epipe-4"
- EVPN R-VPLS "BD-5" in IES "IES-55"

Figure 19: Example topology with services on PEs



Initial configuration

The initial configuration on the nodes in the example topology includes the following:

- cards, MDAs, ports
- router interfaces
- IS-IS between PE-1, RR-5, and ASBR-2 in AS 64496 and between PE-4, RR-6, and ASBR-3 in AS 64511, but not between the ASBRs
- SR-ISIS between PE-1 and ASBR-2 in AS 64496 and between PE-4 and ASBR-3 in AS 64511
- BGP for the VPN-IPv4, VPN-IPv6, and EVPN address families:
 - IBGP in AS 64496 with route reflector RR-5 and clients PE-1 and ASBR-2
 - IBGP in AS 64511 with route reflector RR-6 and clients PE-4 and ASBR-3
 - EBGp between ASBR-2 and ASBR-3

The BGP configuration on PE-1 is as follows:

```
# on PE-1:
configure
  router Base
    autonomous-system 64496
    bgp
      rapid-withdrawal
      split-horizon
      rapid-update vpn-ipv4 vpn-ipv6 evpn
      group "internal"
        peer-as 64496
        neighbor 192.0.2.5
          family vpn-ipv4 vpn-ipv6 evpn
      exit
    exit
```

The BGP configuration on RR-5 is as follows:

```
# on RR-5:
configure
  router Base
    autonomous-system 64496
    bgp
      rapid-withdrawal
      split-horizon
      rapid-update vpn-ipv4 vpn-ipv6 evpn
      group "internal"
        cluster 192.0.2.5
        peer-as 64496
        neighbor 192.0.2.1
          family vpn-ipv4 vpn-ipv6 evpn
      exit
      neighbor 192.0.2.2
        family vpn-ipv4 vpn-ipv6 evpn
      exit
    exit
```

The BGP configuration on ASBR-2 is as follows:

```
# on ASBR-2:
```

```
configure
router Base
  autonomous-system 64496
  bgp
    enable-inter-as-vpn      # required for inter-AS VPRN model B
    rapid-withdrawal
    split-horizon
    rapid-update vpn-ipv4 vpn-ipv6 evpn
    next-hop-resolution
      labeled-routes
        transport-tunnel
        family vpn
        resolution any
      exit
    exit
  exit
  group "external"
    type external
    peer-as 64511
    neighbor 192.168.23.2
      family vpn-ipv4 vpn-ipv6 evpn
    exit
  exit
  group "internal"
    peer-as 64496
    neighbor 192.0.2.5
      family vpn-ipv4 vpn-ipv6 evpn
    exit
  exit
exit
```

The BGP configuration on the nodes in AS 64511 is similar.

Services configuration

The following services are configured on PE-1:

```
# on PE-1:
configure
service
  vprn 1 name "VPRN-1" customer 1 create
  interface "int-test-1" create
    address 10.1.1.1/24
    ipv6
      address 2001:db8::10:1:1:1/120
    exit
    sap 1/1/cl0/1:1 create
    exit
  exit
  bgp-ipvpn
  mpls
    auto-bind-tunnel
    resolution any
    exit
    route-distinguisher 192.0.2.1:1
    vrf-target target:64496:1
    no shutdown
  exit
  exit
  no shutdown
exit
```

```
vrpn 2 name "VPRN-2" customer 1 create
  interface "int-test-2" create
    address 10.2.1.1/24
    ipv6
      address 2001:db8::10:2:1:1/120
    exit
    sap 1/1/c10/1:2 create
    exit
  exit
  bgp-evpn
    mpls
      auto-bind-tunnel
      resolution any
    exit
    route-distinguisher 192.0.2.1:2
    vrf-target target:64496:2
    no shutdown
  exit
exit
no shutdown
exit
vpls 3 name "VPLS-3" customer 1 create
  bgp
    route-target export target:64496:3 import target:64496:3
  exit
  bgp-evpn
    evi 3
    mpls
      auto-bind-tunnel
      resolution any
    exit
    no shutdown
  exit
exit
sap 1/1/c10/1:3 create
  no shutdown
exit
no shutdown
exit
pipe 4 name "Epipe-4" customer 1 create
  bgp
    route-target export target:64496:4 import target:64496:4
  exit
  bgp-evpn
    local-attachment-circuit PE1 create
      eth-tag 1
    exit
    remote-attachment-circuit PE4 create
      eth-tag 4
    exit
    evi 4
    mpls
      auto-bind-tunnel
      resolution any
    exit
    no shutdown
  exit
exit
sap 1/1/c10/1:4 create
  description "SAP to CE-41"
  no shutdown
exit
no shutdown
exit
```

```
vpls 5 name "BD-5" customer 1 create
  allow-ip-int-bind
  exit
  bgp
    route-target export target:64496:5 import target:64496:5
  exit
  bgp-evpn
    evi 5
    mpls bgp 1
    auto-bind-tunnel
    resolution any
    exit
    no shutdown
  exit
exit
stp
  shutdown
exit
sap 1/1/c10/1:5 create
  no shutdown
exit
no shutdown
exit
ies 55 name "IES-55" customer 1 create
  interface "int-BD-5" create
    address 172.16.5.1/24
    ipv6
      address 2001:db8::16:5:1/120
    exit
    vpls "BD-5"
    exit
  exit
  no shutdown
exit
```

The configuration of the services on PE-4 in AS 64511 is similar.

Inter-AS option B services using trusted interfaces

By default, IP interfaces are trusted. With trusted interfaces between ASBR-2 and ASBR-3, traffic can be sent from the services or the CEs connected to the services on PE-1 to the corresponding services on PE-4.

Inter-AS option B services using untrusted interfaces with default-forwarding forward option

It is possible to configure the interface from ASBR-3 to ASBR-2 as **untrusted** with the **default-forwarding** argument set to the **forward** option, or even without this **default-forwarding** argument, because the default option is **forward**:

```
# on ASBR-3:
configure
  router Base
    interface "int-ASBR-3-ASBR-2"
      address 192.168.23.2/30
      port 1/1/c2/1:1000
      untrusted default-forwarding forward    # default option forward
      no shutdown
```

```
exit
```

With this configuration where packets on the untrusted interfaces are forwarded by default, it is possible to send traffic between the services on PE-1 and the services on PE-4:

```
*A:PE-1# ping router-instance "VPRN-1" 10.1.4.4 rapid          # VPN-IPv4
PING 10.1.4.4 56 data bytes
!!!!
---- 10.1.4.4 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.29ms, avg = 2.59ms, max = 3.27ms, stddev = 0.351ms

*A:PE-1# ping router-instance "VPRN-1" 2001:db8::10:1:4:4 rapid # VPN-IPv6
PING 2001:db8::10:1:4:4 56 data bytes
!!!!
---- 2001:db8::10:1:4:4 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.29ms, avg = 2.57ms, max = 3.18ms, stddev = 0.312ms

*A:PE-1# ping router-instance "VPRN-2" 10.2.4.4 rapid          # EVPN IFL
PING 10.2.4.4 56 data bytes
!!!!
---- 10.2.4.4 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.14ms, avg = 2.56ms, max = 3.20ms, stddev = 0.357ms

*A:PE-1# ping router-instance "VPRN-2" 2001:db8::10:2:4:4 rapid # EVPN IFL
PING 2001:db8::10:2:4:4 56 data bytes
!!!!
---- 2001:db8::10:2:4:4 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.19ms, avg = 2.49ms, max = 2.96ms, stddev = 0.259ms

*A:PE-1# ping router-instance "CE-31" 172.16.3.4 rapid          # EVPN VPLS
PING 172.16.3.4 56 data bytes
!!!!
---- 172.16.3.4 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.86ms, avg = 3.08ms, max = 3.44ms, stddev = 0.208ms

*A:PE-1# ping router-instance "CE-31" 2001:db8::16:3:4 rapid     # EVPN VPLS
PING 2001:db8::16:3:4 56 data bytes
!!!!
---- 2001:db8::16:3:4 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.91ms, avg = 3.19ms, max = 3.60ms, stddev = 0.238ms

*A:PE-1# ping router-instance "CE-41" 172.16.4.4 rapid          # EVPN VPWS
PING 172.16.4.4 56 data bytes
!!!!
---- 172.16.4.4 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.84ms, avg = 3.23ms, max = 3.95ms, stddev = 0.381ms

*A:PE-1# ping router-instance "CE-41" 2001:db8::16:4:4 rapid     # EVPN VPWS
PING 2001:db8::16:4:4 56 data bytes
!!!!
---- 2001:db8::16:4:4 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.62ms, avg = 2.94ms, max = 3.58ms, stddev = 0.334ms

*A:PE-1# ping router-instance "CE-51" 172.16.5.54 rapid          # EVPN R-VPLS
PING 172.16.5.54 56 data bytes
```

```
!!!!
---- 172.16.5.54 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 3.10ms, avg = 3.32ms, max = 3.45ms, stddev = 0.120ms

*A:PE-1# ping router-instance "CE-51" 2001:db8::16:5:54 rapid    # EVPN R-VPLS
PING 2001:db8::16:5:54 56 data bytes
!!!!
---- 2001:db8::16:5:54 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 3.21ms, avg = 3.41ms, max = 3.87ms, stddev = 0.241ms
```

All traffic is forwarded, so there is no need to configure the **neighbor-trust** command. If the **neighbor-trust** command is configured for VPN-IPv4, VPN-IPv6, EVPN, or any combination of these, this command has no effect. As an example, the **neighbor-trust** command is configured for the VPN-IPv4 and EVPN address families, as follows:

```
# on ASBR-3:
configure
  router Base
    bgp
      neighbor-trust vpn-ipv4 evpn
```

The datapath forwards all traffic for the corresponding services, regardless of this **neighbor-trust** configuration:

```
*A:PE-1# ping router-instance "VPRN-1" 10.1.4.4 rapid          # VPN-IPv4
PING 10.1.4.4 56 data bytes
!!!!
---- 10.1.4.4 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.00ms, avg = 2.32ms, max = 2.98ms, stddev = 0.339ms

*A:PE-1# ping router-instance "VPRN-1" 2001:db8::10:1:4:4 rapid # VPN-IPv6
PING 2001:db8::10:1:4:4 56 data bytes
!!!!
---- 2001:db8::10:1:4:4 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.08ms, avg = 2.49ms, max = 3.21ms, stddev = 0.381ms

*A:PE-1# ping router-instance "VPRN-2" 10.2.4.4 rapid          # EVPN IFL
PING 10.2.4.4 56 data bytes
!!!!
---- 10.2.4.4 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.16ms, avg = 2.41ms, max = 2.99ms, stddev = 0.305ms

*A:PE-1# ping router-instance "VPRN-2" 2001:db8::10:2:4:4 rapid # EVPN IFL
PING 2001:db8::10:2:4:4 56 data bytes
!!!!
---- 2001:db8::10:2:4:4 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.14ms, avg = 2.47ms, max = 2.99ms, stddev = 0.308ms

*A:PE-1# ping router-instance "CE-31" 172.16.3.4 rapid         # EVPN VPLS
PING 172.16.3.4 56 data bytes
!!!!
---- 172.16.3.4 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.87ms, avg = 3.80ms, max = 5.55ms, stddev = 0.973ms
```

```
*A:PE-1# ping router-instance "CE-31" 2001:db8::16:3:4 rapid      # EVPN VPLS
PING 2001:db8::16:3:4 56 data bytes
!!!!
---- 2001:db8::16:3:4 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 3.03ms, avg = 3.30ms, max = 3.90ms, stddev = 0.306ms

*A:PE-1# ping router-instance "CE-41" 172.16.4.4 rapid          # EVPN VPWS
PING 172.16.4.4 56 data bytes
!!!!
---- 172.16.4.4 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 3.05ms, avg = 3.41ms, max = 4.26ms, stddev = 0.444ms

*A:PE-1# ping router-instance "CE-41" 2001:db8::16:4:4 rapid    # EVPN VPWS
PING 2001:db8::16:4:4 56 data bytes
!!!!
---- 2001:db8::16:4:4 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.63ms, avg = 3.17ms, max = 3.67ms, stddev = 0.349ms

*A:PE-1# ping router-instance "CE-51" 172.16.5.54 rapid        # EVPN R-VPLS
PING 172.16.5.54 56 data bytes
!!!!
---- 172.16.5.54 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.89ms, avg = 3.37ms, max = 3.84ms, stddev = 0.338ms

*A:PE-1# ping router-instance "CE-51" 2001:db8::16:5:54 rapid   # EVPN R-VPLS
PING 2001:db8::16:5:54 56 data bytes
!!!!
---- 2001:db8::16:5:54 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.73ms, avg = 3.11ms, max = 3.56ms, stddev = 0.284ms
```

When **no untrusted** is configured on the interface, the interface is trusted and the connectivity remains.

Inter-AS option B services using untrusted interfaces with default-forwarding drop option

The following command on ASBR-2 configures the IP interface "int-ASBR-2-ASBR-3" as **untrusted** with **default-forwarding** argument set to **drop**:

```
# on ASBR-2:
configure
  router Base
    interface "int-ASBR-2-ASBR-3"
      address 192.168.23.1/30
      port 1/1/c1/1:1000
      untrusted default-forwarding drop
      no shutdown
```

When no **neighbor-trust** command is configured, the datapath drops all traffic for the configured services, as follows:

```
*A:PE-1# ping router-instance "VPRN-1" 10.1.4.4 rapid          # VPN-IPv4
PING 10.1.4.4 56 data bytes
....
---- 10.1.4.4 PING Statistics ----
```

```

5 packets transmitted, 0 packets received, 100% packet loss

*A:PE-1# ping router-instance "VPRN-1" 2001:db8::10:1:4:4 rapid # VPN-IPv6
PING 2001:db8::10:1:4:4 56 data bytes
.....
---- 2001:db8::10:1:4:4 PING Statistics ----
5 packets transmitted, 0 packets received, 100% packet loss

*A:PE-1# ping router-instance "VPRN-2" 10.2.4.4 rapid # EVPN IFL
PING 10.2.4.4 56 data bytes
.....
---- 10.2.4.4 PING Statistics ----
5 packets transmitted, 0 packets received, 100% packet loss

*A:PE-1# ping router-instance "VPRN-2" 2001:db8::10:2:4:4 rapid # EVPN IFL
PING 2001:db8::10:2:4:4 56 data bytes
.....
---- 2001:db8::10:2:4:4 PING Statistics ----
5 packets transmitted, 0 packets received, 100% packet loss

*A:PE-1# ping router-instance "CE-31" 172.16.3.4 rapid # EVPN VPLS
PING 172.16.3.4 56 data bytes
.....
---- 172.16.3.4 PING Statistics ----
5 packets transmitted, 0 packets received, 100% packet loss

*A:PE-1# ping router-instance "CE-31" 2001:db8::16:3:4 rapid # EVPN VPLS
PING 2001:db8::16:3:4 56 data bytes
.....
---- 2001:db8::16:3:4 PING Statistics ----
5 packets transmitted, 0 packets received, 100% packet loss

*A:PE-1# ping router-instance "CE-41" 172.16.4.4 rapid # EVPN VPWS
PING 172.16.4.4 56 data bytes
.....
---- 172.16.4.4 PING Statistics ----
5 packets transmitted, 0 packets received, 100% packet loss

*A:PE-1# ping router-instance "CE-41" 2001:db8::16:4:4 rapid # EVPN VPWS
PING 2001:db8::16:4:4 56 data bytes
.....
---- 2001:db8::16:4:4 PING Statistics ----
5 packets transmitted, 0 packets received, 100% packet loss

*A:PE-1# ping router-instance "CE-51" 172.16.5.54 rapid # EVPN R-VPLS
PING 172.16.5.54 56 data bytes
.....
---- 172.16.5.54 PING Statistics ----
5 packets transmitted, 0 packets received, 100% packet loss

*A:PE-1# ping router-instance "CE-51" 2001:db8::16:5:54 rapid # EVPN R-VPLS
PING 2001:db8::16:5:54 56 data bytes
.....
---- 2001:db8::16:5:54 PING Statistics ----
5 packets transmitted, 0 packets received, 100% packet loss

```

When **neighbor-trust** is configured for the VPN-IPv4 address family, the datapath allows IPv4 traffic in VPRN-1 between PE-1 and PE-4 (but not traffic for services using the other address families):

```

# on ASBR-2:
configure
router Base
  bgp

```

neighbor-trust vpn-ipv4

```
*A:PE-1# ping router-instance "VPRN-1" 10.1.4.4 rapid          # VPN-IPv4
PING 10.1.4.4 56 data bytes
!!!!
---- 10.1.4.4 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.34ms, avg = 2.57ms, max = 3.10ms, stddev = 0.274ms

*A:PE-1# ping router-instance "VPRN-1" 2001:db8::10:1:4:4 rapid # VPN-IPv6
PING 2001:db8::10:1:4:4 56 data bytes
.....
---- 2001:db8::10:1:4:4 PING Statistics ----
5 packets transmitted, 0 packets received, 100% packet loss

*A:PE-1# ping router-instance "VPRN-2" 10.2.4.4 rapid          # EVPN IFL
PING 10.2.4.4 56 data bytes
.....
---- 10.2.4.4 PING Statistics ----
5 packets transmitted, 0 packets received, 100% packet loss

*A:PE-1# ping router-instance "VPRN-2" 2001:db8::10:2:4:4 rapid # EVPN IFL
PING 2001:db8::10:2:4:4 56 data bytes
.....
---- 2001:db8::10:2:4:4 PING Statistics ----
5 packets transmitted, 0 packets received, 100% packet loss

*A:PE-1# ping router-instance "CE-31" 172.16.3.4 rapid          # EVPN VPLS
PING 172.16.3.4 56 data bytes
.....
---- 172.16.3.4 PING Statistics ----
5 packets transmitted, 0 packets received, 100% packet loss

*A:PE-1# ping router-instance "CE-31" 2001:db8::16:3:4 rapid    # EVPN VPLS
PING 2001:db8::16:3:4 56 data bytes
.....
---- 2001:db8::16:3:4 PING Statistics ----
5 packets transmitted, 5 packets bounced, 0 packets received, 100% packet loss

*A:PE-1# ping router-instance "CE-41" 172.16.4.4 rapid          # EVPN VPWS
PING 172.16.4.4 56 data bytes
.....
---- 172.16.4.4 PING Statistics ----
5 packets transmitted, 0 packets received, 100% packet loss

*A:PE-1# ping router-instance "CE-41" 2001:db8::16:4:4 rapid    # EVPN VPWS
PING 2001:db8::16:4:4 56 data bytes
.....
---- 2001:db8::16:4:4 PING Statistics ----
5 packets transmitted, 5 packets bounced, 0 packets received, 100% packet loss

*A:PE-1# ping router-instance "CE-51" 172.16.5.54 rapid          # EVPN R-VPLS
PING 172.16.5.54 56 data bytes
.....
---- 172.16.5.54 PING Statistics ----
5 packets transmitted, 0 packets received, 100% packet loss

*A:PE-1# ping router-instance "CE-51" 2001:db8::16:5:54 rapid    # EVPN R-VPLS
PING 2001:db8::16:5:54 56 data bytes
.....
---- 2001:db8::16:5:54 PING Statistics ----
5 packets transmitted, 5 packets bounced, 0 packets received, 100% packet loss
```

When **neighbor-trust** is configured for the VPN-IPv4 and VPN-IPv6 address families, the datapath allows IPv4 and IPv6 traffic in VPRN-1 between PE-1 and PE-4 (but not traffic for services using the EVPN address family):

```
# on ASBR-2:
configure
  router Base
    bgp
      neighbor-trust vpn-ipv4 vpn-ipv6

*A:PE-1# ping router-instance "VPRN-1" 10.1.4.4 rapid          # VPN-IPv4
PING 10.1.4.4 56 data bytes
!!!!
---- 10.1.4.4 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.72ms, avg = 2.98ms, max = 3.58ms, stddev = 0.306ms

*A:PE-1# ping router-instance "VPRN-1" 2001:db8::10:1:4:4 rapid # VPN-IPv6
PING 2001:db8::10:1:4:4 56 data bytes
!!!!
---- 2001:db8::10:1:4:4 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.54ms, avg = 2.79ms, max = 3.37ms, stddev = 0.294ms

*A:PE-1# ping router-instance "VPRN-2" 10.2.4.4 rapid          # EVPN IFL
PING 10.2.4.4 56 data bytes
.....
---- 10.2.4.4 PING Statistics ----
5 packets transmitted, 0 packets received, 100% packet loss

*A:PE-1# ping router-instance "VPRN-2" 2001:db8::10:2:4:4 rapid # EVPN IFL
PING 2001:db8::10:2:4:4 56 data bytes
.....
---- 2001:db8::10:2:4:4 PING Statistics ----
5 packets transmitted, 0 packets received, 100% packet loss

*A:PE-1# ping router-instance "CE-31" 172.16.3.4 rapid          # EVPN VPLS
PING 172.16.3.4 56 data bytes
.....
---- 172.16.3.4 PING Statistics ----
5 packets transmitted, 0 packets received, 100% packet loss

*A:PE-1# ping router-instance "CE-31" 2001:db8::16:3:4 rapid    # EVPN VPLS
PING 2001:db8::16:3:4 56 data bytes
.....
---- 2001:db8::16:3:4 PING Statistics ----
5 packets transmitted, 5 packets bounced, 0 packets received, 100% packet loss

*A:PE-1# ping router-instance "CE-41" 172.16.4.4 rapid          # EVPN VPWS
PING 172.16.4.4 56 data bytes
.....
---- 172.16.4.4 PING Statistics ----
5 packets transmitted, 0 packets received, 100% packet loss

*A:PE-1# ping router-instance "CE-41" 2001:db8::16:4:4 rapid    # EVPN VPWS
PING 2001:db8::16:4:4 56 data bytes
.....
---- 2001:db8::16:4:4 PING Statistics ----
5 packets transmitted, 5 packets bounced, 0 packets received, 100% packet loss

*A:PE-1# ping router-instance "CE-51" 172.16.5.54 rapid          # EVPN R-VPLS
PING 172.16.5.54 56 data bytes
```

```

.....
---- 172.16.5.54 PING Statistics ----
5 packets transmitted, 0 packets received, 100% packet loss

*A:PE-1# ping router-instance "CE-51" 2001:db8::16:5:54 rapid      # EVPN R-VPLS
PING 2001:db8::16:5:54 56 data bytes

---- 2001:db8::16:5:54 PING Statistics ----
5 packets transmitted, 5 packets bounced, 0 packets received, 100% packet loss

```

When **neighbor-trust** is configured for the EVPN address family only, the datapath allows traffic in VPRN-2, VPLS-3, Epipe-4, and EVPN R-VPLS BD-5 between PE-1 and PE-4, but not in IP-VPN VPRN-1 (which does not use the EVPN address family):

```

# on ASBR-2:
configure
  router Base
    bgp
      neighbor-trust evpn

```

```

*A:PE-1# ping router-instance "VPRN-1" 10.1.4.4 rapid            # VPN-IPv4
PING 10.1.4.4 56 data bytes
.....
---- 10.1.4.4 PING Statistics ----
5 packets transmitted, 0 packets received, 100% packet loss

*A:PE-1# ping router-instance "VPRN-1" 2001:db8::10:1:4:4 rapid # VPN-IPv6
PING 2001:db8::10:1:4:4 56 data bytes
.....
---- 2001:db8::10:1:4:4 PING Statistics ----
5 packets transmitted, 0 packets received, 100% packet loss

*A:PE-1# ping router-instance "VPRN-2" 10.2.4.4 rapid            # EVPN IFL
PING 10.2.4.4 56 data bytes
!!!!!!
---- 10.2.4.4 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 1.98ms, avg = 2.35ms, max = 2.89ms, stddev = 0.343ms

*A:PE-1# ping router-instance "VPRN-2" 2001:db8::10:2:4:4 rapid # EVPN IFL
PING 2001:db8::10:2:4:4 56 data bytes
!!!!!!
---- 2001:db8::10:2:4:4 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.16ms, avg = 2.30ms, max = 2.78ms, stddev = 0.239ms

*A:PE-1# ping router-instance "CE-31" 172.16.3.4 rapid           # EVPN VPLS
PING 172.16.3.4 56 data bytes
!!!!!!
---- 172.16.3.4 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.56ms, avg = 2.93ms, max = 3.52ms, stddev = 0.332ms

*A:PE-1# ping router-instance "CE-31" 2001:db8::16:3:4 rapid      # EVPN VPLS
PING 2001:db8::16:3:4 56 data bytes
!!!!!!
---- 2001:db8::16:3:4 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.53ms, avg = 3.64ms, max = 7.28ms, stddev = 1.83ms

*A:PE-1# ping router-instance "CE-41" 172.16.4.4 rapid           # EVPN VPWS
PING 172.16.4.4 56 data bytes

```

```

!!!!
---- 172.16.4.4 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.77ms, avg = 3.02ms, max = 3.48ms, stddev = 0.279ms

*A:PE-1# ping router-instance "CE-41" 2001:db8::16:4:4 rapid      # EVPN VPWS
PING 2001:db8::16:4:4 56 data bytes
!!!!
---- 2001:db8::16:4:4 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.55ms, avg = 3.74ms, max = 8.13ms, stddev = 2.20ms

*A:PE-1# ping router-instance "CE-51" 172.16.5.54 rapid          # EVPN R-VPLS
PING 172.16.5.54 56 data bytes
!!!!
---- 172.16.5.54 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.82ms, avg = 3.14ms, max = 3.77ms, stddev = 0.330ms


*A:PE-1# ping router-instance "CE-51" 2001:db8::16:5:54 rapid    # EVPN R-VPLS
PING 2001:db8::16:5:54 56 data bytes
!!!!
---- 2001:db8::16:5:54 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.38ms, avg = 3.69ms, max = 7.19ms, stddev = 1.77ms

```

Inter-AS option B label security with services configured on PEs and on ASBR

BGP neighbor trust is not supported on PE-ASBRs for VPLS or Epipe services, as shown for ASBR-2 in the following example. [Figure 20: Example topology with services on PEs and on ASBR-2](#) shows the topology with services on ASBR-2 as well as on the PEs.

Figure 20: Example topology with services on PEs and on ASBR-2



The service configuration on ASBR-2 is similar to the service configuration on PE-1 and PE-4. Epipe-44 is an Epipe between ASBR-2 and PE-4, but the other services are the same as in the PEs. The interface between ASBR-2 and ASBR-3 remains untrusted with **default-forwarding** set to **drop**. The **neighbor-trust** command on ASBR-2 is configured for VPN-IPv4, VPN-IPv6, and EVPN, as follows:

```
# on ASBR-2:
```

```
configure
router Base
  bgp
    neighbor-trust vpn-ipv4 vpn-ipv6 evpn
```

The datapath allows traffic for the VPRN services on ASBR-2 and PE-4 (using VPN-IPv4, VPN-IPv6, or EVPN-IFL), but the traffic between the EVPN VPLS and EVPN VPWS services on ASBR-2 and PE-4 is dropped, as follows:

```
*A:ASBR-2# ping router-instance "VPRN-1" 10.1.4.4 rapid          # VPN-IPv4
PING 10.1.4.4 56 data bytes
!!!!
---- 10.1.4.4 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 1.56ms, avg = 1.93ms, max = 2.32ms, stddev = 0.260ms

*A:ASBR-2# ping router-instance "VPRN-1" 2001:db8::10:1:4:4 rapid # VPN-IPv6
PING 2001:db8::10:1:4:4 56 data bytes
!!!!
---- 2001:db8::10:1:4:4 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 1.75ms, avg = 2.00ms, max = 2.44ms, stddev = 0.238ms

*A:ASBR-2# ping router-instance "VPRN-2" 10.2.4.4 rapid          # EVPN IFL
PING 10.2.4.4 56 data bytes
!!!!
---- 10.2.4.4 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 1.79ms, avg = 2.00ms, max = 2.22ms, stddev = 0.152ms

*A:ASBR-2# ping router-instance "VPRN-2" 2001:db8::10:2:4:4 rapid # EVPN IFL
PING 2001:db8::10:2:4:4 56 data bytes
!!!!
---- 2001:db8::10:2:4:4 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 1.51ms, avg = 1.96ms, max = 2.41ms, stddev = 0.291ms

*A:ASBR-2# ping router-instance "CE-32" 172.16.3.4 rapid          # EVPN VPLS
PING 172.16.3.4 56 data bytes
.....
---- 172.16.3.4 PING Statistics ----
5 packets transmitted, 0 packets received, 100% packet loss

*A:ASBR-2# ping router-instance "CE-32" 2001:db8::16:3:4 rapid    # EVPN VPLS
PING 2001:db8::16:3:4 56 data bytes
.....
---- 2001:db8::16:3:4 PING Statistics ----
5 packets transmitted, 5 packets bounced, 0 packets received, 100% packet loss

*A:ASBR-2# ping router-instance "CE-442" 172.16.44.4 rapid        # EVPN VPWS
PING 172.16.44.4 56 data bytes
.....
---- 172.16.44.4 PING Statistics ----
5 packets transmitted, 0 packets received, 100% packet loss

*A:ASBR-2# ping router-instance "CE-442" 2001:db8::16:44:4 rapid  # EVPN VPWS
PING 2001:db8::16:44:4 56 data bytes
.....
---- 2001:db8::16:44:4 PING Statistics ----
5 packets transmitted, 5 packets bounced, 0 packets received, 100% packet loss

*A:ASBR-2# ping router-instance "CE-52" 172.16.5.54 rapid        # EVPN R-VPLS
PING 172.16.5.54 56 data bytes
```

```
.....
---- 172.16.5.54 PING Statistics ----
5 packets transmitted, 0 packets received, 100% packet loss

*A:ASBR-2# ping router-instance "CE-52" 2001:db8::16:5:54 rapid    # EVPN R-VPLS
PING 2001:db8::16:5:54 56 data bytes

---- 2001:db8::16:5:54 PING Statistics ----
5 packets transmitted, 5 packets bounced, 0 packets received, 100% packet loss
```

The datapath allows traffic between PE-1 and PE-4 for all services, but drops the traffic to and from the local EVPN VPLS and EVPN VPWS on the ASBR. BGP neighbor trust is not supported for EVPN-IFF routes on a PE-ASBR.

Conclusion

BGP neighbor trust prevents label spoofing in inter-AS option B for the VPN-IPv4, VPN-IPv6, and EVPN address families.

Network Group Encryption Helper

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

Topics in this chapter include:

- [Applicability](#)
- [Overview](#)
- [Configuration](#)
- [Conclusion](#)

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 21: General architecture using an NGE helper](#) shows the general architecture using an NGE helper.

Figure 21: General architecture using an NGE helper



37325

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

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:

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

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, as shown in the [VPRN or NG-MVPN using MP-BGP](#) section
- 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

Group encryption configuration

In this example, the following two encryption keygroups are configured manually on NGE-1:

```
# on NGE-1:
configure
  group-encryption
    group-encryption-label 100
    encryption-keygroup 1 create
      keygroup-name "KG1"
      security-association spi 1 authentication-key 0x1111111100000000
        111111110000000001111111000000000111111100000000 encryption-key
        0x11111111000000000111111100000000
      security-association spi 2 authentication-key 0x2222222200000000
        22222222000000000222222200000000022222200000000 encryption-key
        0x22222222000000000222222200000000
      security-association spi 3 authentication-key 0x3333333300000000
        333333330000000003333333000000000333333300000000 encryption-key
```

```

0x33333333000000000333333300000000
security-association spi 4 authentication-key 0x4444444400000000
444444440000000004444444000000004444444400000000 encryption-key
0x44444444000000000444444400000000
active-outbound-sa 1
exit
encryption-keygroup 2 create
keygroup-name "KG2"
security-association spi 5 authentication-key 0x5555555500000000
555555550000000005555555000000005555555500000000 encryption-key
0x55555555000000000555555500000000
security-association spi 6 authentication-key 0x6666666600000000
666666660000000006666666000000006666666600000000 encryption-key
0x66666666000000000666666600000000
security-association spi 7 authentication-key 0x7777777700000000
777777770000000007777777000000007777777700000000 encryption-key
0x77777777000000000777777700000000
security-association spi 8 authentication-key 0x8888888800000000
888888880000000008888888000000008888888800000000 encryption-key
0x88888888000000000888888800000000
active-outbound-sa 5
exit

```

In this example, the authentication key and the encryption key are entered 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.

SDP configuration

On NGE-1, LDP SDP 1 is configured with encryption keygroup 1 and RSVP SDP 3 is configured with encryption keygroup 2:

```

# on NGE-1:
configure
service
sdp 1 mpls create
description "LDP SDP with NGE"
far-end 192.0.2.5
ldp
keep-alive
shutdown
exit
encryption-keygroup 1 direction inbound
encryption-keygroup 1 direction outbound
no shutdown
exit
sdp 3 mpls create
description "RSVP SDP with NGE"
far-end 192.0.2.5
lsp "LSP-NGE-1-NGE-2"
keep-alive
shutdown
exit
encryption-keygroup 2 direction inbound
encryption-keygroup 2 direction outbound
no shutdown
exit

```

PW-template configuration

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

```
# on NGE-1:
configure
service
  pw-template 2 name "2" auto-gre-sdp create
  description "PW template with NGE"
  vc-type vlan
  split-horizon-group "SHG"
  exit
  encryption-keygroup 1 direction inbound
  encryption-keygroup 1 direction outbound
exit
```


BGP configuration

BGP must be enabled on the router and the NGE helper 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)
- MP-BGP VPRN with auto-bind LDP or RSVP-TE (where NGE is configured under the **vprn** context)
- NG-MVPN with MLDP tunnels (where NGE is configured under the **vprn** context)

[Figure 22: BGP topology for learning BGP label routes](#) shows the BGP topology for learning BGP label routes for these services.

Figure 22: BGP topology for learning BGP label routes



37326

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

```
# on PE-1:
configure
router Base
  bgp
    rapid-withdrawal
    group "core-RR"
      family vpn-ipv4 l2-vpn mvpn-ipv4
      peer-as 64496
      neighbor 192.0.2.3      # RR
    exit
  exit
  group "PE-1-NGE-1-RR"
    family vpn-ipv4 l2-vpn mvpn-ipv4
    cluster 192.0.2.1
    peer-as 64496
    neighbor 192.0.2.4      # NGE-1
  exit
  exit
  no shutdown
exit
```

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

```
# on PE-2:
configure
router Base
  bgp
    rapid-withdrawal
    group "core-RR"
      family vpn-ipv4 l2-vpn mvpn-ipv4
      peer-as 64496
      neighbor 192.0.2.3      # RR
    exit
  exit
  group "PE-2-NGE-2-RR"
    family vpn-ipv4 l2-vpn mvpn-ipv4
    cluster 192.0.2.2
    peer-as 64496
    neighbor 192.0.2.5      # NGE-2
  exit
  exit
  no shutdown
exit
```

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

```
# on NGE-1:
configure
router Base
  bgp
    rapid-withdrawal
    group "RR-PE-1"
      family vpn-ipv4 l2-vpn mvpn-ipv4
      peer-as 64496
      neighbor 192.0.2.1      # PE-1
    exit
  exit
  no shutdown
exit
```

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

```
# on NGE-2:
configure
  router Base
    bgp
      rapid-withdrawal
      group "RR-PE-2"
        family vpn-ipv4 l2-vpn mvpn-ipv4
        peer-as 64496
        neighbor 192.0.2.2    # PE-2
      exit
    exit
  no shutdown
exit
```

Operators can enable PE-CE control plane functionality such as EBGp from the NGE helper to learn routes from the CE and advertise them within the VPRN. The optional configuration required for PE-CE functionality is included in this chapter.


Services configuration

VPRN or NG-MVPN using MP-BGP

For these services, NGE is configured under the **vpn** context.

[Figure 23: Operation of NGE helper for MP-BGP auto-bind VPRN or NG-MVPN multicast](#) shows the operation of the NGE helper for MP-BGP auto-bind VPRN-based services or NG-MVPN multicast services.

Figure 23: Operation of NGE helper for MP-BGP auto-bind VPRN or NG-MVPN multicast



37327

VPRN SAPs are typically configured on the router; however, in this case the VPRN and VPRN SAP are configured on the NGE helper. On PE-1, a local Epipe is configured that originates from the customer facing SAP1 and terminates on SAP-A1, connected to the access port on the NGE-1 helper. Traffic on this access port is not encrypted. In this example, Epipe 100301 is configured on PE-1 as follows:

```
# on PE-1:
configure
service
  epipe 100301 name "Epipe-100301" customer 1 create
  sap lag-1:301 create
  description "toward NGE-1 VPRN 301"
  no shutdown
exit
sap lag-11:301.1 create
description "toward CE"
no shutdown
exit
no shutdown
```

```
exit
```

In the VPRN on the NGE-1 helper, the traffic is encrypted. Traffic on the network port is encrypted.

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

- optional RSVP-TE tunnels with fast reroute (FRR) to other remote PEs
 - If RSVP-TE tunnels are configured, then T-LDP sessions with tunneling enabled must also be configured to these same PEs. These sessions allow LDP packets from the NGE helper to use LDP to hop onto RSVP-TE tunnels.
- optional LDP, including MLDP, tunnels on core network interfaces for unicast and multicast traffic to other PEs
- BGP sessions for the VPN-IPv4 and MVPN-IPv4 address families, as described in the [BGP configuration](#) section
- LDP, including MLDP, is configured on the network interface to the NGE helper

On the NGE-1 helper, configuration is minimal and includes:

- VPRN SAPN1 where, optionally, PE-CE IGP protocols can be configured to learn routes from CE-1
- VPRN NG-MVPN for multicast services
- LDP, including MLDP, on the network interface to PE-1
- BGP session for the VPN-IPv4 and MVPN-IPv4 address families, as described in the [BGP configuration](#) section
- NGE enabled on the VPRN for encrypting unicast and multicast services

In this example, the configuration of VPRN 301 on NGE-1 is as follows:

```
# on NGE-1:
configure
  service
    vprn 301 name "VPRN-301" customer 1 create
      description "MP-BGP, NG MVPN, auto-bind LDP, VPRN NGE"
      autonomous-system 64501
      interface "toCE-1" create
        address 172.16.11.2/24
        sap lag-1:301 create
      exit
    exit
  bgp-ipvpn
    mpls
      auto-bind-tunnel
      resolution-filter
      ldp
      exit
      resolution filter
    exit
    route-distinguisher 301:1
    vrf-target target:301:1
    no shutdown
  exit
exit
bgp
  group "CE"
    export "exportBGP"
    neighbor 172.16.11.1
    family ipv4
```

```

        type external
        peer-as 64502
    exit
    exit
    no shutdown
exit
pim
    interface "toCE-1"
    exit
    rp
        static
        exit
        bsr-candidate
        shutdown
        exit
        rp-candidate
        shutdown
        exit
    exit
    no shutdown
exit
mvpn
    auto-discovery default # default auto-discovery via BGP
    c-mcast-signaling bgp
    provider-tunnel
        inclusive
        mldp
        no shutdown
        exit
    exit
    exit
    vrf-target unicast
    exit
exit
encryption-keygroup 1 direction inbound
encryption-keygroup 1 direction outbound
no shutdown
exit

```

T-LDP signaled Epipe or VPLS services

For these services, NGE is configured under the **sdp** context. On NGE-1, LDP SDP 1 is configured with encryption keygroup 1 and RSVP SDP 3 is configured with encryption keygroup 2, as follows:

```


# on NGE-1:
configure
    service
        sdp 1 mpls create
        description "LDP SDP with NGE"
        far-end 192.0.2.5
        ldp
        keep-alive
        shutdown
        exit
        encryption-keygroup 1 direction inbound
        encryption-keygroup 1 direction outbound
        no shutdown
    exit
    sdp 3 mpls create
    description "RSVP SDP with NGE"
    far-end 192.0.2.5

```

```
lsp "LSP-NGE-1-NGE-2"
keep-alive
shutdown
exit
encryption-keygroup 2 direction inbound
encryption-keygroup 2 direction outbound
no shutdown
exit
```

Figure 24: NGE helper for T-LDP signaled Epipe or VPLS services shows the operation of the NGE helper for T-LDP signaled Epipe or VPLS services.

Figure 24: NGE helper for T-LDP signaled Epipe or VPLS services



37328

Similar to the VPRN scenario, the service SAPN1 of the Epipe or VPLS is configured on the NGE helper. On PE-1, a local Epipe is configured that is originating from the customer facing SAP1 and terminating on SAP-A1 connected to the NGE-1 helper on the access port where SAPN1 is configured. For example,

Epipe 100401 toward Epipe 101 on NGE-1 is configured as follows. Similar Epipes are configured toward other services on NGE-1, such as VPLS 501 and VPLS 601.

```
# on PE-1:
configure
service
    epipe 100401 name "Epipe-100401" customer 1 create
        sap lag-1:401 create
            description "toward NGE-1 Epipe 401"
            no shutdown
        exit
        sap lag-11:401.1 create
            description "toward CE"
            no shutdown
        exit
    no shutdown
exit
```

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

- optional RSVP-TE tunnels with FRR to other remote PEs
 - If RSVP-TE tunnels are configured, then T-LDP sessions with tunneling enabled are also configured to these same PEs. These sessions allow LDP packets from the NGE-1 helper to use LDP to hop onto RSVP-TE tunnels.
- optional LDP tunnels if RSVP-TE tunnels are not used
- LDP on each network interface to the NGE-1 helper

On the NGE-1 helper, the configuration is minimal and includes:

- Epipe or VPLS SAPN1 configured on the NGE helper
- T-LDP configured from the NGE helper to each remote PE that needs to participate in the Epipe or VPLS service
- SDPs configured on the NGE helper toward each PE that is participating in the Epipe or VPLS service
- LDP configured on the network interface
- NGE enabled on the SDPs for encrypting the Epipe or VPLS services using the SDPs

Epipe 401 is configured with LDP SDP 1, which uses encryption keygroup 1:

```
# on NGE-1:
configure
service
    epipe 401 name "Epipe-401" customer 1 create
        description "Epipe, LDP SDP, SDP NGE"
        sap lag-1:401 create
            no shutdown
        exit
        spoke-sdp 1:401 create
            no shutdown
        exit
    no shutdown
exit
```

Likewise, VPLS 501 is configured with LDP SDP 1, which uses encryption keygroup 1:

```
# on NGE-1:
```

```
configure
service
  vpls 501 name "VPLS-501" customer 1 create
  description "VPLS, LDP SDP, SDP NGE"
  sap lag-1:501 create
  no shutdown
  exit
  spoke-sdp 1:501 create
  no shutdown
  exit
  no shutdown
exit
```

VPLS 601 is configured with RSVP SDP 3, which uses encryption keygroup 2:

```
# on NGE-1:
configure
service
  vpls 601 name "VPLS-601" customer 1 create
  description "VPLS, RSVP SDP, SDP NGE"
  sap lag-1:601 create
  no shutdown
  exit
  mesh-sdp 3:601 create
  no shutdown
  exit
  no shutdown
exit
```


BGP VPLS or BGP VPWS with auto-GRE SDP

For these services, NGE is configured under the **pw-template** context, as in the following example:

```
# on NGE-1:
configure
service
  pw-template 2 name "2" auto-gre-sdp create
  description "PW template with NGE"
  vc-type vlan
  split-horizon-group "SHG"
  exit
  encryption-keygroup 1 direction inbound
  encryption-keygroup 1 direction outbound
exit
```

Figure 25: 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 25: NGE helper for BGP VPLS or BGP VPWS using GRE SDPs with auto-GRE SDP



37329

Similar to the VPRN 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. The configuration is similar to the preceding configuration of Epipe 100401 on PE-1.

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
    epipe 101 name "Epipe-101" customer 1 create
      description "BGP VPWS auto-gre SDP_PW template 2"
      bgp
        route-distinguisher 101:1
        route-target export target:101:1 import target:101:1
        pw-template-binding 2
      exit
    exit
  bgp-vpws
    ve-name "pe-1"
    ve-id 1
    exit
    remote-ve-name "pe-2"
    ve-id 2
    exit
    no shutdown
  exit
  sap lag-1:101 create
  exit
  no shutdown
exit
```

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 201 name "VPLS-201" customer 1 create
      description "BGP VPLS auto-gre SDP_PW template 2"
      bgp
        route-distinguisher 201:1
        route-target export target:201:1 import target:201:1
        pw-template-binding 2
      exit
    exit
  bgp-vpls
    max-ve-id 10
    ve-name "pe-1"
    ve-id 1
    exit
    no shutdown
  exit
  sap lag-1:201 create
  no shutdown
  exit
  no shutdown
exit
```

Configuration overview

Configuration on NGE-1 helper

On the NGE-1 helper, the configuration of the control plane and services for all preceding services is as follows:

```
#-----
echo "Card Configuration"
#-----
card 1
  card-type iom-v
  mda 1
    mda-type m20-v
    no shutdown
  exit
  mda 2
    mda-type m20-v
    no shutdown
  exit
  mda 3
    mda-type m20-v
    no shutdown
  exit
  mda 4
    mda-type m20-v
    no shutdown
  exit
  no shutdown
exit
#-----
echo "Port Configuration"
#-----
port 1/1/1
  ethernet
    mode hybrid
    encap-type dot1q
  exit
  no shutdown
exit
port 1/1/2
  ethernet
    mode hybrid
    encap-type dot1q
  exit
  no shutdown
exit
---snip---
#-----
echo "LAG Configuration"
#-----
lag 1
  description "LAG to PE-1"
  mode hybrid
  encap-type dot1q
  port 1/1/1
  port 1/1/2
  lacp active administrative-key 32768
  no shutdown
exit
```

```
#-----
echo "Group Encryption Configuration"
#-----
group-encryption
  group-encryption-label 100
  encryption-keygroup 1 create
    keygroup-name "KG1"
    security-association spi 1 authentication-key 0x4669dcf53c34b8138a27
      09022ee24a9b342777047ddfa833e43a5ff9917cde901a6f76bc0cc01cb363a3a77
      9916aa0b8 encryption-key 0x5e172b1138812340ddcdc604ea3f4214bbf7d564
      56cabbab018006d6ac92bc8f crypto
    security-association spi 2 authentication-key 0x731da9633f8496f52a5e
      f240f674b4122cdea4460a24968f8591e4ba0cc713f272b2eeee6b260cb791eedf4
      77f24ad7a encryption-key 0xe7e24975f3168fdaa9f57fcb248d2948cf8154a3
      915a004b261f4b4850b38e1e crypto
    security-association spi 3 authentication-key 0x6c9ab2e6ff1cfa69daef
      d2e2d8107dc96ec5ebf49eb6cb2c75a4f0d7a122e31dd728b9ddc97e4afc31f2c97
      1cfacea34 encryption-key 0x70590aacb24913a3f04afa38ecb929fc9c6f32da
      d6d4f18e891a883b08d8f806 crypto
    security-association spi 4 authentication-key 0x90c67c848bdb9b7ac0c1
      2e42390da7ea7de09002e84af569222072f6dd88a6f8e8d461c04cb044fc1d3df69
      97090d5a5 encryption-key 0x7cc12d7118409173905478f639d623e689e6f313
      7baf91abdcc843725d4d14c6 crypto
    active-outbound-sa 1
  exit
  encryption-keygroup 2 create
    keygroup-name "KG2"
    security-association spi 5 authentication-key 0xae8e620a56288524d2cd
      210b09fad464a3214ce3ce7e79422b385e44cc896acbf933f7ac73cd2c5fa4a683
      a3db75d4d encryption-key 0x97e6dee7ad9ecb03b9e726b1291f9aca88d06200
      bb8218fe0bf378f3b682a3a0 crypto
    security-association spi 6 authentication-key 0xe62e5f59e416bbf27352
      a676dd21b3c7da08a126fb373c8cb7e5ec4f8b95e70f8a99cbd177f2537d4a48a42
      44aebf2e8 encryption-key 0x42d4424316861834a9e8a94688521a623b580c7b
      730d8c37aa825a0d92e9bb80 crypto
    security-association spi 7 authentication-key 0xa4b7d14a16d2e93187c0
      0eb8704001aa588e6b56927bd7a9791878da78ca6c8d7bc35d62b8de0f077451874
      9b257db96 encryption-key 0x7e315a24e9e1f58abbab02ace4fd9099932416e3
      8021c9204866327b580118b0 crypto
    security-association spi 8 authentication-key 0x6a1e474cf8bd552cbb28
      805e22962ddfl1e0e13b478e74be0cabf81c4ea2903a4834d1c64e2aae60e199fac5
      a0c21f6fa encryption-key 0xd7082b7c5d7a7a2f7d139f8dcc9a3921422aab10
      01acb18346e2c63b3b9db7b8 crypto
    active-outbound-sa 5
  exit
exit
exit
#-----
echo "Router (Network Side) Configuration"
#-----
router Base
  interface "int-NGE-1-PE-1"
    address 192.168.14.2/30
    port lag-1:1000
    no shutdown
  exit
  interface "system"
    address 192.0.2.4/32
    no shutdown
  exit
  autonomous-system 64496
  router-id 192.0.2.4
#-----
echo "OSPFv2 Configuration"
#-----
```

```
ospf 0
  asbr
  traffic-engineering
  timers
    lsa-arrival 200
    lsa-generate 5000 lsa-initial-wait 200 lsa-second-wait 1000
    spf-wait 1000 spf-initial-wait 10 spf-second-wait 500
  exit
  disable-ldp-sync
  area 0.0.0.0
    interface "system"
      no shutdown
    exit
    interface "int-NGE-1-PE-1"
      interface-type point-to-point
      no advertise-subnet
      hello-interval 1
      dead-interval 4
      no shutdown
    exit
  exit
  no shutdown
exit
#-----
echo "PIM Configuration"
#-----
pim
  interface "system"
  exit
  interface "int-NGE-1-PE-1"
  exit
  rp
    static
    exit
    bsr-candidate
      shutdown
    exit
    rp-candidate
      shutdown
    exit
  exit
  no shutdown
exit
#-----
echo "MPLS Configuration"
#-----
mpls
  interface "system"
    no shutdown
  exit
  interface "int-NGE-1-PE-1"
    no shutdown
  exit
exit
#-----
echo "RSVP Configuration"
#-----
rsvp
  interface "system"
    no shutdown
  exit
  interface "int-NGE-1-PE-1"
    no shutdown
  exit
```

```
        no shutdown
    exit
#-----
echo "MPLS LSP Configuration"
#-----
    mpls
        path "path-NGE-1-NGE-2"
        no shutdown
    exit
    lsp "LSP-NGE-1-NGE-2"
        to 192.0.2.5
        primary "path-NGE-1-NGE-2"
    exit
    no shutdown
exit
no shutdown
exit
#-----
echo "LDP Configuration"
#-----
    ldp
        import-pmsi-routes
    exit
    tcp-session-parameters
    exit
    interface-parameters
        interface "int-NGE-1-PE-1" dual-stack
        ipv4
            no shutdown
        exit
        no shutdown
    exit
    exit
    targeted-session
        peer 192.0.2.5
        no shutdown
    exit
    exit
    no shutdown
exit
exit
#-----
echo "Service Configuration"
#-----
    service
        sdp 1 mpls create
        description "LDP SDP with NGE"
        far-end 192.0.2.5
        ldp
        keep-alive
        shutdown
    exit
    encryption-keygroup 1 direction inbound
    encryption-keygroup 1 direction outbound
    no shutdown
exit
    sdp 3 mpls create
    description "RSVP SDP with NGE"
    far-end 192.0.2.5
    lsp "LSP-NGE-1-NGE-2"
    keep-alive
    shutdown
exit
```

```
        encryption-keygroup 2 direction inbound
        encryption-keygroup 2 direction outbound
        no shutdown
    exit
    customer 1 name "1" create
        description "Default customer"
    exit
    pw-template 2 name "2" auto-gre-sdp create
        vc-type vlan
        split-horizon-group "SHG"
    exit
        encryption-keygroup 1 direction inbound
        encryption-keygroup 1 direction outbound
    exit
    vprn 301 name "VPRN-301" customer 1 create
        interface "toCE-1" create
    exit
    exit
    epipe 101 name "Epipe-101" customer 1 create
        description "BGP VPWS auto-gre SDP_PW template 2"
        bgp
            route-distinguisher 101:1
            route-target export target:101:1 import target:101:1
            pw-template-binding 2
        exit
    exit
    bgp-vpws
        ve-name "pe-1"
        ve-id 1
    exit
        remote-ve-name "pe-2"
        ve-id 2
    exit
        no shutdown
    exit
    sap lag-1:101 create
        no shutdown
    exit
    no shutdown
    exit
    vpls 201 name "VPLS-201" customer 1 create
        description "BGP VPLS auto-gre SDP_PW template 2"
        bgp
            route-distinguisher 201:1
            route-target export target:201:1 import target:201:1
            pw-template-binding 2
        exit
    exit
    bgp-vpls
        max-ve-id 10
        ve-name "pe-1"
        ve-id 1
    exit
        no shutdown
    exit
    stp
        shutdown
    exit
    sap lag-1:201 create
        no shutdown
    exit
    no shutdown
    exit
    vprn 301 name "VPRN-301" customer 1 create
```

```
description "MP-BGP, NG MVPN, auto-bind LDP, VPRN NGE"
autonomous-system 64501
interface "toCE-1" create
    address 172.16.11.2/24
    sap lag-1:301 create
    exit
exit
bgp-ipvpn
    mpls
        auto-bind-tunnel
        resolution-filter
        ldp
        exit
        resolution filter
    exit
    route-distinguisher 301:1
    vrf-target target:301:1
    no shutdown
exit
exit
bgp
    group "CE"
        export "exportBGP"
        neighbor 172.16.11.1
            family ipv4
            type external
            peer-as 64502
        exit
    exit
    no shutdown
exit
pim
    interface "toCE-1"
    exit
    rp
        static
        exit
        bsr-candidate
        shutdown
        exit
        rp-candidate
        shutdown
        exit
    exit
    no shutdown
exit
mvpn
    auto-discovery default
    c-mcast-signaling bgp
    provider-tunnel
        inclusive
        mldp
        no shutdown
        exit
    exit
    exit
    vrf-target unicast
    exit
exit
encryption-keygroup 1 direction inbound
encryption-keygroup 1 direction outbound
no shutdown
exit
epipe 401 name "Epipe-401" customer 1 create
```

```
description "Epipe, LDP SDP, SDP NGE"
sap lag-1:401 create
no shutdown
exit
spoke-sdp 1:401 create
no shutdown
exit
no shutdown
exit
vpls 501 name "VPLS-501" customer 1 create
description "VPLS, LDP SDP, SDP NGE"
stp
    shutdown
exit
sap lag-1:501 create
no shutdown
exit
spoke-sdp 1:501 create
no shutdown
exit
no shutdown
exit
vpls 601 name "VPLS-601" customer 1 create
description "VPLS, RSVP SDP, SDP NGE"
stp
    shutdown
exit
sap lag-1:601 create
no shutdown
exit
mesh-sdp 3:601 create
no shutdown
exit
no shutdown
exit
exit
#-----
---snip---
#-----
echo "Policy Configuration"
#-----
    policy-options
    begin
    policy-statement "exportBGP"
        entry 10
            from
                protocol bgp-vpn
            exit
            action accept
            exit
        exit
    exit
    commit
    exit
#-----
echo "BGP Configuration"
#-----
    bgp
        rapid-withdrawal
        group "RR-PE-1"
            family vpn-ipv4 l2-vpn mvpn-ipv4
            peer-as 64496
            neighbor 192.0.2.1
            exit
```

```
        exit
    no shutdown
    exit
exit
#-----
```

Configuration on PE-1

The configuration on PE-1 is as follows:

```
---snip---
#-----
echo "LAG Configuration"
#-----
    lag 1
        description "LAG to NGE-1"
        mode hybrid
        encap-type dot1q
        port 1/1/c1/3
        port 1/1/c1/4
        lacp passive administrative-key 1
        no shutdown
    exit
    lag 11
        description "LAG to CE-1_access"
        mode access
        encap-type qinq
        port 1/1/c2/1
        port 1/1/c2/2
        lacp passive administrative-key 11
        no shutdown
    exit
    lag 12
        description "LAG to core"
        mode hybrid
        encap-type dot1q
        port 1/1/c1/1
        port 1/1/c1/2
        lacp active administrative-key 12
        no shutdown
    exit
---snip---
#-----
echo "Router (Network Side) Configuration"
#-----
    router Base
        interface "int-PE-1-NGE-1"
            address 192.168.14.1/30
            port lag-1:1000
            no shutdown
        exit
        interface "int-PE-1-core"
            address 192.168.12.1/30
            port lag-12:1000
            no shutdown
        exit
        interface "system"
            address 192.0.2.1/32
            no shutdown
    exit
```

```

autonomous-system 64496
router-id 192.0.2.1
#-----
echo "OSPFv2 Configuration"
#-----
    ospf 0
    asbr
    traffic-engineering
    ldp-over-rsvp      # only if LDPoRSVP is used in the core
    area 0.0.0.0
        interface "system"
            no shutdown
        exit
        interface "int-PE-1-core"
            interface-type point-to-point
            no advertise-subnet
            hello-interval 1
            dead-interval 4
            authentication-type message-digest
            message-digest-key 10 md5 "qBlAj0UBDKLgnvWaw9ifX+l6Nfo=" hash2
            no shutdown
        exit
        interface "int-PE-1-NGE-1"
            interface-type point-to-point
            no advertise-subnet
            hello-interval 1
            dead-interval 4
            no shutdown
        exit
    exit
    no shutdown
exit
#-----
echo "PIM Configuration"
#-----
    pim
        interface "system"
        exit
        interface "int-PE-1-core"
        exit
        interface "int-PE-1-NGE-1"
        exit
        rp
            static
            exit
            bsr-candidate
                shutdown
            exit
            rp-candidate
                shutdown
            exit
        exit
    exit
    no shutdown
exit
#-----
echo "MPLS Configuration"
#-----
    mpls
        interface "system"
            no shutdown
        exit
        interface "int-PE-1-core"
            no shutdown
        exit

```

```
        interface "int-PE-1-NGE-1"
            no shutdown
        exit
    exit
#-----
echo "RSVP Configuration"
#-----
    rsvp
        interface "system"
            no shutdown
        exit
        interface "int-PE-1-core"
            no shutdown
        exit
        interface "int-PE-1-NGE-1"
            no shutdown
        exit
        no shutdown
    exit
#-----
echo "MPLS LSP Configuration"
#-----
    mpls
        path "path-PE-1-PE-2"      # only if LDPoRSVP is used in the core
            no shutdown
        exit
        lsp "LSP-PE-1-PE-2"      # only if LDPoRSVP is used in the core
            to 192.0.2.2
            primary "path-PE-1-PE-2"
        exit
        no shutdown
    exit
    no shutdown
exit
#-----
echo "LDP Configuration"
#-----
    ldp
        prefer-mcast-tunnel-in-tunnel
        import-pmsi-routes
        exit
        tcp-session-parameters
        exit
        interface-parameters
            interface "int-PE-1-core" dual-stack
                ipv4
                    no shutdown
                exit
                no shutdown
            exit
            interface "int-PE-1-NGE-1" dual-stack
                ipv4
                    transport-address system
                    no shutdown
                exit
                no shutdown
            exit
        exit
        exit
        targeted-session
            peer 192.0.2.2      # only if LDPoRSVP is used in the core
            tunneling
                lsp "LSP-PE-1-PE-2"
            exit
            no shutdown
```

```
        exit
    exit
    no shutdown
exit
exit

#-----
echo "Service Configuration"
#-----

service
    customer 1 name "1" create
        multi-service-site "bras" create
        exit
        description "Default customer"
    exit
    epipe 100101 name "Epipe-100101" customer 1 create
        sap lag-1:101 create
            description "toward NGE-1 Epipe 101"
            no shutdown
        exit
        sap lag-11:101.1 create
            description "toward CE"
            no shutdown
        exit
        no shutdown
    exit
    epipe 100201 name "Epipe-100201" customer 1 create
        sap lag-1:201 create
            description "toward NGE-1 VPLS 201"
            no shutdown
        exit
        sap lag-11:201.1 create
            description "toward CE"
            no shutdown
        exit
        no shutdown
    exit
    epipe 100301 name "Epipe-100301" customer 1 create
        sap lag-1:301 create
            description "toward NGE-1 VPRN 301"
            no shutdown
        exit
        sap lag-11:301.1 create
            description "toward CE"
            no shutdown
        exit
        no shutdown
    exit
    epipe 100401 name "Epipe-100401" customer 1 create
        sap lag-1:401 create
            description "toward NGE-1 Epipe 401"
            no shutdown
        exit
        sap lag-11:401.1 create
            description "toward CE"
            no shutdown
        exit
        no shutdown
    exit
    epipe 100501 name "Epipe-100501" customer 1 create
        sap lag-1:501 create
            description "toward NGE-1 VPLS 501"
            no shutdown
        exit
```

```

        sap lag-11:501.1 create
            description "toward CE"
            no shutdown
        exit
    no shutdown
exit
epipe 100601 name "Epipe-100601" customer 1 create
    sap lag-1:601 create
        description "toward NGE-1 VPLS 601"
        no shutdown
    exit
    sap lag-11:601.1 create
        description "toward CE"
        no shutdown
    exit
    no shutdown
exit
exit
exit
#-----
---snip---
#-----
echo "BGP Configuration"
#-----
    bgp
        rapid-withdrawal
        group "core-RR"
            family vpn-ipv4 l2-vpn mvpn-ipv4
            peer-as 64496
            neighbor 192.0.2.3
        exit
    exit
    group "PE-1-NGE-1-RR"
        family vpn-ipv4 l2-vpn mvpn-ipv4
        cluster 192.0.2.1
        peer-as 64496
        neighbor 192.0.2.4
    exit
    no shutdown
exit
exit
#-----
---snip---

```

The Epipes are the connections between the CE and the NGE helper for each service.

Verification

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

```

*A:NGE-1# show service id 101 base

=====
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: 03/29/2023 07:23:33
Last Mgmt Change : 03/29/2023 07:23:33
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: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: 03/29/2023 07:21:39
Last Mgmt Change : 03/29/2023 07:23:33
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

*A:NGE-1# show service id 301 base

Service Basic Information

```
Service Id       : 301                               Vpn Id          : 0
Service Type     : VPRN
MACSec enabled   : no
Name             : VPRN-301
Description      : MP-BGP, NG MVPN, auto-bind LDP, VPRN NGE
Customer Id      : 1                               Creation Origin : manual
Last Status Change: 03/29/2023 07:21:39
Last Mgmt Change : 03/29/2023 07:21:39
Admin State      : Up                               Oper State       : Up
---snip---
```

SAP Count : 1 SDP Bind Count : 0

```

-----
Service Access & Destination Points
-----
Identifier                Type            AdmMTU  OprMTU  Adm  Opr
-----
sap:lag-1:301             q-tag          8936    8936    Up   Up
=====

*A:NGE-1# show service id 401 base

=====
Service Basic Information
=====
Service Id       : 401                Vpn Id          : 0
Service Type     : Epipe
MACSec enabled   : no
Name             : Epipe-401
Description      : Epipe, LDP SDP, SDP NGE
Customer Id      : 1                Creation Origin  : manual
Last Status Change: 03/29/2023 07:22:05
Last Mgmt Change : 03/29/2023 07:21:39
Test Service     : No
Admin State      : Up                Oper State       : Up
MTU              : 1514
Vc Switching     : False
SAP Count        : 1                SDP Bind Count   : 1
---snip---

-----
Service Access & Destination Points
-----
Identifier                Type            AdmMTU  OprMTU  Adm  Opr
-----
sap:lag-1:401             q-tag          8936    8936    Up   Up
sdp:1:401 S(192.0.2.5)    Spok           0        8910    Up   Up
=====

*A:NGE-1# show service id 501 base

=====
Service Basic Information
=====
Service Id       : 501                Vpn Id          : 0
Service Type     : VPLS
MACSec enabled   : no
Name             : VPLS-501
Description      : VPLS, LDP SDP, SDP NGE
Customer Id      : 1                Creation Origin  : manual
Last Status Change: 03/29/2023 07:21:39
Last Mgmt Change : 03/29/2023 07:21:39
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-1:501          q-tag      8936      8936      Up      Up
sdp:1:501 S(192.0.2.5) Spok        0        8910      Up      Up
=====

*A:NGE-1# show service id 601 base

=====
Service Basic Information
=====
Service Id       : 601          Vpn Id       : 0
Service Type     : VPLS
MACSec enabled   : no
Name            : VPLS-601
Description      : VPLS, RSVP SDP, SDP NGE
Customer Id     : 1            Creation Origin : manual
Last Status Change: 03/29/2023 07:21:39
Last Mgmt Change : 03/29/2023 07:21:39
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-1:601          q-tag      8936      8936      Up    Up
sdp:3:601 M(192.0.2.5) Mesh        0        8910      Up    Up
=====
```

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

```
*A:NGE-1# show group-encryption encryption-keygroup 1

=====
Encryption Keygroup Configuration Detail
=====
Keygroup Id       : 1
Keygroup Name     : KG1
Description       : None
Authentication Algo: sha256
Encryption Algo   : aes128
Active Outbound SA : 1
Activation Time   : 03/29/2023 09:14:59

-----
Security Associations
-----
Spi              : 1
Install Time     : 03/29/2023 09:14:59
Key CRC          : 0xf57dcffc

Spi              : 2
Install Time     : 03/29/2023 09:14:59
Key CRC          : 0x26134d07

Spi              : 3
Install Time     : 03/29/2023 09:14:59
```

```

Key CRC           : 0xde19ce91
Spi              : 4
Install Time     : 03/29/2023 09:14:59
Key CRC         : 0x5bbf4eb0
-----
Encryption Keygroup Forwarded Statistics
-----
Encrypted Pkts      : 164          Encrypted Bytes      : 15624
Decrypted Pkts     : 149          Decrypted Bytes     : 14204
-----
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
-----
1                   Inbound   Outbound
32766               Inbound   Outbound
32767               Inbound   Outbound
-----
Inbound Keygroup SDP Association Count: 3
Outbound Keygroup SDP Association Count: 3
-----
VPRN Keygroup Association Table
-----
VPRN SVC ID         Direction
-----
301                 Inbound   Outbound
-----
Inbound Keygroup VPRN Association Count: 1
Outbound Keygroup VPRN Association Count: 1
-----
Network Interface Association Table
-----
No entries found
-----
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.

Seamless BFD Application — Auto-bind tunnel

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

Topics in this chapter include:

- [Applicability](#)
- [Overview](#)
- [Configuration](#)
- [Conclusion](#)

Applicability

This chapter was initially written based on SR OS Release 19.10.R3, but the 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 *7750 SR and 7950 XRS Segment Routing and PCE Advanced Configuration Guide for Classic CLI*.


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 26: 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).

Figure 26: S-BFD session establishment – continuity check



35629


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 *7750 SR and 7950 XRS Segment Routing and PCE Advanced Configuration Guide for Classic CLI*, 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 27: [Example topology](#) shows the example topology. The VPRN and EVPN services will be configured on PE-2 and PE-5.

Figure 27: Example topology



35836

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"
    address 192.168.23.1/30
    port 1/1/c2/1:1000
  exit
  interface "int-PE-2-PE-4"
    address 192.168.24.1/30
    port 1/1/c1/1:1000
  exit
  interface "system"
    address 192.0.2.2/32
  exit
  mpls-labels
    sr-labels start 32000 end 32999
  exit
  isis 0
    area-id 49.0001
```

```

traffic-engineering
advertise-router-capability area
segment-routing
    prefix-sid-range global
    no shutdown
exit
interface "system"
    ipv4-node-sid index 2
exit
interface "int-PE-2-PE-3"
    interface-type point-to-point
exit
interface "int-PE-2-PE-4"
    interface-type point-to-point
exit
no shutdown
exit

```

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"
            discriminator 524291
            local-state up
            no shutdown
        exit
    exit

```

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
    router Base
        bfd
            begin
            bfd-template "bfd-cpm-np-1s"
                type "cpm-np"
                transmit-interval 1000    # minimum value is 10 ms
                receive-interval 1000    # minimum value is 10 ms
            exit
        commit

```



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
    path "empty"
    no shutdown
  exit
  lsp "LSP-PE-2-PE-5_empty_localCSPF" sr-te
    to 192.0.2.5
    path-computation-method local-cspf
    bfd
      bfd-template "bfd-cpm-np-ls"
      bfd-enable
      failure-action failover-or-down
    exit
    primary "empty"
  exit
  no shutdown
exit
no shutdown
```

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: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
=====
```

The SR-TE LSP with tunnel ID 655362 is "LSP-PE-2-PE-5_empty_localCSPF":

```
*A: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 Index   : 65536
From        : 192.0.2.2
To          : 192.0.2.5
Adm State   : Up
---snip---
```

LSP Tunnel ID : 1
TTM Tunnel Id : 655362

Oper State : Up

The S-BFD session for the SR-TE LSP is up, as follows:

```
*A: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 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
  router Base
    autonomous-system 64496
    bgp
      vpn-apply-import
      vpn-apply-export
      rapid-withdrawal
      split-horizon
      rapid-update vpn-ipv4 evpn
      group "internal"
        family vpn-ipv4 evpn
        peer-as 64496
        neighbor 192.0.2.5
      exit
    exit
  exit
exit
service
```

```

vprn 1 name "VPRN-1" customer 1 create
  interface "int-VPRN-1_PE-2_CE-11" create
    address 172.31.2.2/30
    mac 00:00:5e:00:53:11
    sap 1/1/c4/1:1 create
    exit
  exit
  bgp-ipvpn
  mpls
    auto-bind-tunnel
    resolution-filter
    sr-isis
    sr-te
    exit
    resolution filter
  exit
  route-distinguisher 64496:1
  vrf-target target:64496:1
  no shutdown
  exit
exit
no shutdown
exit
vpls 2 name "VPLS-2" customer 1 create
  bgp
  exit
  bgp-evpn
  evi 2
  mpls bgp 1
    auto-bind-tunnel
    resolution-filter
    sr-isis
    sr-te
    exit
    resolution filter
  exit
  no shutdown
  exit
exit
stp
  shutdown
exit
sap 1/1/c3/1:2 create
  no shutdown
exit
no shutdown
exit

```

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: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   00h00m15s    0
int-VPRN-1_PE-2_CE-11      0
172.31.5.4/30               Remote BGP VPN 00h00m09s    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: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 07:41:50
2	00:00:5e:00:53:62	mpls-1: 192.0.2.5:524284	Evpn	07/05/23 07:41:50

```

      sr-te:655362
-----
No. of MAC Entries: 2
-----
Legend: L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
=====

```

```

*A: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)	00h00m33s

```

-----
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: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 28: Primary path of SR-TE LSP via PE-4 shows the primary path of the SR-TE LSP.

Figure 28: Primary path of SR-TE LSP via PE-4



35837


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
  shutdown
exit
```

Figure 29: 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.

Figure 29: Remote failure in the primary path of the SR-TE LSP



35838

The S-BFD session goes operationally down, as follows:

```
*A: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      : 51200
Fec Type         : srTe
Oper State       : Down
Last Up Time     : 0d 00:04:45
Down Time        : 0d 00:00:01
Protocols        : mplsLsp
Up Transitions   : 1
Down Transitions : 1
Version Mismatch : 0

Forwarding Information

Local Discr      : 1
Local Diag       : 1 (Detect time expired)
Local Mode       : Demand
Local Min Tx     : 1000
Last Sent (ms)   : 0
Type             : cpm-np
Remote           : Unheard
Local State      : Down
Local Mult       : 3
Local Min Rx     : 0
Remote Discr     : 524291
=====
```

When the S-BFD session goes down, the SR-TE LSP goes operationally down, as follows:

```
*A:PE-2# show router mpls sr-te-lsp

=====
MPLS SR-TE LSPs (Originating)
=====
```

LSP Name To	Tun Id	Protect Path	Adm	Opr
LSP-PE-2-PE-5_empty_localCSPF 192.0.2.5	1	N/A	Up	Dwn
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: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: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  00h01m43s 0
int-VPRN-1_PE-2_CE-11  0
172.31.5.4/30           Remote BGP VPN 00h00m13s 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: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 07:41:50
2          00:00:5e:00:53:62 mpls-1:                  Evpn    07/05/23 07:41:50
                        192.0.2.5:524284
isis:524293

-----
No. of MAC Entries: 2
-----
Legend:  L=Learned O=0am 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: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 30: SR-TE LSP reconnects after retry timer expires show that the primary path of the SR-TE tunnel goes via PE-3.

Figure 30: SR-TE LSP reconnects after retry timer expires



35839

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:PE-2# show router tunnel-table 192.0.2.5/32

=====
IPv4 Tunnel Table (Router: Base)
=====
```

Destination Color	Owner	Encap	TunnelId	Pref	Nexthop	Metric
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.

Customer document and product support



Customer documentation

[Customer documentation welcome page](#)



Technical support

[Product support portal](#)



Documentation feedback

[Customer documentation feedback](#)