



7450 Ethernet Service Switch 7750 Service Router Virtualized Service Router

Releases up to 24.10.R2

Triple Play Service Delivery Architecture Advanced Configuration Guide for Classic CLI

3HE 20810 AAAC TQZZA
Edition: 01
March 2025

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.

© 2025 Nokia.

Table of contents

List of tables..... 6

List of figures.....8

Preface..... 17

ARP Hosts..... 18

Bridged CO.....38

DHCP Server Failover States..... 69

DHCPv4 Server Basics..... 93

Diameter Application NASREQ..... 124

Diameter Base Protocol: Establishing a Diameter Peer Connection..... 137

ESM Basics..... 150

ESM 128-bit Mode for DHCPv6 IA_NA WAN Hosts..... 180

ESM IPv4: Multicast in a Wholesale/Retail Scenario..... 185

ESM IPv4: Multicast with Redirection..... 198

ESM IPv4: Multicast with SRRP.....229

ESM SLAAC Prefix Assignment via Local Address Server..... 257

ESMv4: PPPoE Hosts.....275

ESMv6: IPoE Dual Stack Hosts..... 316

ESMv6: PPPoE Dual Stack Hosts.....350

Flexible Authentication Model in ESM.....	377
GTP Access.....	420
High Scale QoS IOM in ESM Context: Expanded SLA Mode.....	445
High Scale QoS IOM in ESM Context: Single SLA Mode.....	479
Ingress Multicast Path Management.....	512
IPoE Sessions.....	545
IPv4 DHCP Hosts.....	570
L2TP for Subscriber Access — LAC.....	606
Local User Database Basics.....	654
Local User Database for DHCPv4 Server.....	682
Local User Database for Enhanced Subscriber Management.....	699
Managed SAPs with Routed CO.....	728
Multi-Chassis Ring Layer 2 with Enhanced Subscriber Management.....	750
Python Cache Support for ESM Applications.....	774
RADIUS-Triggered Dynamic Data Service Provisioning.....	790
Raw Formatting of DHCPv4/v6 Options in ESM.....	825
Routed CO.....	852
Subscriber Redundancy for Routed CO.....	887
Virtual Residential Gateway Authentication Scenarios.....	916

Virtual Residential Gateway Home LAN Extension.....	940
Virtual Residential Gateway Home Pool Management.....	960
WiFi Aggregation and Offload — Basic Open SSID.....	975
WiFi Aggregation and Offload — Basic Secure SSID with Distributed RADIUS Proxy.....	994
WiFi Aggregation and Offload — IPv4/v6 Dual-Stack UEs.....	1014
WiFi Aggregation and Offload — Migrant User Support.....	1028
WiFi Aggregation and Offload — Open SSID with DSM and Lawful Intercept.....	1038

List of tables

Table 1: ARP Host Time-Related Parameters.....	28
Table 2: Correlation of Hosts and BSA/BSR Services.....	42
Table 3: BSA/BSR Configuration for Host-1 Operation.....	46
Table 4: BSA/BSR Configuration for Host-2 Operation.....	47
Table 5: BSA/BSR Configuration for Host-3 Operation.....	48
Table 6: 64-bit WAN Mode Versus 128-bit WAN Mode.....	181
Table 7: Reserved PPPoE tags.....	278
Table 8: LCP and IPCP code.....	280
Table 9: Applicable Subscriber-Prefix Parameters.....	326
Table 10: Timer Parameters.....	327
Table 11: Router Advertisements Parameters.....	329
Table 12: RADIUS AVPs.....	331
Table 13: Local User Database Parameters.....	333
Table 14: DHCP Lease State Information.....	333
Table 15: Subscriber Prefix Parameters.....	356
Table 16: Subscriber Prefix Subnetting for SLAAC.....	356
Table 17: Subscriber-Prefix Parameters.....	359
Table 18: Prefix Subnetting for delegated-prefix-length /56.....	359
Table 19: RADIUS AVPs.....	360
Table 20: SLAAC-Related Parameters.....	363
Table 21: IPv6CP Nack Message Format.....	375

Table 22: Input/Output Rates.....	473
Table 23: Input and output rates throughout the subscriber QoS hierarchy.....	504
Table 24: Supported DHCP Option 82 Sub-Options.....	576
Table 25: Information in DHCP Lease State.....	581
Table 26: L2TPv2 header fields and descriptions.....	609
Table 27: AVP header fields and descriptions.....	610
Table 28: Generic L2TP RADIUS attributes.....	616
Table 29: Nokia Specific L2TP RADIUS attributes.....	616
Table 30: Masking examples.....	661
Table 31: Dynamic Service Attribute List for Setup, Modify and Teardown.....	797
Table 32: Dynamic Service Actions on Control- and Data-Channel.....	798
Table 33: Function and Dictionary Relationship.....	801
Table 34: RADIUS inserted raw options.....	826
Table 35: Python Modified DHCP Fields.....	827
Table 36: CLI Inserted DHCP Options.....	827
Table 37: DHCP options inserted via RADIUS.....	838
Table 38: ARP/MAT/BD-MAC-Prefix Possible Combinations.....	953

List of figures

Figure 1: Bridged CO and Routed CO Example.....	19
Figure 2: ARP Hosts in a Bridged CO Environment Example.....	20
Figure 3: ARP Hosts in a Routed CO Environment Example.....	22
Figure 4: ARP Host Session Timeout Example.....	29
Figure 5: Trap Generation Example.....	31
Figure 6: Throttling Toward RADIUS Example.....	33
Figure 7: ARP Host Mobility Example.....	34
Figure 8: Bridged CO Network Topology.....	38
Figure 9: Key Concepts of Bridged CO Model.....	39
Figure 10: Flow Chart for Subscriber-Profile Identification Algorithm.....	41
Figure 11: Flowchart for SLA-Profile Identification Algorithm.....	42
Figure 12: Sample Topology.....	44
Figure 13: Functionality of Each Node.....	44
Figure 14: Host-1 Setup Process.....	53
Figure 15: Host-2 Setup Process.....	56
Figure 16: Host-3 Setup Process.....	62
Figure 17: General Redundancy Model.....	70
Figure 18: DHCP Relay Agent and Server Redundancy Model.....	71
Figure 19: Access-Driven and Local-Remote Model.....	72
Figure 20: Local-Remote Model – Active-Standby.....	73
Figure 21: Local-Remote Model – Subnet-Based Load Sharing.....	74

Figure 22: Local-Remote Model – Range-Based Load Sharing.....	74
Figure 23: Failover State Transition Diagram.....	75
Figure 24: VPRN-1 Service Configuration.....	78
Figure 25: Accessing a DHCP server.....	94
Figure 26: Addresses, Subnets, and Pools in a DHCPv4 Server.....	98
Figure 27: General Address Allocation for DHCP.....	100
Figure 28: Baseline Service Configuration.....	102
Figure 29: NASREQ trigger.....	125
Figure 30: Diameter protocol stack.....	138
Figure 31: Diameter network topology.....	138
Figure 32: Bridged RGW Scenario.....	152
Figure 33: Routed RGW Scenario.....	152
Figure 34: SLA-Profile and Sub-Profile.....	153
Figure 35: Subscriber Host Identification and Instantiation Process.....	156
Figure 36: Direct Address Assignment using LUDB/RADIUS.....	160
Figure 37: Indirect Address Assignment using a DHCP Server.....	161
Figure 38: Indirect Address Assignment using LAA.....	163
Figure 39: Wholesale/Retail Model 1.....	185
Figure 40: Wholesale/Retail Model 2.....	186
Figure 41: Layer 3 Wholesale/Retail.....	187
Figure 42: L2TP Wholesale-Retail Multicast.....	193
Figure 43: Network Topology Overview.....	199
Figure 44: Single BNG Setup with Multicast Redirection.....	200

Figure 45: Network Topology with MC-LAG.....	206
Figure 46: IPoE Multicast Message Flow.....	215
Figure 47: PPPoE Multicast Flow.....	219
Figure 48: Network Topology Overview.....	230
Figure 49: Example Topology.....	231
Figure 50: IPoE Subscriber Multicast Flow.....	240
Figure 51: PPPoE Multicast Flow.....	245
Figure 52: TPSDA Network Topology.....	257
Figure 53: Routed CO network topology.....	276
Figure 54: Discovery stage messages.....	279
Figure 55: LCP phase messages.....	281
Figure 56: CHAP handshaking overview process.....	282
Figure 57: PAP overview process.....	282
Figure 58: IPCP phase messages.....	283
Figure 59: Keepalive messages.....	283
Figure 60: Link termination phase.....	284
Figure 61: Authentication flow chart.....	294
Figure 62: Pado-Delay Scenario.....	314
Figure 63: Stateful IPoE Dual Stack Subscriber Hosts.....	317
Figure 64: Dual Stack IPoE Routed Gateway Service.....	318
Figure 65: DHCPv6 Lease Process (Part A).....	319
Figure 66: DHCPv6 Lease Process (Part B).....	320
Figure 67: Prefix Delegation.....	321

Figure 68: IPv6 Address/Prefix Timers.....	327
Figure 69: PPPoE Dual Stack Hosts.....	351
Figure 70: Dual Stack PPPoE Bridged Gateway Service Example.....	351
Figure 71: Dual Stack PPPoE Routed Gateway Service Example.....	352
Figure 72: Message Flow for a Dual Stack PPPoE Host.....	354
Figure 73: Dual Stack PPPoE for Routed Gateway.....	357
Figure 74: DHCPv6 Renewals.....	365
Figure 75: Topology.....	378
Figure 76: GTP Access to the BNG.....	420
Figure 77: EPS Bearer Across the Different Interfaces.....	421
Figure 78: GTP-C and GTP-U Encapsulation.....	422
Figure 79: GTP-U in Up and Downstream.....	423
Figure 80: GTP-C Control - GTP Host Creation.....	425
Figure 81: GTP-C Control - GTP Host Deletion.....	426
Figure 82: GTP Access Topology.....	426
Figure 83: GTP Tunnel and Subscriber Termination Configuration Logic.....	429
Figure 84: Test Environment Example.....	447
Figure 85: QoS Hierarchy in Expanded SLA Mode.....	448
Figure 86: Managing Congestion on HSQ in Expanded SLA Mode.....	472
Figure 87: Test environment example.....	481
Figure 88: QoS hierarchy in single SLA mode.....	482
Figure 89: Managing Congestion on HSQ in Single SLA Mode.....	503
Figure 90: IOM/IMM Paths Connecting to Switch Fabric Planes.....	513

Figure 91: Dynamic Bandwidth Rate Management.....	524
Figure 92: Falling-Percent-Reset.....	525
Figure 93: Admin-Bw Rate Management.....	526
Figure 94: IPoE Session.....	545
Figure 95: IPoE Session Key.....	547
Figure 96: IPoE Session Creation Flow.....	548
Figure 97: IPoE Session Creation via AAA/RADIUS.....	549
Figure 98: Configuring IPoE session authentication.....	554
Figure 99: Baseline configuration.....	555
Figure 100: Bridged CO Network Topology.....	571
Figure 101: Routed CO Network Topology.....	571
Figure 102: DHCP Lease Process.....	573
Figure 103: Subscriber Host Connectivity Verification.....	597
Figure 104: DHCP Proxy Server: Lease Split Operation.....	602
Figure 105: DHCP Proxy Server: Lease Split Operation, DHCP Client Disconnected.....	603
Figure 106: DHCP Host Mobility.....	604
Figure 107: PPP access architectures.....	607
Figure 108: Supported L2TP reachability options.....	608
Figure 109: RADIUS-triggered tunnel/session setup without LNS renegotiation.....	612
Figure 110: RADIUS-triggered tunnel/session setup with LNS renegotiation.....	613
Figure 111: Running multiple PPP sessions over a single L2TP tunnel.....	614
Figure 112: PPP user-initiated release/terminate.....	614
Figure 113: L2TP tunnel and session state diagram.....	615

Figure 114: Base router hosted LAC with single endpoint/single tunnel.....	618
Figure 115: Base router hosted LAC with multiple endpoints.....	619
Figure 116: VRF hosted LAC.....	620
Figure 117: RADIUS returns L2TP tunnel group.....	622
Figure 118: LUDB returns L2TP tunnel group.....	623
Figure 119: L2TP keepalive mechanism.....	644
Figure 120: Floating peers accept.....	647
Figure 121: Floating peers ignore.....	647
Figure 122: Floating peers reject.....	648
Figure 123: LUDB applications.....	655
Figure 124: Processing an LUDB lookup request.....	656
Figure 125: Creating LUDBs and LUDB entries.....	657
Figure 126: Host matching examples.....	662
Figure 127: Host matching examples (continued).....	663
Figure 128: LUDB access via a DHCPv4 server.....	683
Figure 129: Example configuration.....	687
Figure 130: Decoding the ESM user option.....	694
Figure 131: LUDB authentication.....	700
Figure 132: Direct and indirect LUDB authentication.....	701
Figure 133: LUDB parameters for IPoE.....	703
Figure 134: LUDB parameters for PPPoE.....	704
Figure 135: LUDB authentication for regular SAPs.....	705
Figure 136: LUDB authentication for capture and managed SAPs.....	706

Figure 137: Baseline setup.....	708
Figure 138: Network Topology.....	730
Figure 139: MC-Ring Layer 2 CO Dual Homing.....	751
Figure 140: Dual homing Under Steady-State Condition.....	752
Figure 141: Broken Ring State.....	753
Figure 142: Network Topology.....	754
Figure 143: Unicast Services — Logical Topology.....	761
Figure 144: Multicast Service — Logical Setup.....	769
Figure 145: Test topology.....	775
Figure 146: Principle Model of Dynamic Data Services.....	791
Figure 147: Test Topology.....	793
Figure 148: Building Blocks of Dynamic Data Services.....	793
Figure 149: Hierarchy of Snippets.....	802
Figure 150: DHCPv4 Lease-Time Inserted by RADIUS and DHCPv4 Server.....	829
Figure 151: Python Injected Hint for Lease-Time.....	830
Figure 152: Format of the IA-NA Option.....	831
Figure 153: Format of the IA Address Option.....	831
Figure 154: Topology.....	832
Figure 155: Components of the Routed CO Model.....	853
Figure 156: Numbered Scenario For IES 1.....	860
Figure 157: Unnumbered Scenario for IES 1.....	870
Figure 158: Hybrid Configuration.....	878
Figure 159: Network Redundancy Components for ESM Routed CO.....	888

Figure 160: BRG and home device management.....	917
Figure 161: Explicit BRG authentication.....	918
Figure 162: Implicit BRG authentication.....	919
Figure 163: Example service configuration for explicit and implicit BRG authentication.....	921
Figure 164: vRGW-HLE.....	941
Figure 165: BD Connections in the Data Plane.....	941
Figure 166: ARP Requests Flooded with AAR Disabled.....	950
Figure 167: No ARP Request Flooding with AAR Enabled.....	951
Figure 168: MAT - Access to Network Direction.....	952
Figure 169: MAT - Network to Access Direction.....	952
Figure 170: Virtual Residential Gateway in the Network with Bridged Residential Gateway at Home.....	961
Figure 171: Services Configuration Overview.....	962
Figure 172: Call Flow for Open SSID.....	977
Figure 173: WiFi Offload Scenario with Open SSID and Local DHCP Server.....	978
Figure 174: WiFi Offload Scenario with Secure SSID and L2-Aware NAT.....	996
Figure 175: Call Flow for Secure SSID with DSM.....	1006
Figure 176: DHCPv4 + SLAAC/64 — Open SSID.....	1015
Figure 177: DHCPv4 + SLAAC/64 model — closed SSID.....	1016
Figure 178: DHCPv4 + SLAAC/64 with DHCPv4 linking model — closed SSID.....	1017
Figure 179: DHCPv4 + DHCPv6/128 IA_NA model — closed SSID.....	1018
Figure 180: DHCPv4 + SLAAC/64 with DHCPv4 linking model — DTA.....	1019
Figure 181: Sequence of events to establish and authenticate a migrant User.....	1033
Figure 182: Sequence of events to establish and authenticate a migrant user (continued).....	1034

Figure 183: WiFi Offload Scenario with Open SSID, DSM and LI.....1039

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 24.10.R2. 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.

ARP Hosts

This chapter describes advanced ARP host configurations.

Topics in this chapter include:

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

Applicability

This section describes ARP hosts and is applicable to the 7450 ESS and 7750 SR series and was tested on SR OS 13.0 R3.

Overview

In business access areas, CPEs typically get their IP address information through DHCP and PPPoE. However, CPE network facing interfaces can be configured statically. In such cases, the first packet the network receives from the user is an ARP to the Broadband Service Aggregator (BSA) or Broadband Service Router (BSR) interface. In order to accommodate such configurations, Enhanced Subscriber Management (ESM) feature set supports RP hosts.

In practice, this means that authentication, self-provisioning and Service Level Agreement (SLA) enforcement can be triggered by reception of ARP packets.

The BRAS node will learn the IP-MAC association based on the received ARP request packet and will instantiate subscriber hosts based on results from RADIUS authentication, the same way as this would happen through DHCP or PPPoE.

This chapter provides configuration and troubleshooting commands for ARP hosts. Features common with other host types and not unique to ARP hosts are not described in this chapter. (Not exhaustive list: RADIUS managed routes, routed subscriber with dynamic BGP peering, Wholesale/Retail, Managed SAPs configurations, ESM related host limitation mechanisms, host High-Availability, multi-chassis peer synchronization).

Knowledge of the Triple Play Service Delivery Architecture (TPSDA) concepts is assumed throughout this chapter.



WARNING:

Enhanced Subscriber Management and RADIUS authentication are mandatory for the use of ARP hosts.

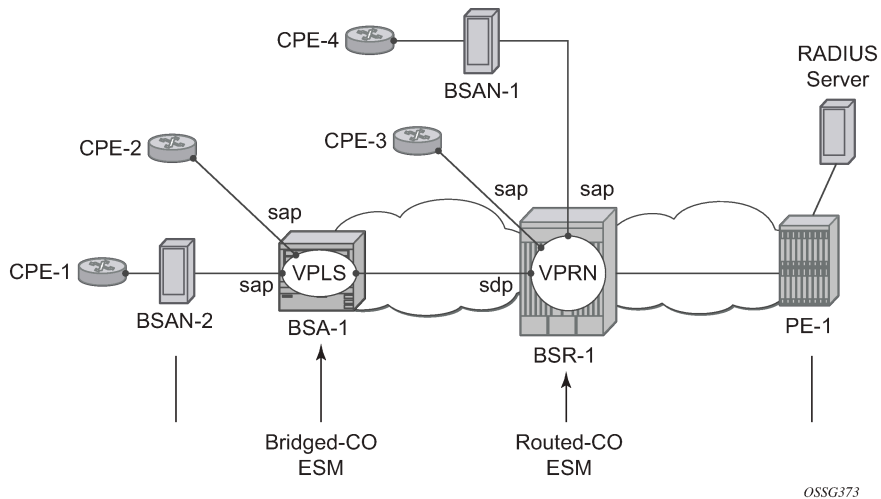
ARP host is supported in the bridged CO (VPLS) and the routed CO (Subscriber Interface) model. It is triggered by the first ARP packet received from the host. ARP host is also supported in a wholesale/retail context and on managed SAPs (MSAP).

The IP and MAC addresses are extracted from the ARP request. They are copied in the access-request message sent to the RADIUS server:

- RADIUS attribute [1] Username = IP address
- VSA [26][27] Client Hardware Address = MAC address

On successful authentication RADIUS will reply with an access-accept message, and ESM will create the ARP host. ESM string assignment options are out of the scope for this scenario.

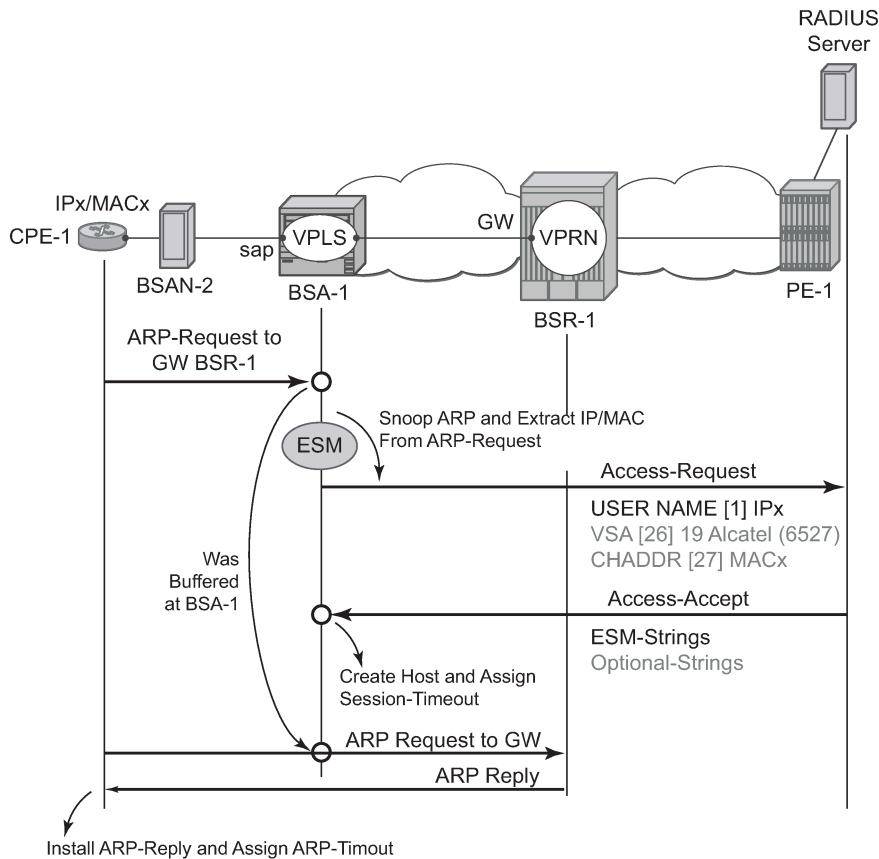
Figure 1: Bridged CO and Routed CO Example



Configuration

ARP Hosts in a Bridged CO Environment

Figure 2: ARP Hosts in a Bridged CO Environment Example



OSSG374

Enabling ARP-host for the Bridged CO model requires a composite service; a VPLS on the BSA node and a VPRN/IES on the BSR node. RADIUS authentication and subscriber management, which mandates IP-MAC or NH-MAC type anti-spoofing, are mandatory for ARP hosts.

```
# on BSA-1
configure
service
  vpls 2 customer 1 create
  description "ARP host - Bridged CO"
  stp
  shutdown
exit
sap 1/1/1:1 create
  authentication-policy "authentication-1"
  anti-spoof ip-mac
  sub-sla-mgmt
  sub-ident-policy "sub-id-default"
```



```

        multi-sub-sap 10
        no shutdown
    exit
    arp-host
    no shutdown
    exit
    exit
    spoke-sdp 12:2 create
    exit
    no shutdown
    exit
    exit
    exit

```

The RADIUS authentication policy does not require any specific parameter settings. The RADIUS username attribute will always contain the host IP address, meaning that the authentication policy parameter `user-name-format` is irrelevant for ARP hosts.

```

configure
  subscriber-mgmt
    authentication-policy "authentication-1" create
    password ALU
    radius-authentication-server
      server 1 address 172.16.1.1 secret ALU
    exit
    # re-authentication is optional
    re-authentication
    # only required when RADIUS Disconnect is needed (optional)
    accept-authorization-change
  exit
exit

```

The CPE ARPs are snooped and the first CPE ARP triggers a RADIUS accept-request. Any subsequent ARPs will trigger RADIUS re-authentication only if the ARP host configurable `min-auth-interval` is expired and if the previously defined re-authentication parameter is set. The initial ARP is only forwarded to the BSR-1 upon successful RADIUS authentication by means of a RADIUS access-accept message. The same RADIUS access-accept message and passing the several session limit checks, triggers the instantiation of the host.

The BSR-1 node requires a VPRN/IES as part of the composite service. No ARP-host-specific parameters are required on the BSR-1 for the bridged CO model.

```

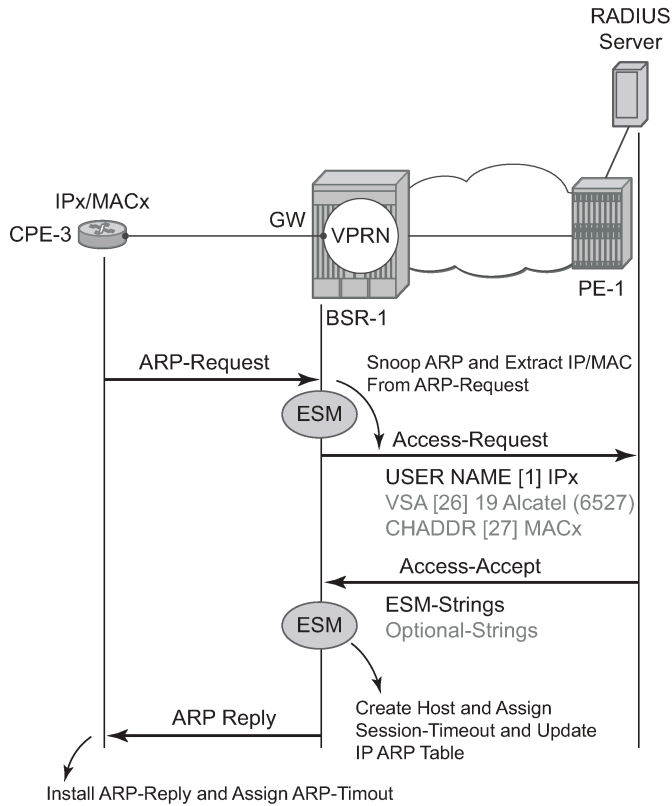
# on BSR-1
configure service
  vprn 1 customer 1 create
  route-distinguisher 64496:1
  auto-bind-tunnel
  resolution-filter
    ldp
  exit
  resolution filter
  exit
  vrf-target target:64496
  interface "int-BSA1-p2mp-1" create
  description "ARP host - Bridged CO" address 10.2.0.6/29
  ip-mtu 1500
  spoke-sdp 21:2 create
  exit
  exit
  exit

```

```
exit
```

ARP Hosts in a Routed CO Environment

Figure 3: ARP Hosts in a Routed CO Environment Example



OSSG375

Enabling ARP-host for the routed CO model is identical for VPRN and IES services. RADIUS authentication and subscriber management, which mandates IP-MAC or NH-MAC type anti-spoofing, are mandatory for ARP hosts.

The initial ARP will, only upon successful RADIUS authentication and passing the several sessions limit checks, create the ARP host. The ARP reply or update of the IP ARP table is not performed on any unsuccessful RADIUS authentication.

```
# on BSR-1
configure service
  vprn 1 customer 1 create
    route-distinguisher 64496:1
    auto-bind-tunnel
    resolution-filter
      ldp
    exit
  resolution filter
  exit

  vrf-target target:64496:1
```

```
subscriber-interface "sub-int-1" create
description "ARP host - Routed CO" address 10.1.0.6/29
group-interface "group-int-1" create
authentication-policy "authentication-1"
sap 1/1/1:1 create
anti-spoof ip-mac
sub-sla-mgmt
sub-ident-policy "sub-id-default"
no shutdown
exit
exit
arp-host
no shutdown
exit
exit
exit
exit
exit
```

RADIUS User Configuration Bridged/Routed CO

The username in the RADIUS access request is always the statically configured IP address from the CPE and configured as key in the RADIUS users file. The RADIUS Framed-Route attribute is not required and is silently ignored (if returned to BSA/BSR node).

```
"10.1.0.1"      Auth-Type := Local, User-Password == ALU
                Alc-Subsc-ID-Str = "arp-host-routed-%{User-name}",
                Alc-Subsc-Prof-Str = "sub-profile-1",
                Alc-SLA-Prof-Str = "sla-profile-1"

"10.2.0.1"      Auth-Type := Local, User-Password == ALU
                Alc-Subsc-ID-Str = "arp-host-bridged-%{User-name}",
                Alc-Subsc-Prof-Str = "sub-profile-1",
                Alc-SLA-Prof-Str = "sla-profile-1"
```

Setup and Debugging of ARP Host

Identical methodologies are used to debug/setup and troubleshoot ARP hosts for the bridged or the Routed CO model. The Routed CO model is used as an example through the rest of this section on ARP hosts.

There are two modes of ARP host debugging: all and dropped-only. The dropped-only mode shows all cases where the creation of the ARP host will be unsuccessful.

By default, all enabled ARP hosts on a service will be monitored. More specific filtering on a particular IP, MAC or SAP is optional.

All main traps are by default cyclically logged in log-id 99 and can be viewed anytime.

```
debug service id 1 arp-host mode all
```

ARP host mandate RADIUS authentication and a separate debug option is available for RADIUS interaction.

```
debug radius detail
```

CPE-3 with statically configured IP 10.1.0.1 sends an ARP to the BSR-1 gateway.

```
1 2015/06/22 15:48:00.72 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
Access-Request(1) 172.16.1.1:1812 id 2 len 79 vrid 1 pol authentication-1
  USER NAME [1] 8 10.1.0.1
  PASSWORD [2] 16 gy3yhtT5dF9YYilHtiNNk
  NAS IP ADDRESS [4] 4 192.0.2.2
  VSA [26] 19 Alcatel(6527)
  CHADDR [27] 17 00:00:0a:01:00:01
"

2 2015/06/22 15:48:00.74 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
Access-Accept(2) id 2 len 94 from 172.16.1.1:1812 vrid 1 pol authentication-1
  VSA [26] 26 Alcatel(6527)
  SUBSC ID STR [11] 24 arp-host-routed-10.1.0.1
  VSA [26] 15 Alcatel(6527)
  SUBSC PROF STR [12] 13 sub-profile-1
  VSA [26] 15 Alcatel(6527)
  SLA PROF STR [13] 13 sla-profile-1
"

3 2015/06/22 15:48:00.75 CEST MINOR: DEBUG #2001 vprn1 ARP Host
"ARP Host: Created ARP host
  VPRN 1, SAP 1/1/1:1

  IP: 10.1.0.1
  MAC: 00:00:0a:01:00:01
"

A:BSR-1# show log log-id 99
---snip---
58 2015/06/22 15:48:00.73 CEST WARNING: SVCNMR #2500 Base Subscriber created
"Subscriber arp-host-routed-10.1.0.1 has been created in the system"
```

The user name in the RADIUS access-request contains the CPE's IP address independent from the user-name-format defined in the authentication policy. The MAC address of the ARP host is included in the RADIUS access-request as VSA (Alc-Client-Hardware-Addr) independent on the include-radius-attribute mac-address parameter from the authentication policy.

The **show service id 1 arp-host** command displays all active ARP hosts on this service.

```
A:BSR-1# show service id 1 arp-host
=====
ARP host table, service 1
=====
IP Address      Mac Address      Sap Id      Remaining      MC
                  Time                               Stdby
-----
10.1.0.1        00:00:0a:01:00:01 1/1/1:1      03h59m23s
-----
Number of ARP hosts : 1
=====
A:BSR-1#
```

More specific filters such as **sap**, **ip-address**, **mac** and others can be used to show dedicated ARP hosts created on the BSR.

```
A:BSR-1# show service id 1 arp-host ip-address 10.1.0.1 detail
=====
```

```

ARP hosts for service 1
=====
Service ID          : 1
IP Address          : 10.1.0.1
MAC Address         : 00:00:0a:01:00:01
Subscriber-interface : sub-int-1
Group-interface     : group-int-1
SAP                 : 1/1/1:1
Remaining Time      : 03h59m15s

Sub-Ident           : "arp-host-routed-10.1.0.1"
Sub-Profile-String  : "sub-profile-1"
SLA-Profile-String  : "sla-profile-1"
App-Profile-String  : ""
ARP host ANCP-String : ""
ARP host Int Dest Id : ""
Category-Map-Name   : ""

RADIUS-User-Name    : "10.1.0.1"

Session Timeout (s) : 14400
Start Time          : 06/22/2015 15:48:00
Last Auth           : 06/22/2015 15:48:00
Last Refresh        : 06/22/2015 15:48:00
Persistence Key     : N/A
-----
Number of ARP hosts : 1
=====
A:BSR-1#

```

Dynamically created ARP hosts are added as /32 addresses to the routing table and marked with protocol type Sub Mgmt. Routes of the Sub Mgmt protocol type are not exported into vpn-ipv4 by the default vrf-target policy. A separate, dedicated vrf-export policy is required to achieve this.

```

A:BSR-1# show router 1 route-table 10.1.0.0/24 longer
=====
Route Table (Service: 1)
=====
Dest Prefix[Flags]          Type  Proto   Age           Pref
Next Hop[Interface Name]    Metric
-----
10.1.0.0/29                 Local  Local    00h04m21s    0
sub-int-1                   0
10.1.0.1/32                 Remote Sub Mgmt  00h00m56s    0
[group-int-1]               0
-----
No. of Routes: 2
Flags: n = Number of times nexthop is repeated
      B = BGP backup route available
      L = LFA nexthop available
      S = Sticky ECMP requested
=====
A:BSR-1#

```

Specific ARP host counters can be shown or cleared using the CLI command **show/clear service id 1 ARP host statistics**.

```

A:BSR-1# show service id 1 arp-host statistics
=====
ARP host statistics
=====
Num Active Hosts           : 1

```

```
Received Triggers      : 5
Ignored Triggers      : 3
Ignored Triggers (overload) : 0
SHCV Checks Forced    : 0
Hosts Created         : 1
Hosts Updated         : 1
Hosts Deleted         : 0
Authentication Requests Sent : 4
=====
A:BSR-1#
```

The ARP hosts mandate Enhanced Subscriber managed (ESM) and therefore an anti-spoofing configuration (IP-MAC or NH-MAC). The anti-spoofing table with active hosts can be viewed with the command **show service id 1 subscriber-hosts**.

```
A:BSR-1# show service id 1 subscriber-hosts
=====
Subscriber Host table
=====
Sap      Subscriber
IP Address
MAC Address      PPPoE-SID Origin      Fwding State
-----
1/1/1:1      arp-host-routed-10.1.0.1
10.1.0.1
00:00:0a:01:00:01      N/A      ARP-Host      Fwding
-----
Number of subscriber hosts : 1
=====
A:BSR-1#
```

An ARP host can be manually deleted from the system using one of the two following methods:

- clear service id 1 arp-host
- RADIUS disconnect message

Using the first method, **clear service id 1 arp-host** and omitting any more specific parameter results in the removal of all ARP hosts in this service. Extra filters such as **ip-address**, **mac** or **sap-id** are required to remove a specific ARP host.

```
*A:BSR-1# clear service id 1 arp-host
- arp-host {all | mac <ieee-address> | sap <sap-id> | ip-address <ip-address[/mask]> }
- arp-host {port <port-id> | {inter-dest-id <intermediate-destination-id> | no-inter-dest-id}
[port <port-id>] }
- arp-host statistics [sap <sap-id> | interface <interface-name>]

A:BSR-1# clear service id 1 arp-host ip-address 10.1.0.1
```

Using the second method, RADIUS disconnect always result in the removal of a unique host because **nas-port-id** and **framed-ip-address** are mandatory parameters in the RADIUS disconnect message. This RADIUS disconnect message is used also for other host-types.

```
nas-port-id = 1/1/1:1
framed-ip-address=10.1.0.1
```

RADIUS disconnect messages are, for security reasons, rejected by default. The RADIUS disconnect messages can be accepted by enabling the **accept-authorization-change** parameter in the authentication

policy. The **debug radius detail** command and **show subscriber-mgmt authentication coa-statistics** can be used during troubleshooting.

```
10 2015/06/22 15:51:08.43 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Disconnect Request(40) id 247 len 44 from 172.16.1.1:46749 vrid 1
  SESSION ID [44] 22 02DAFF0000000255881288
"

11 2015/06/22 15:51:08.42 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Disconnect Ack(41) 172.16.1.1:46749 id 247 len 26 vrid 1 pol authentication-1
  TERMINATE CAUSE [49] 4 Admin Reset(6)
"

12 2015/06/22 15:51:08.43 CEST MINOR: DEBUG #2001 vprn1 ARP Host
"ARP Host: Removed ARP host
  VPRN 1, SAP 1/1/1:1

  IP: 10.1.0.1
  MAC: 00:00:0a:01:00:01
"
```

In both cases the ARP host with an IP address is removed from the system together with all related state information (such as an anti-spoof filter and an IP ARP entry).

ARP Host Session Timeout

The remaining time is initialized at the ARP host session timeout value (300s to 14400s) and starts counting down when an ARP host is instantiated. Ultimately the host is removed from the system when the remaining time becomes zero. Any subsequent arp-request or arp-reply for this host results in the remaining value to be reset to the host session timeout value defined.

The default assigned session timeout at ARP host creation time is 14400 seconds, a value which can be overruled by the optional RADIUS attribute session-Timeout but not by the node group-interface arp-timeout parameter.

RADIUS values lower than 300 seconds will be silently adjusted to 300 seconds and values above 14400 seconds are topped silently to 14400 seconds.

```
"10.1.0.1"    Auth-Type := Local, User-Password == ALU
              Alc-Subsc-ID-Str = "arp-host-routed-%{User-name}",
              Alc-Subsc-Prof-Str = "sub-profile-1",
              Alc-SLA-Prof-Str = "sla-profile-1",
              Session-Timeout = 300 # value in seconds
```

```
A:BSR-1# show service id 1 arp-host
=====
ARP host table, service 1
=====
IP Address      Mac Address      Sap Id           Remaining      MC
                  Stdby
-----
10.1.0.1        00:00:0a:01:00:01 1/1/1:1         00h04m56s
-----
Number of ARP hosts : 1
=====
```

```
A:BSR-1#
```

```
A:BSR-1# show service id 1 arp-host ip-address 10.1.0.1 detail
```

```
=====
```

```
ARP hosts for service 1
```

```
=====
```

```
Service ID      : 1
IP Address      : 10.1.0.1
MAC Address     : 00:00:0a:01:00:01
Subscriber-interface : sub-int-1
Group-interface : group-int-1
SAP             : 1/1/1:1
Remaining Time  : 00h04m39s

Sub-Ident       : "arp-host-routed-10.1.0.1"
Sub-Profile-String : "sub-profile-1"
SLA-Profile-String : "sla-profile-1"
App-Profile-String : ""
ARP host ANCP-String : ""
ARP host Int Dest Id : ""
Category-Map-Name : ""
```

```
RADIUS-User-Name : "10.1.0.1"
```

```
Session Timeout (s) : 300
Start Time          : 06/22/2015 15:53:11
Last Auth           : 06/22/2015 15:53:11
Last Refresh        : 06/22/2015 15:53:11
Persistence Key     : N/A
```

```
-----
```

```
Number of ARP hosts : 1
```

```
=====
```

```
A:BSR-1#
```

Typical time related parameters of the ARP host are:

Table 1: ARP Host Time-Related Parameters

Parameter	Comment
Session Timeout	Time value in seconds and retrieved by default or by the RADIUS Accept message and pasted into the remaining time at the moment of ARP host creation or RADIUS re-authentication.
Remaining Time	The remaining time before the ARP host is deleted from the system (updated each time an ARP request/reply is seen for this host).
Start Time	Time and date when this host was created (first ARP received).
Last Auth	Time and date when this host was last RADIUS authenticated.
Last Refresh	Time and date when last ARP was received for this host.

ARP hosts do not have an expiry timer in the ARP table and have type **managed**.

```
A:BSR-1# show service id 1 arp 10.1.0.1
```

```
=====
```

```
ARP Table
```

```
=====
```

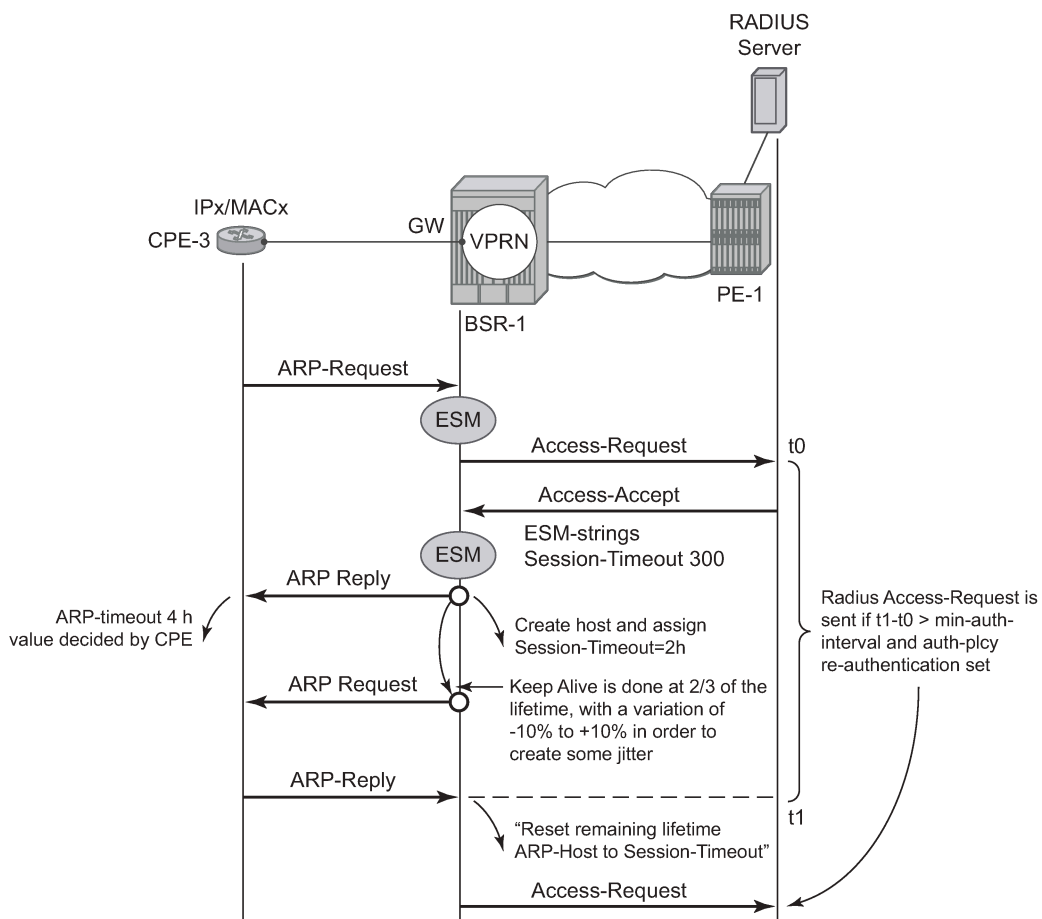

IP Address	MAC Address	Type	Expiry	Interface	SAP
10.1.0.1	00:00:0a:01:00:01	Managed	00h00m00s	group-int-1	1/1/1:1

A:BSR-1#

An automatic mechanism is implemented to handle the possible asynchronicity between the ARP session timeout values installed on the BSR and the ARP timeouts installed on the CPE. This mechanism is mostly effective in case the timeout on the CPE exceeds the timeout on the BSR. In this case, the BSR session would expire, resulting in a host removal with a deletion of the corresponding anti-spoof entry because the CPE ARP request arrives too late. This CPE ARP request will however recreate the session but requires the complete setup of the host RADIUS authentication included. This mechanism causes unwanted service interruption for this ARP host.

A better approach, which is implemented in an automatic way, and illustrated in [Figure 4: ARP Host Session Timeout Example](#) is an ARP request triggered by the BSR towards the CPE prior to the session timeout. The ARP reply sent by the CPE then will reset the remaining lifetime for the ARP host to the session timeout. If the ARP reply is received outside the **min-auth-interval** window and the parameter re-authentication in the authentication policy is set, then RADIUS re-authentication is executed. This re-authentication mechanism is described later in the [Throttling Toward RADIUS](#) section.

Figure 4: ARP Host Session Timeout Example



OSSG376

This mechanism, also known as automatic Subscriber Host Connectivity Verification (SHCV), will prevent that the host will be deleted and re-created, resulting in undesired service interruptions, in case asynchronous CPE-BSR ARP session values would be used.

The **debug service id 1 host-connectivity-verify** command shows the sequence of events and can be used for troubleshooting. Debugging and ARP host counters show the automatic SHCV mechanism with an active CPE.

```

4 2015/06/25 16:32:45.21 CEST MINOR: DEBUG #2001 vprn1 SHCV
"SHCV: Forced Check Scheduled
  1/1/1:1
  ARP host 10.1.0.1 00:00:0a:01:00:01"

5 2015/06/25 16:32:46.11 CEST MINOR: DEBUG #2001 vprn1 SHCV
"SHCV: Forced Check
  1/1/1:1
  ARP host 10.1.0.1 00:00:0a:01:00:01"

6 2015/06/25 16:32:46.12 CEST MINOR: DEBUG #2001 vprn1 SHCV
"SHCV: Received Reply
  1/1/1:1
  ARP host 10.1.0.1 00:00:0a:01:00:01"

7 2015/06/25 16:32:46.12 CEST MINOR: DEBUG #2001 vprn1 ARP Host
"ARP Host: Updated ARP host
  VPRN 1, SAP 1/1/1:1

  IP: 10.1.0.1
  MAC: 00:00:0a:01:00:01
"

8 2015/06/25 16:34:29.34 CEST MINOR: DEBUG #2001 vprn1 SHCV
"SHCV: Received Reply
  1/1/1:1
  ARP host 10.1.0.1 00:00:0a:01:00:01"

9 2015/06/25 16:34:29.34 CEST MINOR: DEBUG #2001 vprn1 ARP Host
"ARP Host: Updated ARP host
  VPRN 1, SAP 1/1/1:1

  IP: 10.1.0.1
  MAC: 00:00:0a:01:00:01
"

```

```

A:BSR-1# show service id 1 arp-host statistics

```

```

=====
ARP host statistics
=====

```

```

Num Active Hosts      : 1
Received Triggers     : 3
Ignored Triggers      : 0
Ignored Triggers (overload) : 0
SHCV Checks Forced    : 1
Hosts Created         : 1
Hosts Updated         : 2
Hosts Deleted         : 0
Authentication Requests Sent : 1
=====

```

```

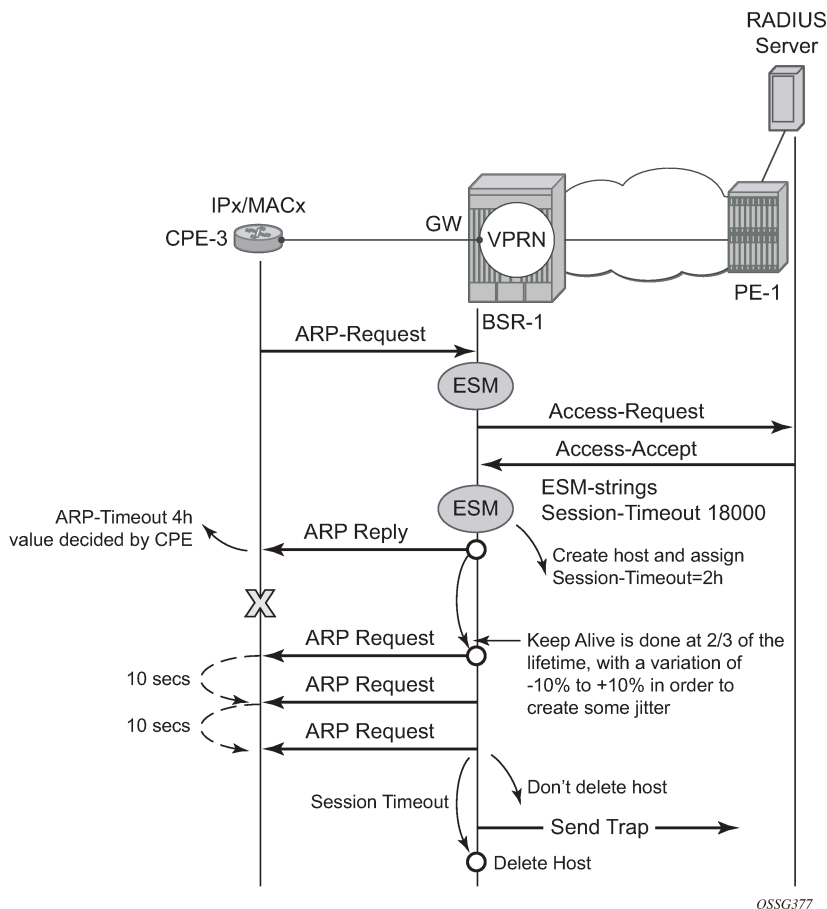
A:BSR-1#

```

CPEs that are not active and which therefore do not respond to ARP requests as part of the automatic SHCV check will be rechecked three times with 10 second intervals.

The number of retries and the interval cannot be changed. A trap is generated, but the ARP host is not removed and will remain until the session-timeout expires or until the host revives. This mechanism is displayed in [Figure 5: Trap Generation Example](#).

Figure 5: Trap Generation Example



OSSG377

```

16 2015/06/25 16:42:38.21 CEST MINOR: DEBUG #2001 vprn1 SHCV
"SHCV: Forced Check Scheduled
 1/1/1:1
 ARP host 10.1.0.1 00:00:0a:01:00:01"

17 2015/06/25 16:42:39.11 CEST MINOR: DEBUG #2001 vprn1 SHCV
"SHCV: Forced Check
 1/1/1:1
 ARP host 10.1.0.1 00:00:0a:01:00:01"

18 2015/06/25 16:42:49.11 CEST MINOR: DEBUG #2001 vprn1 SHCV
"SHCV: Forced Check
 1/1/1:1
 ARP host 10.1.0.1 00:00:0a:01:00:01"

19 2015/06/25 16:42:59.11 CEST MINOR: DEBUG #2001 vprn1 SHCV
"SHCV: Forced Check

```

```
1/1/1:1
ARP host 10.1.0.1 00:00:0a:01:00:01"

20 2015/06/25 16:43:09.11 CEST MINOR: DEBUG #2001 vprn1 SHCV
"SHCV: Connectivity Lost
1/1/1:1
ARP host 10.1.0.1 00:00:0a:01:00:01"

21 2015/06/25 16:43:10.21 CEST MINOR: DEBUG #2001 vprn1 ARP Host
"ARP Host: Removed ARP host
VPRN 1, SAP 1/1/1:1

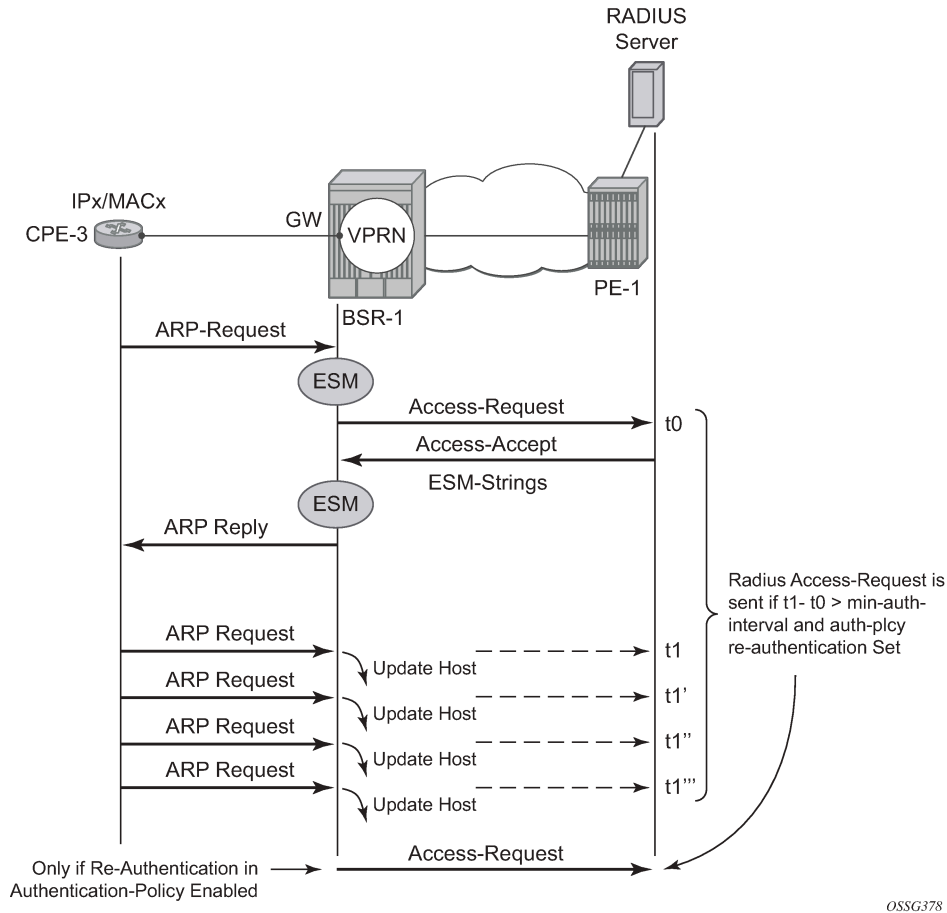
IP: 10.1.0.1
MAC: 00:00:0a:01:00:01
"
```

Throttling Toward RADIUS

A new ARP request from the ARP host will trigger RADIUS re-authentication only when the min-auth-interval has expired. The minimum RADIUS authentication interval between two consecutive authentication attempts for the same ARP host is by default 15 minutes but can range between 1 and 6000 minutes.

```
configure
  service
    vprn 1 customer 1 create
      ---snip---
      arp-host
        min-auth-interval 60 # value in minutes
        no shutdown
      exit
    exit
  exit
```

Figure 6: Throttling Toward RADIUS Example



```
A:BSR-1# show service id 1 arp-host detail
```

```
=====
```

```
ARP hosts for service 1
```

```
=====
```

```
Service ID       : 1
IP Address       : 10.1.0.1
MAC Address      : 00:00:0a:01:00:01
Subscriber-interface : sub-int-1
Group-interface  : group-int-1
SAP              : 1/1/1:1
Remaining Time   : 00h04m31s

Sub-Ident        : "arp-host-routed-10.1.0.1"
Sub-Profile-String : "sub-profile-1"
SLA-Profile-String : "sla-profile-1"
App-Profile-String : ""
ARP host ANCP-String : ""
ARP host Int Dest Id : ""
Category-Map-Name : ""

RADIUS-User-Name : "10.1.0.1"

Session Timeout (s) : 300
Start Time          : 06/22/2015 15:59:32
```

A:BSR-1#

```

persistence
  subscriber-mgmt
    location cf3:
  exit
exit

```

The persistence key and the index into the persistency file are linked to the ARP host at host creation time.

```

A:BSR-1# show service id 1 arp-host detail
=====
ARP hosts for service 1
=====
Service ID           : 1
IP Address           : 10.1.0.1
MAC Address          : 00:00:0a:01:00:01
Subscriber-interface : sub-int-1
Group-interface      : group-int-1
SAP                  : 1/1/1:1
Remaining Time       : 00h04m32s

Sub-Ident            : "arp-host-routed-10.1.0.1"
Sub-Profile-String   : "sub-profile-1"
SLA-Profile-String   : "sla-profile-1"
App-Profile-String   : ""
ARP host ANCP-String : ""
ARP host Int Dest Id : ""
Category-Map-Name    : ""

RADIUS-User-Name     : "10.1.0.1"

Session Timeout (s)  : 300
Start Time           : 06/22/2015 15:59:32
Last Auth            : 06/22/2015 15:59:32
Last Refresh         : 06/22/2015 16:03:33
Persistence Key       : 0x00000000
-----
Number of ARP hosts : 1
=====
*A:BSR-1#

```

The content of the stored record is viewed with the **tools dump persistency** command using the **persistency** key as a record number.

```

*A:BSR-1# tools dump persistence submgt record 0x00000000
-----
Persistence Record
-----
Client       : submgt
Persist-Key  : 0x00000000
Filename     : cf3:\submgt.011
Entries      : Index  FedHandle  Last Update          Action Valid
               000040  0x00000000  2015/06/22 14:01:31 (UTC) ADD    Yes
Data         : 243 bytes

Host Type    : ARP host
Service ID   : 1
SAP ID       : 1/1/1:1
NH MAC       : 00:00:0a:01:00:01
Created      : 2015/06/22 13:59:32 (UTC)
IP           : 10.1.0.1
Session Timeout: 300 (seconds)
RADIUS Fallback: NO

```

```
Acct-Sess-Id   : 02DAFF00000008558814C4
Multi-Sess-Id  : 02DAFF00000009558814C4
Class Attr     : 0 bytes
User-Name      : "10.1.0.1"
host is authenticated by radius: true
Subscriber-Id  : "arp-host-routed-10.1.0.1"
Sub-Profile-Str: "sub-profile-1"
SLA-Profile-Str: "sla-profile-1"
```

*A:BSR-1#

Session Limitation Options

The maximum number of allowed arp-hosts in a bridged CO model can be configured with the per SAP parameter **host-limit** in the range of 1 to 32767.

```
configure
  service
    vpls 2
      ---snip---
      sap 1/1/3:1
        arp-host
          # default value for host-limit is 1
          host-limit 1
          no shutdown
        exit
      exit
    exit
  exit
```

The maximum number of allowed arp-hosts in a routed CO model can be configured with the per group interface parameter **host-limit** in the range of 1 to 32767 and/or by the **sap-host-limit** parameter.

```
configure
  service
    vprn 1
      ---snip---
      arp-host
        # default value for host-limit is 1
        host-limit 1
        # default value for sap-host-limit is 1
        sap-host-limit 1
        no shutdown
      exit
    exit
```



WARNING:

Specific ESM-related host limit mechanisms such as **sla-profile host-limit** and **sub-sla-mgmt multi-sub-sap** also apply to ARP hosts but are not further elaborated in this section.

Debugging **arp-host mode dropped-only** indicates the dropped reason and a logging trap is included in the standard log 99.

```
56 2015/06/22 16:08:35.37 CEST MINOR: DEBUG #2001 vprn1 ARP Host
"ARP Host: Dropped trigger
  VPRN 1, SAP 1/1/1:1

  Problem: Interface limit (1) of ARP hosts reached
```



```
IP: 10.1.0.2
MAC: 00:00:0a:01:00:02
"

*A:BSR-1# show log log-id 99

---snip---

78 2015/06/22 16:08:55.52 CEST WARNING: SVCGR #2520 vprn1 ARP Host Population Error
"ARP host table population error on SAP 1/1/1:1 in service 1 - Interface limit (1) of ARP hosts
reached"
```

Increasing the **sap-host-limit** to 100 and the **host-limit** to 2000 results in the following summary:

```
*A:BSR-1# show service id 1 arp-host summary
=====
ARP host Summary, service 1
=====
Interface Name           Used      Provided  Admin State
-----
group-int-1              1         2000      inService
-----
Interfaces: 1            1
-----
=====
*A:BSR-1#
```

Conclusion

This chapter provides configuration and troubleshooting commands for dynamic ARP hosts. ARP hosts can be instantiated in a Layer 2 bridged CO (VPLS) environment as well as in a Layer 3 Routed CO (IES/ VPRN subscriber interface) context.

Bridged CO

This chapter provides information about bridged CO model of Triple Play Service Delivery Architecture (TPSDA).

Topics in this chapter include:

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

Applicability

The information and configuration in this chapter are based on SR OS Release 7.0.R4. This chapter is related only to the use of IPv4 DHCP hosts.

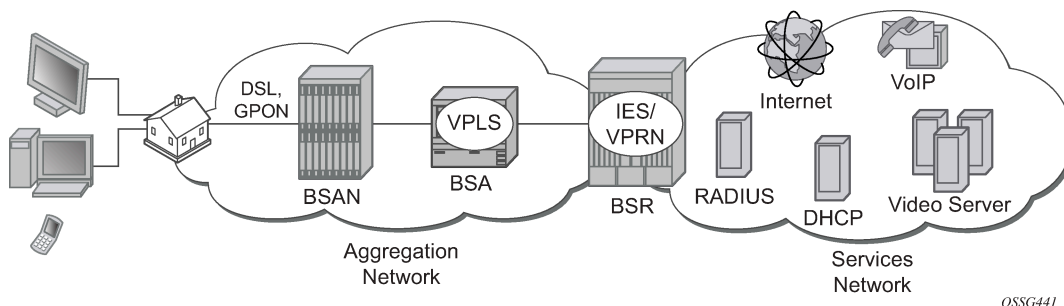
Knowledge of the Triple Play Service Delivery Architecture (TPSDA) concepts is assumed throughout this chapter.

Overview

This chapter provides information about basic technology, network topology and configuration examples which are used in Bridged CO model of Triple Play Service Delivery Architecture (TPSDA). Regardless of aggregation technologies which are used by customers, Nokia offers flexible and easy to use methodology to manage DHCP subscribers in Layer 2 domain and distribute subscriber management intelligence across multiple nodes.

Bridged CO is a basic TPSDA model and implies that access nodes are united in one Layer 2 aggregation network and VPLS is used as a primary technology for these purposes. This fact allows the use of subscriber management functionality on BSA nodes. Bridged CO network topology is shown in [Figure 8: Bridged CO Network Topology](#).

Figure 8: Bridged CO Network Topology



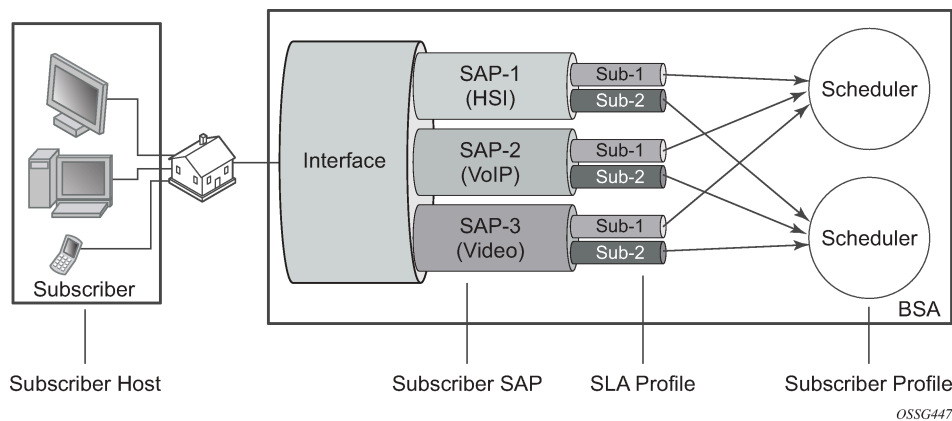
Following types of nodes are defined in Bridged CO model:

- Broadband Service Access Node (BSAN) — Access node connected to Layer 2 domain to aggregate all traffic from subscribers (IP DSLAM, ethernet switch).
- Broadband Services Aggregator (BSA) — Layer 2 node, which is capable for subscriber management in VPLS service (7450 ESS).
- Broadband Service Router (BSR) — Layer 3 node, which is capable for routing and service allocation (7750 SR).

As any model, Bridged CO introduces several key concepts that must be determined in advance. Major ones are presented in [Figure 9: Key Concepts of Bridged CO Model](#) and include:

- Subscriber— A set of hosts belonging to a single connection line (switch port, DSL line)
- Subscriber host — Unique customer device (could be PC, IP phone, STB, routed CPE).
- Subscriber-profile — Configured entity which defines the aggregate QoS for all hosts within a subscriber context.
- SLA-profile — Configured entity which defines QoS policies and filters for a subset of hosts within a subscriber context.
- Subscriber identification policy — Configured entity which defines the python script for dynamic subscriber host identification
- Authentication policy — Configured entity which defines the RADIUS servers to use for dynamic subscriber host identification
- Subscriber identification string — 32 characters identification string which uniquely identifies a subscriber on a node.

Figure 9: Key Concepts of Bridged CO Model



For normal operation each subscriber has to get several parameters / attributes:

- Subscriber-ID — Attribute, which uniquely identifies subscriber on the node and used as index key in subscriber database
- IP parameters — Attributes, which allows host to get access to services
- Subscriber profile and SLA profile for subscriber host — a set of filters and QoS policies.
- Lease time — Period when subscriber parameters are kept in subscriber database on the node.

There are several methods how to get each of these parameters:

- Static
- Python scripts
- RADIUS
- DHCP

Each of the subscriber parameters could be defined in several ways simultaneously. In this case use the following algorithm for selecting:

1. For subscriber profile

- a. A lookup in the **subscriber-explicit-map** is performed with the *sub-ident* string returned by the Python script, RADIUS or statically configured. If a matching entry is found, the sub-profile-name (if defined) is taken. If no entry was found go to [1.b](#).

```
A:BSA>config>subscr-mgmt# info
    explicit-subscriber-map
    entry key "Sub-1" sub-profile "sub-profile-1" sla-profile "sla-profile-1"
```

- b. If a sub-ident-policy is defined on the SAP, a lookup is done on its sub-profile-map with the sub-profile string from the script. The sub-profile-name is taken from the entry. If no entry was found go to [1.c](#).

```
A:BSA>config>service>vpls>sap# info
    sub-sla-mgmt
    sub-ident-policy "sub-ident-policy-1"

A:BSA>config>subscr-mgmt# info
    sub-ident-policy "sub-ident-policy-1" create
    sub-profile-map
    entry key "sub-1" sub-profile "sub-profile-1"
```


- c. If provisioned, the sub-profile-name is taken from the def-sub-profile attribute on the SAP. If no entry was found go to [1.d](#).

```
A:BSA>config>service>vpls>sap# info
    sub-sla-mgmt
    def-sub-profile "sub-profile-1"
```

- d. If a sub-profile with the name *default* is provisioned. If no entry was found DHCP Ack is dropped.

```
A:BSA>config>subscr-mgmt# info
    sub-profile "default" create
```

Figure 10: Flow Chart for Subscriber-Profile Identification Algorithm



2. For SLA profile

- a. The sla-profile-name is taken from the sub-ident string (returned by the Python script, RADIUS or statically configured) in the subscriber-explicit-map. If no entry was found go to [2.b2.a](#).

```
A:BSA>config>subscr-mgmt# info
    explicit-subscriber-map
    entry key "Sub-1" sub-profile "sub-profile-1" sla-profile "sla-profile-1"
```

- b. A lookup with the sla-profile string from the script is done in the sla-profile-map of the sub-profile found earlier. The corresponding sla-profile-name is used. If no entry was found go to [2.c](#):

```
A:BSA>config>subscr-mgmt# info
    sub-profile "sub-profile-1" create
    sla-profile-map
    entry key "sla-1" sla-profile "sla-profile-1"
```

- c. The sla-profile-name is taken from sla-profile-map of the sub-ident-policy configured on the SAP. The corresponding sla-profile-name is used. If no entry was found go to [2.d](#).

```
A:BSA>config>service>vpls>sap# info
    sub-sla-mgmt
    sub-ident-policy "sub-ident-policy-1"

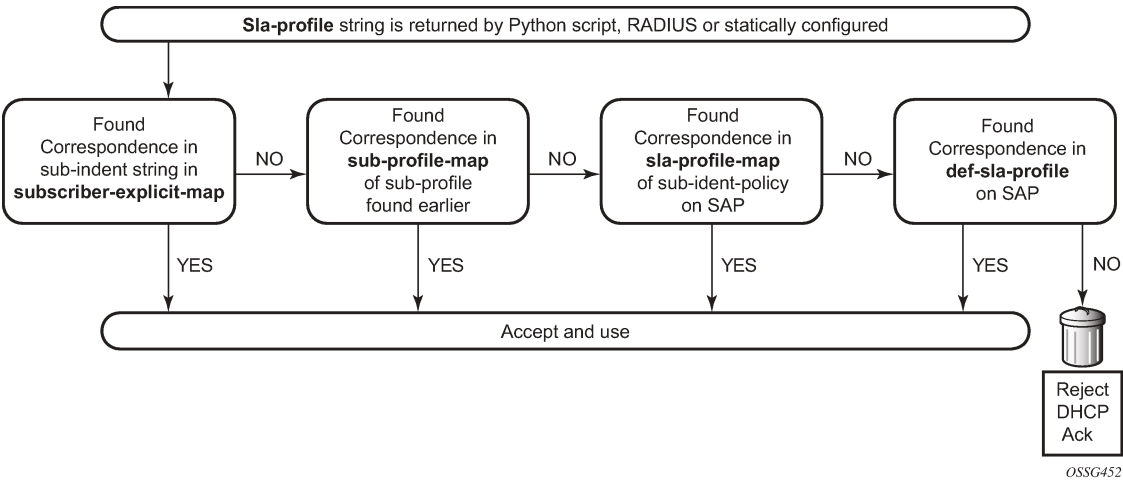
A:BSA>config>subscr-mgmt# info
    sub-ident-policy "sub-ident-policy-1" create
    sla-profile-map
    entry key "sla-1" sla-profile "sla-profile-1"
```



- d. The *sla-profile-name* is taken from the **def-sla-profile** attribute on the SAP. If no entry was found DHCP Ack is dropped.

```
A:BSA>config>service>vpls>sap# info
    sub-sla-mgmt
```

```
def-sla-profile "sla-profile-1"
```

Figure 11: Flowchart for SLA-Profile Identification Algorithm



- **Note:** Static configuration has priority over RADIUS configuration and RADIUS has priority over DHCP/Python scripts.
- **Note:** Each host can have different SLA-profile, while sub-profile applies to whole subscriber. The last definition of sub-profile will force all previously defined hosts to change their sub-profile.

Bridged CO supports typical access node connection models, such as:

- One VLAN per service (ESM for subscriber differentiation and SAP for service)
- One VLAN per subscriber (SAP for subscriber differentiation and QoS flag for service)
- One VLAN per access node (ESM for subscriber differentiation and QoS flag for service)

Each of these modes has its pros and cons, but this is out of scope of this document.

This configuration guide focuses on configuration of one subscriber with three different hosts. VLAN per service is used as mode of subscriber aggregation and mixed RADIUS and DHCP as subscriber identification method. IP termination is done in IES service of BSR.

Correlation of BSA/BSR services and subscriber hosts is presented in [Table 2: Correlation of Hosts and BSA/BSR Services](#) .

Table 2: Correlation of Hosts and BSA/BSR Services

	BSA (Service/Features)	BSR (Service/Features)
Host-1 ca:00:0c:54:00:08	VPLS-100 <ul style="list-style-type: none">• DHCP proxy server• SAP/SDP DHCP snoop• Sub-Ident origin via RADIUS• Sla/Sub-profiles via RADIUS	IES-100

	BSA (Service/Features)	BSR (Service/Features)
	<ul style="list-style-type: none"> IP options via RADIUS 	
Host-2 ca:01:08:10:00:08	VPLS-200 <ul style="list-style-type: none"> SAP/SDP DHCP snoop Sub-Ident origin through RADIUS Sla/Sub-profiles through RADIUS IP options through DHCP 	IES-200 <ul style="list-style-type: none"> DHCP relay
Host-3 ca:02:02:d0:00:08	VPLS-300 <ul style="list-style-type: none"> SAP/SDP DHCP snoop Sub-Ident origin through DHCP Sla/Sub-profiles through DHCP IP options through DHCP 	IES-300 <ul style="list-style-type: none"> DHCP relay

Different methods of authentication and address allocation were chosen for demonstration purposes. The customer is not limited to one method and can use a combination of methods as presented in this guide.

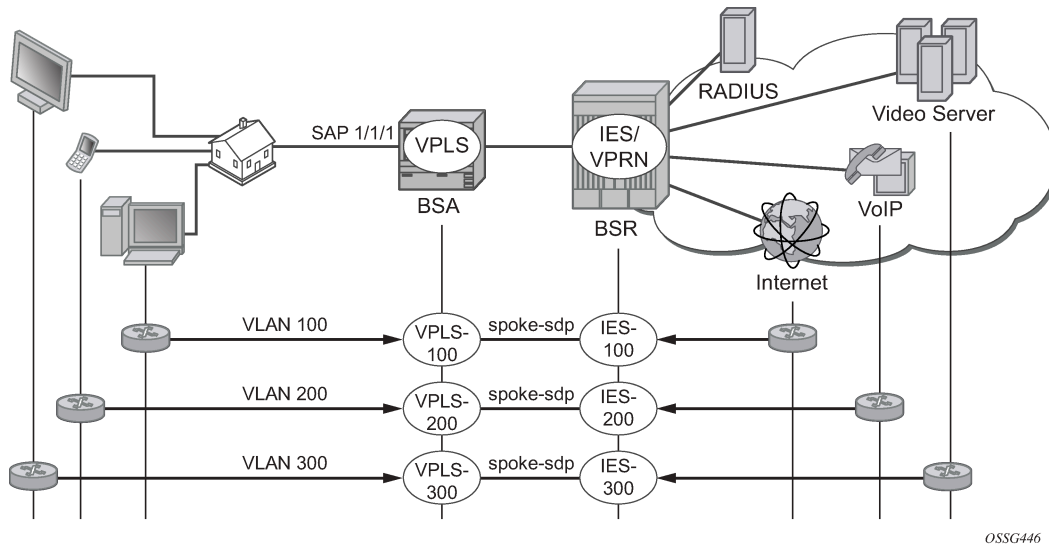
The following entities should be configured in advance. Refer to the appropriate platform user guide for specific information.

- Basic router configuration (interfaces, routing protocols, MPLS)
- External RADIUS server
- External/Local DHCP server

Configuration

A sample topology is presented in [Figure 12: Sample Topology](#).


Figure 12: Sample Topology



OSSG446

Bridged CO model requires certain techniques and features to be used on different nodes. Major methods are presented in [Figure 13: Functionality of Each Node](#).

Figure 13: Functionality of Each Node



OSSG442

The following configuration steps are required:

1. On BSA

- a. Configure subscriber management profiles
 - i. Configure sla profiles
 - ii. Configure subscriber profiles
 - iii. Configure subscriber identification policies
 - iv. Configure authentication and accounting policies if required
- b. Configure VPLS service
 - i. Configure split horizon group
 - ii. Configure SAP
 - 2.1 Configure anti-spoofing filters
 - 2.2 Configure DHCP snooping

- 2.3 Configure optional parameters (lease split, L2 DHCP relay agent, etc.)
 - 2.4 In case of RADIUS authentication apply authentication policy
 - 2.5 Configure ESM
 - iii. Configure SDP
 - 3.1 Configure DHCP snooping
2. On BSR
- a. Configure IES service
 - i. Configure IP interface
 - 1.1 Configure DHCP relay agent if required

Basic ESM Configuration on BSA

Subscriber management is enabled on BSA in Bridged CO model. A relevant configuration is presented below. SLA and subscriber profiles show the default configurations. The authentication policy appeals to RADIUS server 192.0.2.5. The subscriber identification policy is configured to use DHCP Option 254 to transfer custom attributes (subscriber-id, sla-profile, sub-profile, etc.)

```
A:BSA>config>subscr-mgmt# info
authentication-policy "auth-policy-1" create
password password-1
radius-authentication-server
router "management"
server 1 address 192.0.2.5 secret ALU
exit
include-radius-attribute
circuit-id
remote-id
nas-port-id
nas-identifier
exit
exit
sla-profile "sla-profile-1" create
exit
sla-profile "sla-profile-2" create
exit
sla-profile "sla-profile-3" create
exit
sla-profile "sla-profile-default" create
exit
sub-profile "sub-profile-1" create
exit
sub-profile "sub-profile-default" create
exit
sub-ident-policy "sub-ident-policy-1" create
sub-profile-map
use-direct-map-as-default
exit
sla-profile-map
use-direct-map-as-default
exit
strings-from-option 254
exit
```

The **strings-from-option 254** command is shared in-built dhcp-server of BSR. Using this option, the DHCP server could transmit subscriber identification options such the subscriber-id, sla-profile-string, and sub-profile-string.

BSA/BSR Configuration for Host-1 Operation

The test subscriber has three hosts. Host-1 gets all necessary information from RADIUS server.

Table 3: BSA/BSR Configuration for Host-1 Operation

	BSA (Service/Features)	BSR (Service/Features)
Host-1 ca:00:0c:54:00:08	VPLS-100 <ul style="list-style-type: none"> DHCP proxy server SAP/SDP DHCP snoop Sub-Ident origin through RADIUS Sla/Sub-profiles through RADIUS IP options through RADIUS 	IES-100

In this case BSA takes role of DHCP proxy with DHCP server emulation. DHCP snooping on the SAP must be enabled. Anti-spoofing filters on the SAP must be enabled. An authentication policy must be applied on the SAP.

```

vpls 100 customer 1 create
split-horizon-group "RSHG-1" residential-group create
exit
---snip---
sap 1/1/4:100 split-horizon-group "RSHG-1" create
dhcp
snoop
lease-populate 400
proxy-server
emulated-server 10.0.1.253
no shutdown
exit
no shutdown
exit
authentication-policy "auth-policy-1"
anti-spoof ip-mac
sub-sla-mgmt
def-sub-id string "default-subscriber"
def-sub-profile "sub-profile-default"
def-sla-profile "sla-profile-default"
sub-ident-policy "sub-ident-policy-1"
no shutdown
exit
exit
spoke-sdp 12:100 create
exit
no shutdown
exit

```

On BSR IES-100, the service is configured with a pure IP interface, which plays role of DG for host-1.

```
ies 100 customer 1 create
  interface "int-host-1" create
    address 10.0.1.254/24
    spoke-sdp 21:100 create
  exit
exit
no shutdown
exit
```

BSA/BSR Configuration for Host-2 Operation

The test subscriber has three hosts. Host-2 gets subscriber-id and sla/sub-profiles information from RADIUS server and IP options from DHCP server.

Table 4: BSA/BSR Configuration for Host-2 Operation

	BSA (Service/Features)	BSR (Service/Features)
Host-2 ca:01:08:10:00:08	VPLS-200 <ul style="list-style-type: none">SAP/SDP DHCP snoopSub-Ident origin through RADIUSSla/Sub-profiles through RADIUSIP options through DHCP	IES-200 <ul style="list-style-type: none">DHCP relay

DHCP snooping on the SAP and SDP must be enabled. Anti-spoofing filters on the SAP must be enabled.

```
vpls 200 customer 1 create
  split-horizon-group "RSHG-1" residential-group create
  exit
---snip---
  sap 1/1/4:200 split-horizon-group "RSHG-1" create
    dhcp
      snoop
      lease-populate 400
      no shutdown
    exit
    authentication-policy "auth-policy-1"
    anti-spoof ip-mac
    sub-sla-mgmt
      def-sub-id string "default-subscriber"
      def-sub-profile "sub-profile-default"
      def-sla-profile "sla-profile-default"
      sub-ident-policy "sub-ident-policy-1"
      no shutdown
    exit
  exit
  spoke-sdp 12:200 create
    dhcp
      snoop
    exit
  exit
```

```
no shutdown
exit
```

On BSR IES-200, the service is configured with an IP interface which as the DG for Host-2. DHCP relay must be configured to transform broadcast DHCP discover message into unicast and send it to DHCP server for processing.

```
ies 200 customer 1 create
  interface "int-host-2" create
    address 10.0.2.254/24
    dhcp
      server 192.0.2.4
      trusted
      no shutdown
    exit
  spoke-sdp 21:200 create
  exit
exit
no shutdown
exit
```

BSA/BSR Configuration for Host-3 Operation

The test subscriber has three hosts. Host-3 receives all necessary information from the DHCP server.

Table 5: BSA/BSR Configuration for Host-3 Operation

	BSA (Service/Features)	BSR (Service/Features)
Host-3 ca:02:02:d0:00:08	VPLS-300 * SAP/SDP DHCP snoop * Sub-Ident origin through DHCP * Sla/Sub-profiles through DHCP * IP options through DHCP	IES-300 * DHCP relay

DHCP snooping on the SAP and SDP must be enabled. Anti-spoofing filters on the SAP must be enabled.

```
vpls 300 customer 1 create
  split-horizon-group "RSHG-1" residential-group create
  exit
---snip---
  sap 1/1/4:300 split-horizon-group "RSHG-1" create
    dhcp
      snoop
      lease-populate 400
      no shutdown
    exit
    anti-spoof ip-mac
    sub-sla-mgmt
      def-sub-id string "default-subscriber"
      def-sub-profile "sub-profile-default"
      def-sla-profile "sla-profile-default"
      sub-ident-policy "sub-ident-policy-1"
```

```
        no shutdown
    exit
exit
spoke-sdp 12:300 create
    dhcp
        snoop
    exit
exit
no shutdown
exit
```

On BSR IES-300, the service is configured with IP interface, which plays role of DG for host-3. DHCP relay must be configured to transform broadcast DHCP discover message into unicast and send it to DHCP server for processing.

```
ies 300 customer 1 create
interface "int-host-3" create
    address 10.0.3.254/24
    dhcp
        server 192.0.2.4
        trusted
        no shutdown
    exit
    spoke-sdp 21:300 create
    exit
exit
no shutdown
exit
```

RADIUS Configuration Bridged CO

The username in the RADIUS access request is configurable and could be one of the following formats:

- mac — MAC Source Address of the DHCP DISCOVER message
- circuit-id — Taken from option 82 in the received DHCP message. If no circuit-id can be found, the DHCP-msg is rejected.
- tuple — Concatenation of MAC source address and circuit-ID
- ascii-converted-circuit-id — Identical to circuit-id, but the user name will be sent to the RADIUS server as a string of hex digits
- ascii-converted-tuple — Identical to tuple, but the circuit-id part of the user name will be sent to the RADIUS server as a string of hex digits



Note:

Refer to [IPv4 DHCP Hosts](#) for detailed information about how to use different options.

```
A:BSA>config>subscr-mgmt>auth-plcy# user-name-format
- user-name-format <format> [append domain-name]
- no user-name-format

<format>                : mac|circuit-id|tuple|ascii-converted-circuit-id|
                        ascii-converted-tuple
```

For simplicity, MAC format is used in this guide.

There are two hosts configured in the users file on RADIUS server:

- a:00:0c:54:00:08 — The mac address of host-1 host [VPLS/IES 100]. For host-1 all necessary parameters are returned: subscriber-id, sla/sub-profiles, IP parameters and lease time.
- a:01:08:10:00:08 — The mac address of host-2 host [VPLS/IES 200]. For host-2 only subscriber-id, sla/sub-profiles are returned, while ip parameters and lease time are returned from DHCP server.

```
ca:00:0c:54:00:08 Auth-Type := Local, User-Password == "password-1"
                  Alc-Subsc-ID-Str = "sub-id-1",
                  Alc-Subsc-Prof-Str == "sub-profile-1",
                  Alc-SLA-Prof-Str == "sla-profile-1",
                  Framed-IP-Address = 10.0.1.1,
                  Framed-IP-Netmask = 255.255.255.0,
                  Alc-Default-Router = 10.0.1.254,
                  Session-Timeout = 6000

ca:01:08:10:00:08 Auth-Type := Local, User-Password == "password-1"
                  Alc-Subsc-ID-Str = "sub-id-1",
                  Alc-Subsc-Prof-Str == "sub-profile-1",
                  Alc-SLA-Prof-Str == "sla-profile-2"
```

Local DHCP Server Configuration Bridged CO

In the setup local DHCP server is used with reference to local user database.

```
A:BSR>config>router>dhcp# info
    local-dhcp-server "dhcp-server-1" create
    user-db "user-db-1"
    pool "pool-1" create
    subnet 10.0.2.0/24 create
    exit
    subnet 10.0.3.0/24 create
    exit
    exit
    no shutdown
exit
```



Note:

Subnets must be configured, even if all IP parameters are returned from local user DB. Without this option, DHCP server do not return IP parameters.

The local user database is configured on BSR. Identification is done via MAC address of a host, which is taken from DHCP-Discover message. There are several possibilities to identify DHCP host. **match-list** command is used for this purpose.

```
*A:BSR>config>subscr-mgmt>loc-user-db>dhcp# match-list
- no match-list
- match-list <dhcp-match-type-1> [<dhcpmatch-type-2>...(up to 4 max)]

<dhcp-match-type>      : circuit-id|mac|option60|remote-id|sap-id|service-id|
                        string|system-id
```

There are two hosts configured:

- a:01:08:10:00:08 — mac address of host-2 [VPLS/IES 200]. DHCP returns ip address, subnet mask and default route.
- a:02:02:d0:00:08 — mac address of host-3 [VPLS/IES 300]. DHCP returns all necessary parameters: subscriber-id, sla/sub-profiles and all ip options.

```
A:BSR>config>subscr-mgmt# info
    local-user-db "user-db-1" create
    dhcp
        match-list mac
        host "host-2" create
            host-identification
                mac ca:01:08:10:00:08
            exit
            address 10.0.2.1
            options
                subnet-mask 255.255.255.0
                default-router 10.0.2.254
            exit
            no shutdown
        exit
        host "host-3" create
            host-identification
                mac ca:02:02:d0:00:08
            exit
            address 10.0.3.1
            identification-strings 254 create
                subscriber-id "sub-id-1"
                sla-profile-string "sla-profile-3"
                sub-profile-string "sub-profile-1"
            exit
            options
                subnet-mask 255.255.255.0
                default-router 10.0.3.254
            exit
            no shutdown
        exit
    exit
    no shutdown
exit
```

Setup Procedures and Debugging

Subscriber/Host Verification

The initialization of all active subscribers and hosts can be shown using the **show service active-subscribers** command. Different options can be used to filter the output of the command.

```
A:BSA# show service active-subscribers
=====
Active Subscribers
=====
-----
Subscriber sub-id-1 (sub-profile-1)
-----
-----
(1) SLA Profile Instance sap:1/1/4:100 - sla:sla-profile-1
```

IP Address	MAC Address	PPPoE-SID Origin
10.0.1.1	ca:00:0c:54:00:08	N/A DHCP
(2) SLA Profile Instance sap:1/1/4:200 - sla:sla-profile-2		
IP Address	MAC Address	PPPoE-SID Origin
10.0.2.1	ca:01:08:10:00:08	N/A DHCP
(3) SLA Profile Instance sap:1/1/4:300 - sla:sla-profile-3		
IP Address	MAC Address	PPPoE-SID Origin
10.0.3.1	ca:02:02:d0:00:08	N/A DHCP
Number of active subscribers : 1		
=====		
A:BSA#		


Hierarchy of subscriber hosts is represented in a convenient form using following command.

```
A:BSA# show service active-subscribers hierarchy
=====
Active Subscriber hierarchy
=====
-- sub-id-1 (sub-profile-1)
|
|-- sap:1/1/4:100 - sla:sla-profile-1
|   |
|   |-- 10.0.1.1 - ca:00:0c:54:00:08 - N/A (DHCP)
|   |
|
|-- sap:1/1/4:200 - sla:sla-profile-2
|   |
|   |-- 10.0.2.1 - ca:01:08:10:00:08 - N/A (DHCP)
|   |
|
|-- sap:1/1/4:300 - sla:sla-profile-3
|   |
|   |-- 10.0.3.1 - ca:02:02:d0:00:08 - N/A (DHCP)
|   |
|
```

Host-1 Setup Debug

The Host-1 setup process is shown in [Figure 14: Host-1 Setup Process](#).

Figure 14: Host-1 Setup Process



Host-1 sends DHCP discover message in VLAN 100 to BSA. BSA plays role of DHCP proxy server and transforms DHCP discover into RADIUS access-request message. After receiving RADIUS access-accept BSA transforms it to DHCP Ack message. Session setup process could be represented using debug commands:

```
A:BSA# debug service id 100 dhcp mode egr-ingr-and-dropped
A:BSA# debug service id 100 dhcp detail-level medium
A:BSA# debug radius detail
18 2009/12/15 06:31:56.63 UTC MINOR: DEBUG #2001 Base SVCMMGR
"SVCMMGR: RX DHCP Packet
  VPLS 100, SAP 1/1/4:100

  BootRequest to UDP port 67
  ciaddr: 0.0.0.0      yiaddr: 0.0.0.0
  siaddr: 0.0.0.0      giaddr: 0.0.0.0
  chaddr: ca:00:0c:54:00:08  xid: 0xd42

  DHCP options:
  [53] Message type: Discover
---snip---
"
19 2009/12/15 06:31:56.63 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Access-Request
  user ca:00:0c:54:00:08  policy auth-policy-1"
20 2009/12/15 06:31:56.63 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Transmit
  Access-Request(1) 192.0.2.5:1812 id 69  len 85
  USER NAME [1] 17 ca:00:0c:54:00:08
  PASSWORD [2] 16 lkhSVrFePQ0A0Xc4ZyMwMk
  NAS IP ADDRESS [4] 4 192.0.2.1
  NAS PORT TYPE [61] 4 Ethernet(15)
  NAS PORT ID [87] 9 1/1/4:100
  NAS IDENTIFIER [32] 3 BSA
"
21 2009/12/15 06:31:56.73 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Receive
```

```
Access-Accept(2) id 69 len 108 from 138.203.18.79:1812
VSA [26] 10 Alcatel(6527)
  SUBSC ID STR [11] 8 sub-id-1
VSA [26] 15 Alcatel(6527)
  SUBSC PROF STR [12] 13 sub-profile-1
VSA [26] 15 Alcatel(6527)
  SLA PROF STR [13] 13 sla-profile-1
FRAMED IP ADDRESS [8] 4 10.0.1.1
FRAMED IP NETMASK [9] 4 255.255.255.0
VSA [26] 6 Alcatel(6527)
  DEFAULT ROUTER [18] 4 10.0.1.254
  SESSION TIMEOUT [27] 4 6000
"
22 2009/12/15 06:31:56.73 UTC MINOR: DEBUG #2001 Base SVCMMGR
"SVCMMGR: TX DHCP Packet
VPLS 100, SAP 1/1/4:100

BootReply to UDP port 68
ciaddr: 0.0.0.0          yiaddr: 10.0.1.1
siaddr: 10.0.1.253      giaddr: 0.0.0.0
chaddr: ca:00:0c:54:00:08  xid: 0xd42

DHCP options:
[53] Message type: Offer
---snip---
"
23 2009/12/15 06:31:57.57 UTC MINOR: DEBUG #2001 Base SVCMMGR
"SVCMMGR: RX DHCP Packet
VPLS 100, SAP 1/1/4:100

BootRequest to UDP port 67
ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
siaddr: 0.0.0.0          giaddr: 0.0.0.0
chaddr: ca:00:0c:54:00:08  xid: 0xd42

DHCP options:
[53] Message type: Request
---snip---
"
24 2009/12/15 06:31:57.57 UTC MINOR: DEBUG #2001 Base SVCMMGR
"SVCMMGR: TX DHCP Packet
VPLS 100, SAP 1/1/4:100

BootReply to UDP port 68
ciaddr: 0.0.0.0          yiaddr: 10.0.1.1
siaddr: 10.0.1.253      giaddr: 0.0.0.0
chaddr: ca:00:0c:54:00:08  xid: 0xd42

DHCP options:
[53] Message type: Ack
---snip---
```

The number of snooped/forwarded/dropped/proxied DHCP packets can be checked using the **show service id 100 dhcp statistics** command.

```
A:BSA# show service id 100 dhcp statistics
=====
DHCP Statistics, service 100
=====
Client Packets Snooped           : 2
Client Packets Forwarded         : 0
Client Packets Dropped           : 0
```

```
Client Packets Proxied (RADIUS)      : 2
Client Packets Proxied (Lease-Split) : 0
Server Packets Snooped              : 0
Server Packets Forwarded            : 0
Server Packets Dropped              : 0
DHCP RELEASEs Spoofed              : 0
DHCP FORCERENEWs Spoofed           : 0
=====
A:BSA#
```

Connectivity of Host-1 could be checked with the **show service id 100 subscriber-hosts** command. Different options can be used to filter output of the command.

```
A:BSA# show service id 100 subscriber-hosts detail
=====
Subscriber Host table
=====
Sap      IP Address  MAC Address  PPPoE-SID  Origin
Subscriber
-----
1/1/4:100 10.0.1.1    ca:00:0c:54:00:08 N/A        DHCP
sub-id-1
-----
Sub Profile      : sub-profile-1
SLA Profile      : sla-profile-1
App Profile      : N/A
-----
Number of subscriber hosts : 1
```

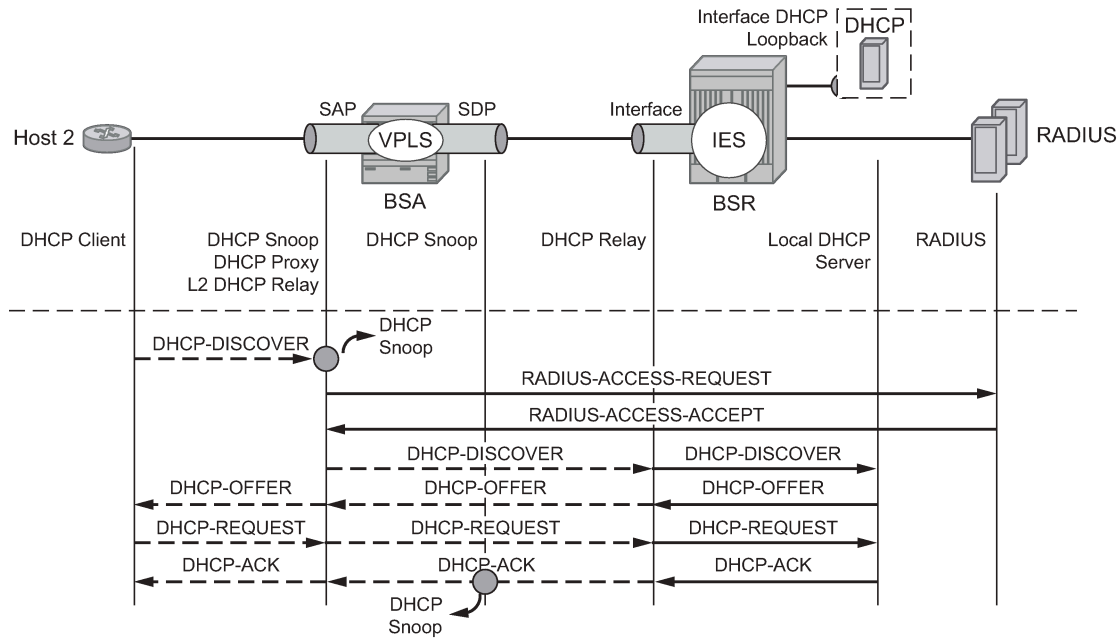
The DHCP lease state can be checked with the **show service id 100 dhcp lease-state** command. Different options can be used to filter output of a command.

```
A:BSA# show service id 100 dhcp lease-state detail
=====
DHCP lease states for service 100
=====
Service ID      : 100
IP Address      : 10.0.1.1
Client HW Address : ca:00:0c:54:00:08
SAP             : 1/1/4:100
Remaining Lifetime : 01h33m41s
Persistence Key  : N/A
Sub-Ident       : "sub-id-1"
Sub-Profile-String : "sub-profile-1"
SLA-Profile-String : "sla-profile-1"
---snip---
Sub-Ident origin : Radius
Strings origin   : Radius
Lease Info origin : Radius
---snip---
Radius User-Name : "ca:00:0c:54:00:08"
-----
Number of lease states : 1
=====
A:BSA#
```

Host-2 Setup Debug

Host-1 setup process is displayed in [Figure 15: Host-2 Setup Process](#).

Figure 15: Host-2 Setup Process



OSSG444

Host-2 sends DHCP discover message in VLAN 200 to BSA. Host-2 is authenticated through RADIUS and gets subscriber-id, sla/sub-profiles. DHCP Discover message is flooded in VPLS service and reaches IP interface on BSA, where DHCP relay is configured. Session setup process could be represented using debug commands:

```
A:BSA# debug service id 200 dhcp mode egr-ingr-and-dropped
A:BSA# debug service id 200 dhcp detail-level medium
A:BSA#
*A:BSA#
18 2009/12/15 13:00:36.28 UTC MINOR: DEBUG #2001 Base SVCMMGR
"SVCMMGR: RX DHCP Packet
  VPLS 200, SAP 1/1/4:200

  BootRequest to UDP port 67
  ciaddr: 0.0.0.0      yiaddr: 0.0.0.0
  siaddr: 0.0.0.0      giaddr: 0.0.0.0
  chaddr: ca:01:08:10:00:08  xid: 0xfc8

  DHCP options:
  [53] Message type: Discover
---snip---
```

```
19 2009/12/15 13:00:36.28 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Access-Request
  user ca:01:08:10:00:08  policy auth-policy-1"
```

```
20 2009/12/15 13:00:36.28 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Transmit
  Access-Request(1) 192.0.2.5:1812 id 80  len 85
  USER NAME [1] 17 ca:01:08:10:00:08
  PASSWORD [2] 16 .czdppt/0qAsqKqStbvnV.
  NAS IP ADDRESS [4] 4 192.0.2.1
  NAS PORT TYPE [61] 4 Ethernet(15)
```

```
NAS PORT ID [87] 9 1/1/4:200
NAS IDENTIFIER [32] 3 BSA
"

21 2009/12/15 13:00:36.34 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Receive
Access-Accept(2) id 80 len 78 from 138.203.18.79:1812
VSA [26] 10 Alcatel(6527)
SUBSC ID STR [11] 8 sub-id-1
VSA [26] 15 Alcatel(6527)
SUBSC PROF STR [12] 13 sub-profile-1
VSA [26] 15 Alcatel(6527)
SLA PROF STR [13] 13 sla-profile-2
"

22 2009/12/15 13:00:36.34 UTC MINOR: DEBUG #2001 Base SVCMMGR
"SVCMMGR: TX DHCP Packet
flooding in VPLS 200

BootRequest to UDP port 67
ciaddr: 0.0.0.0 yiaddr: 0.0.0.0
siaddr: 0.0.0.0 giaddr: 0.0.0.0
chaddr: ca:01:08:10:00:08 xid: 0xfc8

DHCP options:
[53] Message type: Discover
---snip---
"

23 2009/12/15 13:00:36.35 UTC MINOR: DEBUG #2001 Base SVCMMGR
"SVCMMGR: RX DHCP Packet
VPLS 200, spoke-sdp 12:200

BootReply to UDP port 68
ciaddr: 0.0.0.0 yiaddr: 10.0.2.1
siaddr: 192.0.2.4 giaddr: 0.0.0.0
chaddr: ca:01:08:10:00:08 xid: 0xfc8

DHCP options:
[53] Message type: Offer
---snip---
"

24 2009/12/15 13:00:36.34 UTC MINOR: DEBUG #2001 Base SVCMMGR
"SVCMMGR: TX DHCP Packet
VPLS 200, SAP 1/1/4:200

BootReply to UDP port 68
ciaddr: 0.0.0.0 yiaddr: 10.0.2.1
siaddr: 192.0.2.4 giaddr: 0.0.0.0
chaddr: ca:01:08:10:00:08 xid: 0xfc8

DHCP options:
[53] Message type: Offer
---snip---
"

25 2009/12/15 13:00:36.46 UTC MINOR: DEBUG #2001 Base SVCMMGR
"SVCMMGR: RX DHCP Packet
VPLS 200, SAP 1/1/4:200

BootRequest to UDP port 67
ciaddr: 0.0.0.0 yiaddr: 0.0.0.0
siaddr: 0.0.0.0 giaddr: 0.0.0.0
chaddr: ca:01:08:10:00:08 xid: 0xfc8
```

```
DHCP options:
[53] Message type: Request
---snip---
"

26 2009/12/15 13:00:36.46 UTC MINOR: DEBUG #2001 Base SVCNMR
"SVCNMR: TX DHCP Packet
    flooding in VPLS 200

    BootRequest to UDP port 67
    ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
    siaddr: 0.0.0.0          giaddr: 0.0.0.0
    chaddr: ca:01:08:10:00:08  xid: 0xfc8

    DHCP options:
    [53] Message type: Request
    ---snip---
    "

27 2009/12/15 13:00:36.47 UTC MINOR: DEBUG #2001 Base SVCNMR
"SVCNMR: RX DHCP Packet
    VPLS 200, spoke-sdp 12:200

    BootReply to UDP port 68
    ciaddr: 0.0.0.0          yiaddr: 10.0.2.1
    siaddr: 192.0.2.4        giaddr: 0.0.0.0
    chaddr: ca:01:08:10:00:08  xid: 0xfc8

    DHCP options:
    [53] Message type: Ack
    ---snip---
    "

28 2009/12/15 13:00:36.46 UTC MINOR: DEBUG #2001 Base SVCNMR
"SVCNMR: TX DHCP Packet
    VPLS 200, SAP 1/1/4:200

    BootReply to UDP port 68
    ciaddr: 0.0.0.0          yiaddr: 10.0.2.1
    siaddr: 192.0.2.4        giaddr: 0.0.0.0
    chaddr: ca:01:08:10:00:08  xid: 0xfc8

    DHCP options:
    [53] Message type: Ack
    ---snip---
    "

DHCP relay is enabled in service IES-200 on BSR.

A:BSR# debug router ip dhcp mode egr-ingr-and-dropped
A:BSR#
*A:BSR#
17 2009/12/15 13:00:36.34 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 6 (int-host-2),
    received DHCP Boot Request on Interface int-host-2 (1/1/2) Port 67

H/W Type: Ethernet(10Mb)  H/W Address Length: 6
ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
siaddr: 0.0.0.0          giaddr: 0.0.0.0
chaddr: ca:01:08:10:00:08  xid: 0xfc8

DHCP options:
```

```
[53] Message type: Discover
---snip---
"

18 2009/12/15 13:00:36.34 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base),
    transmitted DHCP Boot Request to 192.0.2.4 Port 67

    H/W Type: Ethernet(10Mb)  H/W Address Length: 6
    ciaddr: 0.0.0.0            yiaddr: 0.0.0.0
    siaddr: 0.0.0.0            giaddr: 10.0.2.254
    chaddr: ca:01:08:10:00:08  xid: 0xfc8

    DHCP options:
    [53] Message type: Discover
---snip---
"

19 2009/12/15 13:00:36.35 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base),
    received DHCP Boot Reply on 192.0.2.4 Port 67

    H/W Type: Ethernet(10Mb)  H/W Address Length: 6
    ciaddr: 0.0.0.0            yiaddr: 10.0.2.1
    siaddr: 192.0.2.4          giaddr: 10.0.2.254
    chaddr: ca:01:08:10:00:08  xid: 0xfc8

    DHCP options:
    [53] Message type: Offer
---snip---
"

20 2009/12/15 13:00:36.35 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 6 (int-host-2),
    transmitted DHCP Boot Reply to Interface int-host-2 (spoke-21:200) Port 68

    H/W Type: Ethernet(10Mb)  H/W Address Length: 6
    ciaddr: 0.0.0.0            yiaddr: 10.0.2.1
    siaddr: 192.0.2.4          giaddr: 0.0.0.0
    chaddr: ca:01:08:10:00:08  xid: 0xfc8

    DHCP options:
    [53] Message type: Offer
---snip---
"

21 2009/12/15 13:00:36.47 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 6 (int-host-2),
    received DHCP Boot Request on Interface int-host-2 (1/1/2) Port 67

    H/W Type: Ethernet(10Mb)  H/W Address Length: 6
    ciaddr: 0.0.0.0            yiaddr: 0.0.0.0
    siaddr: 0.0.0.0            giaddr: 0.0.0.0
    chaddr: ca:01:08:10:00:08  xid: 0xfc8

    DHCP options:
    [53] Message type: Request
---snip---
"
```

```
22 2009/12/15 13:00:36.47 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base),
  transmitted DHCP Boot Request to 192.0.2.4 Port 67

  H/W Type: Ethernet(10Mb)  H/W Address Length: 6
  ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
  siaddr: 0.0.0.0           giaddr: 10.0.2.254
  chaddr: ca:01:08:10:00:08  xid: 0xfc8

  DHCP options:
  [53] Message type: Request
---snip---
"

23 2009/12/15 13:00:36.47 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base),
  received DHCP Boot Reply on 192.0.2.4 Port 67

  H/W Type: Ethernet(10Mb)  H/W Address Length: 6
  ciaddr: 0.0.0.0           yiaddr: 10.0.2.1
  siaddr: 192.0.2.4         giaddr: 10.0.2.254
  chaddr: ca:01:08:10:00:08  xid: 0xfc8

  DHCP options:
  [53] Message type: Ack
---snip---
"

24 2009/12/15 13:00:36.47 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 6 (int-host-2),
  transmitted DHCP Boot Reply to Interface int-host-2 (spoke-21:200) Port 68

  H/W Type: Ethernet(10Mb)  H/W Address Length: 6
  ciaddr: 0.0.0.0           yiaddr: 10.0.2.1
  siaddr: 192.0.2.4         giaddr: 0.0.0.0
  chaddr: ca:01:08:10:00:08  xid: 0xfc8

  DHCP options:
  [53] Message type: Ack
---snip---
"
```

The number of snooped/forwarded/dropped/proxied DHCP packets could be checked using the **show service id 200 dhcp statistics** command.

```
A:BSA# show service id 200 dhcp statistics
=====
DHCP Statistics, service 200
=====
Client Packets Snooped           : 2
Client Packets Forwarded         : 2
Client Packets Dropped           : 0
Client Packets Proxied (RADIUS)  : 0
Client Packets Proxied (Lease-Split) : 0
Server Packets Snooped           : 2
Server Packets Forwarded         : 2
Server Packets Dropped           : 0
DHCP RELEASEs Spoofed           : 0
DHCP FORCERENEWs Spoofed        : 0
=====
```


A:BSA#

The connectivity of Host-2 can be verified with the **show service id 200 subscriber-hosts** command. Different options can be used to filter output of the command.

```
A:BSA# show service id 200 subscriber-hosts detail
=====
Subscriber Host table
=====
Sap      IP Address  MAC Address  PPPoE-SID  Origin
Subscriber
-----
1/1/4:200  10.0.2.1    ca:01:08:10:00:08  N/A        DHCP
sub-id-1
-----
Sub Profile      : sub-profile-1
SLA Profile      : sla-profile-2
App Profile      : N/A
-----
Number of subscriber hosts : 1
=====
A:BSA#
```


DHCP lease state can be verified with the **show service id 200 dhcp lease-state** command. Different options can be used to filter output of the command.

```
A:BSA# show service id 200 dhcp lease-state detail
=====
DHCP lease states for service 200
=====
Service ID      : 200
IP Address      : 10.0.2.1
Client HW Address : ca:01:08:10:00:08
SAP             : 1/1/4:200
Remaining Lifetime : 09d23h44m
Persistence Key  : N/A
Sub-Ident       : "sub-id-1"
Sub-Profile-String : "sub-profile-1"
SLA-Profile-String : "sla-profile-2"
---snip---
Sub-Ident origin : Radius
Strings origin   : Radius
Lease Info origin : DHCP
---snip---
Radius User-Name : "ca:01:08:10:00:08"
-----
Number of lease states : 1
=====
A:BSA#
```

Host-3 Setup Debug

The Host-3 setup process is presented in [Figure 16: Host-3 Setup Process](#).

Figure 16: Host-3 Setup Process



OSSG445

Host-3 sends DHCP a discover message in VLAN 300 to BSA. Host-3 receives all parameters from the DHCP server using pre-configured Option 254. A DHCP discover message is flooded in VPLS service and reaches IP interface on BSA where DHCP relay is configured. The session setup process can be represented using debug commands.

```
A:BSA# debug service id 300 dhcp mode egr-ingr-and-dropped
A:BSA# debug service id 300 dhcp detail-level medium
*A:BSA#
33 2009/12/15 13:02:34.39 UTC MINOR: DEBUG #2001 Base SVCMMGR
"SVCMMGR: RX DHCP Packet
  VPLS 300, SAP 1/1/4:300

  BootRequest to UDP port 67
  ciaddr: 0.0.0.0      yiaddr: 0.0.0.0
  siaddr: 0.0.0.0      giaddr: 0.0.0.0
  chaddr: ca:02:02:d0:00:08  xid: 0x2a6

  DHCP options:
  [53] Message type: Discover
  ---snip---
  "
34 2009/12/15 13:02:34.38 UTC MINOR: DEBUG #2001 Base SVCMMGR
"SVCMMGR: TX DHCP Packet
  flooding in VPLS 300

  BootRequest to UDP port 67
  ciaddr: 0.0.0.0      yiaddr: 0.0.0.0
  siaddr: 0.0.0.0      giaddr: 0.0.0.0
  chaddr: ca:02:02:d0:00:08  xid: 0x2a6

  DHCP options:
  [53] Message type: Discover
  ---snip---
  "
35 2009/12/15 13:02:34.40 UTC MINOR: DEBUG #2001 Base SVCMMGR
"SVCMMGR: RX DHCP Packet
  VPLS 300, spoke-sdp 12:300
```

```
BootReply to UDP port 68
ciaddr: 0.0.0.0          yiaddr: 10.0.3.1
siaddr: 192.0.2.4        giaddr: 0.0.0.0
chaddr: ca:02:02:d0:00:08  xid: 0x2a6

DHCP options:
[53] Message type: Offer
---snip---
"
36 2009/12/15 13:02:34.40 UTC MINOR: DEBUG #2001 Base SVCNMR
"SVCNMR: TX DHCP Packet
VPLS 300, SAP 1/1/4:300

BootReply to UDP port 68
ciaddr: 0.0.0.0          yiaddr: 10.0.3.1
siaddr: 192.0.2.4        giaddr: 0.0.0.0
chaddr: ca:02:02:d0:00:08  xid: 0x2a6

DHCP options:
[53] Message type: Offer
---snip---
"
37 2009/12/15 13:02:34.52 UTC MINOR: DEBUG #2001 Base SVCNMR
"SVCNMR: RX DHCP Packet
VPLS 300, SAP 1/1/4:300

BootRequest to UDP port 67
ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
siaddr: 0.0.0.0          giaddr: 0.0.0.0
chaddr: ca:02:02:d0:00:08  xid: 0x2a6

DHCP options:
[53] Message type: Request
---snip---
"
38 2009/12/15 13:02:34.52 UTC MINOR: DEBUG #2001 Base SVCNMR
"SVCNMR: TX DHCP Packet
flooding in VPLS 300

BootRequest to UDP port 67
ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
siaddr: 0.0.0.0          giaddr: 0.0.0.0
chaddr: ca:02:02:d0:00:08  xid: 0x2a6

DHCP options:
[53] Message type: Request
---snip---
"
39 2009/12/15 13:02:34.53 UTC MINOR: DEBUG #2001 Base SVCNMR
"SVCNMR: RX DHCP Packet
VPLS 300, spoke-sdp 12:300

BootReply to UDP port 68
ciaddr: 0.0.0.0          yiaddr: 10.0.3.1
siaddr: 192.0.2.4        giaddr: 0.0.0.0
chaddr: ca:02:02:d0:00:08  xid: 0x2a6

DHCP options:
[53] Message type: Ack
---snip---
"
40 2009/12/15 13:02:34.53 UTC MINOR: DEBUG #2001 Base SVCNMR
"SVCNMR: TX DHCP Packet
```

```
VPLS 300, SAP 1/1/4:300

BootReply to UDP port 68
ciaddr: 0.0.0.0      yiaddr: 10.0.3.1
siaddr: 192.0.2.4    giaddr: 0.0.0.0
chaddr: ca:02:02:d0:00:08  xid: 0x2a6

DHCP options:
[53] Message type: Ack
---snip---
```

DHCP relay is enabled in service IES-300 on the BSR.

```
A:BSR# debug router ip dhcp mode egr-ingr-and-dropped
*A:BSR#
29 2009/12/15 13:02:34.39 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 7 (int-VoD),
  received DHCP Boot Request on Interface int-VoD (1/1/2) Port 67

H/W Type: Ethernet(10Mb)  H/W Address Length: 6
ciaddr: 0.0.0.0      yiaddr: 0.0.0.0
siaddr: 0.0.0.0      giaddr: 0.0.0.0
chaddr: ca:02:02:d0:00:08  xid: 0x2a6

DHCP options:
[53] Message type: Discover
---snip---
30 2009/12/15 13:02:34.39 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base),
  transmitted DHCP Boot Request to 192.0.2.4 Port 67

H/W Type: Ethernet(10Mb)  H/W Address Length: 6
ciaddr: 0.0.0.0      yiaddr: 0.0.0.0
siaddr: 0.0.0.0      giaddr: 10.0.3.254
chaddr: ca:02:02:d0:00:08  xid: 0x2a6

DHCP options:
[53] Message type: Discover
---snip---
"
31 2009/12/15 13:02:34.39 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base),
  received DHCP Boot Reply on 192.0.2.4 Port 67

H/W Type: Ethernet(10Mb)  H/W Address Length: 6
ciaddr: 0.0.0.0      yiaddr: 10.0.3.1
siaddr: 192.0.2.4      giaddr: 10.0.3.254
chaddr: ca:02:02:d0:00:08  xid: 0x2a6

DHCP options:
[53] Message type: Offer
---snip---
"
32 2009/12/15 13:02:34.39 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 7 (int-VoD),
  transmitted DHCP Boot Reply to Interface int-VoD (spoke-21:300) Port 68

H/W Type: Ethernet(10Mb)  H/W Address Length: 6
ciaddr: 0.0.0.0      yiaddr: 10.0.3.1
```

```
siaddr: 192.0.2.4      giaddr: 0.0.0.0
chaddr: ca:02:02:d0:00:08  xid: 0x2a6

DHCP options:
[53] Message type: Offer
---snip---
"
33 2009/12/15 13:02:34.53 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 7 (int-VoD),
    received DHCP Boot Request on Interface int-VoD (1/1/2) Port 67

H/W Type: Ethernet(10Mb) H/W Address Length: 6
ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
siaddr: 0.0.0.0          giaddr: 0.0.0.0
chaddr: ca:02:02:d0:00:08  xid: 0x2a6

DHCP options:
[53] Message type: Request
---snip---
"
34 2009/12/15 13:02:34.53 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base),
    transmitted DHCP Boot Request to 192.0.2.4 Port 67

H/W Type: Ethernet(10Mb) H/W Address Length: 6
ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
siaddr: 0.0.0.0          giaddr: 10.0.3.254
chaddr: ca:02:02:d0:00:08  xid: 0x2a6

DHCP options:
[53] Message type: Request
---snip---
"
35 2009/12/15 13:02:34.53 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base),
    received DHCP Boot Reply on 192.0.2.4 Port 67

H/W Type: Ethernet(10Mb) H/W Address Length: 6
ciaddr: 0.0.0.0          yiaddr: 10.0.3.1
siaddr: 192.0.2.4        giaddr: 10.0.3.254
chaddr: ca:02:02:d0:00:08  xid: 0x2a6

DHCP options:
[53] Message type: Ack
---snip---
36 2009/12/15 13:02:34.53 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 7 (int-VoD),
    transmitted DHCP Boot Reply to Interface int-VoD (spoke-21:300) Port 68

H/W Type: Ethernet(10Mb) H/W Address Length: 6
ciaddr: 0.0.0.0          yiaddr: 10.0.3.1
siaddr: 192.0.2.4        giaddr: 0.0.0.0
chaddr: ca:02:02:d0:00:08  xid: 0x2a6

DHCP options:
[53] Message type: Ack
---snip---
"
```

The number of snooped/forwarded/dropped/proxied DHCP packets can be verified with the using **show service id 300 dhcp statistics** command.

```
A:BSA# show service id 300 dhcp statistics
=====
DHCP Statistics, service 300
=====
Client Packets Snooped           : 2
Client Packets Forwarded         : 2
Client Packets Dropped           : 0
Client Packets Proxied (RADIUS)  : 0
Client Packets Proxied (Lease-Split) : 0
Server Packets Snooped           : 2
Server Packets Forwarded         : 2
Server Packets Dropped           : 0
DHCP RELEASEs Spoofed           : 0
DHCP FORCERENEWs Spoofed        : 0
=====
A:BSA#
```

The connectivity of Host-3 can be verified with the **show service id 300 subscriber-hosts** command. Different options can be used to filter output of a command.

```
A:BSA# show service id 300 subscriber-hosts detail
=====
Subscriber Host table
=====
Sap      IP Address  MAC Address  PPPoE-SID Origin
Subscriber
-----
1/1/4:300 10.0.3.1    ca:02:02:d0:00:08 N/A      DHCP
sub-id-1
-----
Sub Profile      : sub-profile-1
SLA Profile      : sla-profile-3
App Profile      : N/A
-----
Number of subscriber hosts : 1
=====
A:BSA#
```

The DHCP lease state could be checked with the **show service id 300 dhcp lease-state** command. Different options can be used to filter output of a command.

```
A:BSA# show service id 300 dhcp lease-state detail
=====
DHCP lease states for service 300
=====
Service ID      : 300
IP Address      : 10.0.3.1
Client HW Address : ca:02:02:d0:00:08
SAP             : 1/1/4:300
Remaining Lifetime : 09d23h52m
Persistence Key  : N/A
Sub-Ident       : "sub-id-1"
Sub-Profile-String : "sub-profile-1"
SLA-Profile-String : "sla-profile-3"
---snip---
Sub-Ident origin : DHCP
Strings origin   : DHCP
Lease Info origin : DHCP
```

```
---snip---
Radius User-Name      : ""
-----
Number of lease states : 1
=====
A:BSA#
```

Advanced Topics

Limiting Number of Subscribers

This topic is discussed in DHCP hosts. Refer to [IPv4 DHCP Hosts](#) for detailed information.

```
vpls 100 customer 1 create
---snip---
sap 1/1/4:100 split-horizon-group "RSHG-1" create
---snip---
sub-sla-mgmt
---snip---
multi-sub-sap 2
```

Limiting Number of Lease States

This topic is discussed in DHCP hosts. Refer to [IPv4 DHCP Hosts](#) for detailed information.

```
vpls 100 customer 1 create
---snip---
sap 1/1/4:100 split-horizon-group "RSHG-1" create
  dhcp
    lease-populate 400
```

Limiting Number of Host Per SLA-Profile

This topic is discussed in DHCP hosts. Refer to [IPv4 DHCP Hosts](#) for detailed information.

```
subscriber-mgmt
  sla-profile "sla-profile-1" create
    host-limit 1 [remove-oldest]
```

Subscriber Host Connectivity Verification

This topic is discussed in DHCP hosts. Refer to [IPv4 DHCP Hosts](#) for detailed information.

```
vpls 100 customer 1 create
  sap 1/1/4:100 split-horizon-group "RSHG-1" create
    host-connectivity-verify source-ip 10.1.0.253 source-mac 1e:54:ff:00:00:00
  interval 1 action remove

A:BSA# show service id 100 host-connectivity-verify statistics
=====
```

Host connectivity check statistics					
Svc Id	SapId/SdpId	DestIp Address	Timestamp last-reply/conn-lost	Time since Reply/Lost	Oper State
100	1/1/4:100	10.0.1.1	12/15/2009 09:04:06	0d 00:00:11	Up
1 host-connectivity states : 1 Up / 0 Down / 0 Retry pending					
A:BSA#					

Lease Split

This topic is discussed in DHCP hosts. Refer to [IPv4 DHCP Hosts](#) for detailed information.

```
vpls 100 customer 1 create
---snip---
sap 1/1/4:100 split-horizon-group "RSHG-1" create
    dhcp
        proxy-server
            lease-time hrs 1
```

DHCP Option 82

This topic is discussed in DHCP hosts. Refer to [IPv4 DHCP Hosts](#) for detailed information.

Conclusion

This note provides configuration and troubleshooting commands for Bridged CO model.

DHCP Server Failover States

This chapter describes DHCP server failover states.

Topics in this chapter include:

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

Applicability

This chapter is applicable to SR OS routers and is based on SR OS Release 14.0.R7.

Overview

A common way to maintain DHCP service during a partial power loss or partial network outage is to provide DHCP server redundancy, where two DHCP servers in the network serve a common set of subnets. Failover is a mechanism where the second server takes the role of the first server in case of a failure or a planned network outage, thereby providing a backup.

Failover requires a pair of redundant servers, and IP address assignment continuity is ensured in case of a failure of one of the servers, while at same time preventing address duplication. This contributes to a higher availability of service in the network.

Failover can be performed by an operator taking manual actions, but in most networks failover is usually performed automatically, relying on failure detection mechanisms that can trigger the activation of the second server.

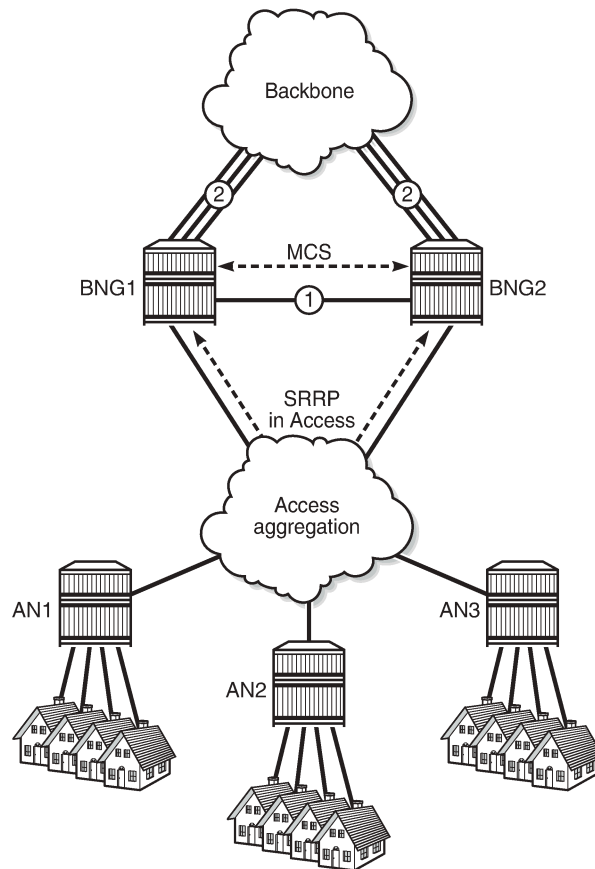
For the peers of a redundant pair to take over the role of each other, they must have the same view of the network and subnets they need to serve, in terms of:

- the definition of the common subnets
- the leases already assigned in these subnets

The definitions of the subnets to be shared are synchronized through configuration, whereas the leases assigned by both DHCP peers are synchronized with each other through Multi Chassis Synchronization (MCS). MCS is also used to detect communication failures between the DHCP servers, but MCS cannot detect whether the cause is a link failure or a server failure.

[Figure 17: General Redundancy Model](#) shows the general redundancy model, where clients connected to the Layer 2 access nodes AN1, AN2, and AN3 get their addresses from the DHCP servers located in BNG1 and BNG2, via relay agents that are also located in BNG1 and BNG2. Access network redundancy can be supported through the Subscriber Routed Redundancy protocol (SRRP) or through Multi Chassis Link Aggregation (MC-LAG) in combination with SRRP.

Figure 17: General Redundancy Model




26352

Figure 18: DHCP Relay Agent and Server Redundancy Model shows the DHCP relay agent and server redundancy model, where the DHCP relay agents R1 and R2 and the DHCP servers S1 and S2 are situated in different nodes in the network. DHCP requests received from the access network are relayed onto S1 and S2 via R1 and R2. MCS is used for synchronizing the lease database between servers S1 and S2. MCS can also be used for synchronizing the subscriber management information between R1 and R2, but that is out of the scope of this chapter.

MCS typically runs over a direct link connecting the two peers of a pair (scenario 1 in [Figure 17: General Redundancy Model](#)), but can also run over backbone links if no direct link is present (scenario 2 in [Figure 17: General Redundancy Model](#)). Regardless of the scenario, this link is referred to as the intercommunication link (ICL), and it should be well protected with multiple underlying physical paths.

DHCP server failover relies on the detection of a failure of the ICL. This link should be disjoint from the access links toward the DHCP clients.

Figure 18: DHCP Relay Agent and Server Redundancy Model



26353

DHCP server failover requires the nodes of a failover pair to have their date and time synchronized. This is commonly implemented using the network time protocol (NTP).

In the configuration section of this chapter, the local-remote deployment model is used, using a single relay agent and two DHCPv4 servers.

For basic DHCPv4 server configuration, see the [DHCPv4 Server Basics](#) chapter.

DHCP Server Failover and Address Management

For DHCP servers to support failover, the redundant servers need to share a set of subnets and address ranges so that one can take the role of the other in case of a failure, at the same time avoiding double allocations.

Following models are supported to achieve these requirements; see [Figure 19: Access-Driven and Local-Remote Model](#):

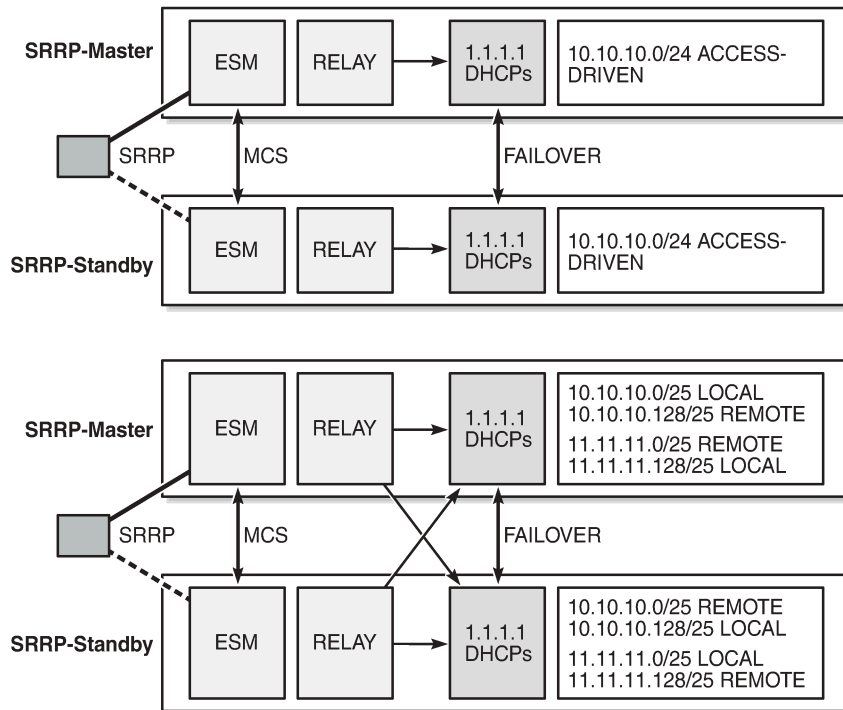
- access-driven model
- local-remote model



Note:

An unsupported local-local model is similar to the access-driven model. They behave the same in terms of their failover, but different in terms of error handling; the local-local model can emit erroneous traps, whereas the access-driven model does not.

Figure 19: Access-Driven and Local-Remote Model



26354

In the access-driven model, the address ranges on both DHCP servers are defined as access-driven. As for the topology, the relay agents are configured to relay the messages to and from one server only, and Nokia recommends that both DHCP servers use the same interface address. This can be achieved by hosting the ESM, relay, and DHCP server functionalities in the same router.

In the local-remote model, address ranges declared as local on one peer must be declared as remote on the other peer, and vice versa. As for the topology, the relay agents are configured to relay the messages to and from both servers, and do not need to be hosted by the same routers as the DHCP servers.

Avoiding Double Allocations

To avoid double allocations in the access-driven model, only one path should be active out of the access network toward the relay agent and the accompanying DHCP server. This is achieved through SRRP or MC-LAG in combination with SRRP in the access network. The relay agent must relay the messages to and from one server only, so this model is effectively an active-standby model.

To avoid double allocations in the local-remote model, where two paths from the access network to the DHCP servers exist and where two DHCP servers work in parallel, the following rules apply:

- Addresses from the local (and access-driven) ranges can always be allocated, even when there is a communication failure between the peer servers, or when the peer server is down.
- Addresses from remote ranges should only be allocated when the peer server is down.

When the ICL between the peers fails, so that the DHCP servers become isolated from each other, new clients connecting are allocated addresses from local (or access-driven) address ranges only. Because the

ranges are declared as local on one peer, and remote on the other, without any overlap, there is no risk of double allocations. Because of the ICL failure, lease synchronization between the peers is not possible.

However, if the ICL failure lasts for an extended time, while both servers are up and running, then both peers consider their partner to be down. Both servers can start allocating addresses from their remote ranges, and there is a risk of assigning the same address to different clients. This situation **must be avoided** by ensuring that, when both peers are up, the ICL is also up.

Double allocations in a network is an indication of either two DHCP servers being isolated for too long or a misconfiguration in the network, and **must never happen**.

The [DHCP Server Failover States and State Transitions](#) section in this chapter provides a more extensive explanation.

Local-Remote Model – Active-Standby Configuration

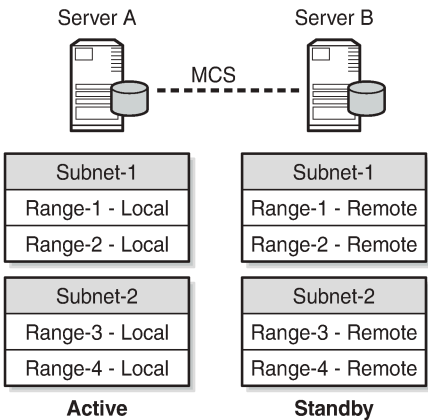
The first example of the local-remote model is an active-standby configuration; see [Figure 20: Local-Remote Model – Active-Standby](#). All the ranges in all pools are declared local on DHCP server A (the active server), and remote on DHCP server B (the standby server).

Usually, leases are allocated by the active DHCP server. The standby DHCP server synchronizes with the active server through MCS, so it can take over all ranges of all subnets in case of a failure.

Even though server B also receives DHCP requests from clients, it will not allocate addresses from its subnets (because they are all declared remote) unless server A is down. Only when server A goes down does server B become active, and connects, rebinds, and renews clients.

Caution must be taken when deploying this model, because fast switchover requires ignoring the maximum client lead time (MCLT) on takeover and a low value for the partner down delay, as described later in this chapter.

Figure 20: Local-Remote Model – Active-Standby



26355


Local-Remote Model – Load Sharing Configuration

The second example of the local-remote model is a load sharing configuration; see [Figure 21: Local-Remote Model – Subnet-Based Load Sharing](#). All subnet-1 ranges are declared local on server A, and remote on server B. For subnet-2, this is the opposite.

Usually, leases are allocated by both server A and server B. Server A is responsible for subnet-1, and thus manages addresses from subnet-1 (allocation, rebind, renew). At the same time, server B is responsible for subnet-2.

The standby DHCP server synchronizes with the active server through MCS, so it can take over all ranges of all subnets in case of a failure. Addresses from the remote ranges are managed only in case the server is in the partner down state.


Figure 21: Local-Remote Model – Subnet-Based Load Sharing



26356

A variant of this model is shown in [Figure 22: Local-Remote Model – Range-Based Load Sharing](#), where a subnet is split into two ranges with the first range declared as local and the second range declared as a remote. The local/remote ratio can be chosen arbitrarily. For example, the ratio for subnet-1 can be defined as 80/20, meaning that for a range of 100 addresses, 80 addresses are available in the local range and 20 in the remote range. At the same time, a 50/50 ratio can be defined for subnet-2.

Figure 22: Local-Remote Model – Range-Based Load Sharing



26357



Note:

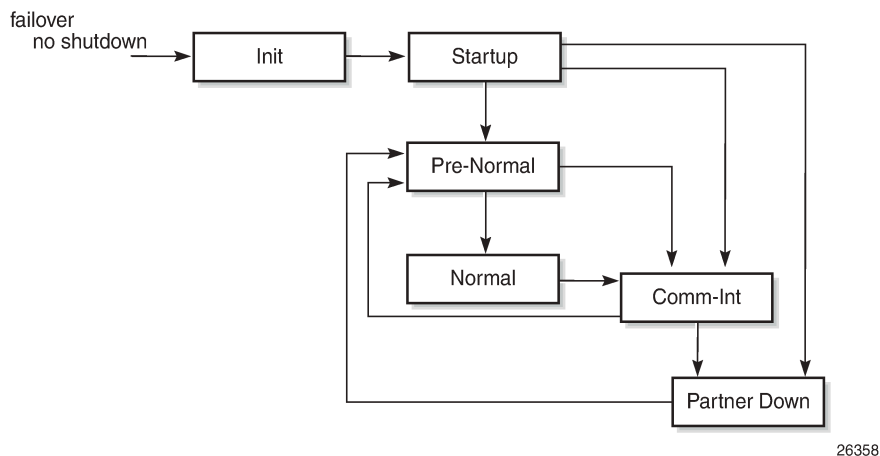
All these deployment models are supported for both IPv4 and IPv6, except for the model shown in [Figure 22: Local-Remote Model – Range-Based Load Sharing](#), because IPv6 prefixes cannot be organized in ranges.

DHCP Server Failover States and State Transitions

When a DHCP server is configured and enabled for failover, SR OS maintains a failover state, see [Figure 23: Failover State Transition Diagram](#). The failover state can have one of the following values:

- **INIT** – the DHCP server is initializing and possibly recovering its leases from the persistency database (if persistency is enabled). In this state, the DHCP server does not respond to any unicast or broadcast messages.
- **STARTUP** – the DHCP server recovers leases through MCS, and does not respond to any unicast or broadcast messages.
- **PRE-NORMAL** – the DHCP server responds to unicast and broadcast messages for addresses in the local range, and to unicast messages for addresses in the remote range.
- **NORMAL** – the DHCP server responds to local addresses only.
- **COMMUNICATIONS-INTERRUPTED** – the DHCP server responds to local (and access-driven) addresses only, and operator intervention is required. In the remainder of this chapter, this state is abbreviated as **COMM-INT**.
- **PARTNER-DOWN** – the DHCP server responds to local and remote ranges.

Figure 23: Failover State Transition Diagram



Enabling failover (no shutdown) triggers the change to the INIT failover state.

The INIT state is used while the DHCP server is recovering its lease database from persistency files when persistency applies. When recovery is completed, the failover state transitions to STARTUP. When persistency does not apply, the failover state transitions immediately to STARTUP.

When entering the STARTUP state, the startup-wait-timer is started to supervise the TCP connection setup to the MCS peer. If this timer expires and the connection is still not established, MCS communication has failed, and the failover state changes to COMM-INT. MCS recovery starts automatically if the TCP connection is established. When MCS recovery finishes, the failover state is changed to PRE-NORMAL.

However, if an MCS state-record indicates that the failover state was PARTNER-DOWN before a reboot, the failover state is set to PARTNER-DOWN immediately.

The DHCP server will not respond to any DHCP messages while in the INIT or STARTUP state.

When the PRE-NORMAL state is reached at power on or reboot (so the previous failover state was STARTUP), the system immediately changes the failover state to NORMAL.

In the NORMAL state, the DHCP server manages the addresses from the local (and access-driven) address ranges. In parallel, MCS keeps the DHCP lease states between the peers synchronized. If an MCS no-sync event is received, the failover state changes to COMM-INT.

Because MCS cannot determine whether a server is down or a server is not reachable because of an ICL failure, an operator must intervene when the COMM-INT state is reached. If no operator intervenes for an extended period of time (defined by the partner-down-delay, default 23h 59m 59s), the failover state changes to PARTNER-DOWN.

This is not a problem if the DHCP peer is down; for example, because of a power failure. The active DHCP server starts managing the addresses from the local and the remote address ranges. When the failing DHCP peer is up and running again, it gets to the NORMAL state through the process previously described.

The situation where two DHCP peers are isolated and running independent of each other, so that both are in the PARTNER-DOWN state, **must be avoided**. It would lead to double allocations, where both servers assign the same addresses to different clients, which is service disrupting to the users involved. Potential duplicates are resolved when the MCLT timer expires, and both peers are synchronized again.

Getting into the COMM-INT status is not service affecting, but should be avoided because DHCP lease synchronization fails. Operators must prevent both DHCP servers getting into the PARTNER-DOWN state, and the time to take corrective actions is bound by the partner-down-delay. If the partner-down-delay is not sufficiently large, ensure that one of the peer DHCP servers is not reachable by any of the clients anymore; for example, by shutting down or removing power from that server.

When in either the COMM-INT or PARTNER-DOWN state, and an MCS sync-event is received because the ICL becomes active again, the DHCP server moves to the PRE-NORMAL state, and starts the pre-normal-timer, which is initialized to the MCLT value, described in the next section.

In the PRE-NORMAL state, the DHCP server recovers the remote leases through MCS. While in this state, the DHCP server will respond to unicast and broadcast DHCP messages from the local ranges, and to unicast DHCP messages from the remote range. Recovery will be finished before the pre-normal timer expires, after which the failover state returns to NORMAL.

Maximum Client Lead Time

When failover does not apply, the DHCP server provides lease durations as defined in the pool or subnet definitions. When a client explicitly requests a lease duration, the server checks and validates the requested lease duration, potentially changing the requested lease duration to match the boundaries.

Regardless of the failover deployment model, it is important that DHCP servers can only allocate or extend a lease for a limited amount of time beyond the lease time known by its peer. The maximum client lead time (MCLT) defines the maximum time that one server can extend the lease for a client's binding, beyond the time known by the partner server, and is a safeguard against potential double allocations.

Nokia recommends using the same value for the MCLT on both partners of a failover pair. If they are different, the larger value is used. The default MCLT value is 10 min.

In the NORMAL state, clients initially get a lease time equal to the MCLT time. Over time, when renewing and rebinding, the allocated lease times are gradually increased to:

- the requested lease times if they are within the configured bounds
- the configured lease timer value

See the *Lease Time Synchronization* chapter in the *7450 ESS, 7750 SR, and VSR Triple Play Service Delivery Architecture Guide* for examples.

In the COMM-INT state, lease durations for existing leases are gradually decreased on renewal and rebinding, down to a minimum value defined by the MCLT. New clients are provided lease times equal to the MCLT.

When the DHCP server returns to the NORMAL state, lease durations start increasing again if clients renew and rebind.

Commands Controlling Failover and Failover State

Commands controlling failover are available at the DHCP server level, or at the pool level, in the base router and in a VPRN context.

These commands are grouped in the failover context:

- peer <ip-address> tag <sync-tag>
- ignore-mclt-on-takeover
- maximum-client-lead-time
- partner-down-delay
- startup-wait-time

The peer address is the IPv4 or IPv6 address of the DHCP failover peer, and is accompanied by a string of up to 32 characters, which serves as the sync-tag. This sync-tag must be the same on both peers.

The use of the other parameters is explained in the [DHCP Server Failover States and State Transitions](#) and [Maximum Client Lead Time](#) sections of this chapter.

A tools command is available, forcing the failover state to PARTNER-DOWN, and should be used with caution:

```
*A:P-2# tree flat | match tools | match force-partner
tools perform router dhcp local-dhcp-server failover force-partner-down
tools perform router dhcp local-dhcp-server pool failover force-partner-down
tools perform router dhcp6 local-dhcp-server failover force-partner-down
tools perform router dhcp6 local-dhcp-server pool failover force-partner-down
*A:P-2#
```

Configuration

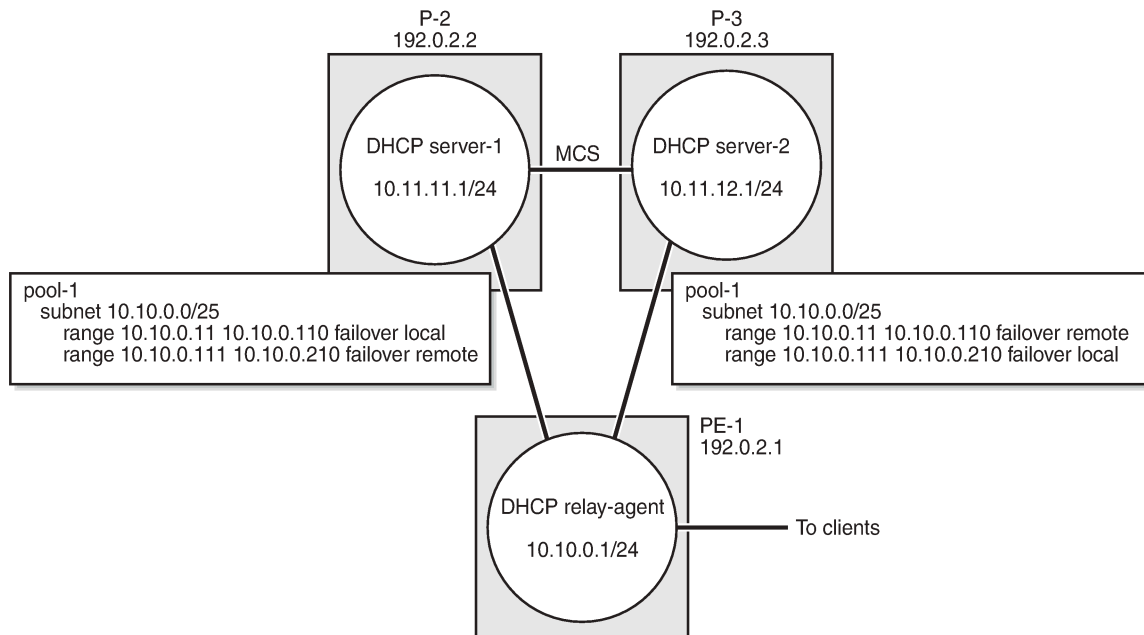
Starting a DHCP server in an SR OS environment requires following steps:

1. Configure the DHCP server.
2. Configure the interfaces for the DHCP server to listen on.
3. Configure one or more relay agents.

The baseline configuration used in this chapter is shown in [Figure 24: VPRN-1 Service Configuration](#), and relies on the relay agent to relay DHCP messages to and from both DHCP servers.

The example scenario uses DHCP clients only.

Figure 24: VPRN-1 Service Configuration



26359

Configure Multi-Chassis Synchronization

MCS must be configured before configuring failover, because the DHCP lease state database is to be synchronized between the failover peers. Therefore, P-2 points to P-3, and vice-versa, as follows:

```

# P-2
configure
  redundancy
    multi-chassis
      peer 192.0.2.3 create
      sync
      local-dhcp-server
      no shutdown
    exit
  no shutdown
exit
exit
exit
exit
exit

```

```

# P-3
configure
  redundancy
    multi-chassis

```

```
        peer 192.0.2.2 create
        sync
        local-dhcp-server
        no shutdown
    exit
    no shutdown
exit
exit
exit
exit
```

For MCS and failover to work, the clocks of the servers must be aligned, which is achieved through NTP. Configuration of NTP is beyond the scope of this chapter.

Configure common DHCP subnets

The 10.10.0.0/24 subnet is shared by P-2 and P-3. The [11-110] range is declared local to P-2 and remote to P-3. The [111-210] range is declared remote to P-2 and local to P-3, as follows:

```
# P-2
configure
service
    vprn 1 customer 1 create
    dhcp
        local-dhcp-server "dhcp-1" create
        use-gi-address scope pool
        pool "pool-1" create
        options
            dns-server 1.1.1.1 1.1.2.2
            lease-time min 20
        exit
        subnet 10.10.0.0/24 create
        options
            subnet-mask 255.255.255.0
            default-router 10.10.0.1
        exit
        address-range 10.10.0.11 10.10.0.110 failover local
        address-range 10.10.0.111 10.10.0.210 failover remote
    exit
exit
```

The subnet and pool definitions on P-3 are as follows:

```
# P-3
configure
service
    vprn 1 customer 1 create
    dhcp
        local-dhcp-server "dhcp-2" create
        use-gi-address scope pool
        pool "pool-1" create
        options
            dns-server 1.1.1.1 1.1.2.2
            lease-time min 20
        exit
        subnet 10.10.0.0/24 create
        options
            subnet-mask 255.255.255.0
            default-router 10.10.0.1
        exit
```

```
        address-range 10.10.0.11 10.10.0.110 failover remote
        address-range 10.10.0.111 10.10.0.210 failover local
    exit
exit
```

Configure Failover

P-2's failover configuration for this example is as follows:

```
# P-2
configure
service
    vprn 1 customer 1 create
    dhcp
        local-dhcp-server "dhcp-1" create
        failover
            peer 192.0.2.3 tag "mytag"
            maximum-client-lead-time min 12
            no shutdown
        exit
    no shutdown
exit
exit
```

P-3's failover configuration for this example is as follows:

```
# P-3
configure
service
    vprn 1 customer 1 create
    dhcp
        local-dhcp-server "dhcp-2" create
        failover
            peer 192.0.2.2 tag "mytag"
            maximum-client-lead-time min 12
            no shutdown
        exit
    no shutdown
exit
exit
```

Configure the Relay Agent

The DHCP relay agent for service VPRN 1 on PE-1 relays the DHCP messages to and from servers 10.11.11.1 and 10.11.12.1, as follows:

```
# PE-1
configure
service
    vprn 1 customer 1 create
    route-distinguisher 64496:1
    auto-bind-tunnel
    resolution-filter
        ldp
    exit
    resolution filter
exit
```

```
vrf-target target:64496:1
subscriber-interface "int-SUB1" create
address 10.10.0.1/24
group-interface "int-GRP1" create
arp-populate
dhcp
    option
        action replace
        circuit-id
        no remote-id
    exit
    server 10.11.11.1 10.11.12.1
    lease-populate 100
    client-applications dhcp ppp
    gi-address 10.10.0.1
    no shutdown
exit
sap 1/1/1:1 create
sub-sla-mgmt
    ---snip---
exit
exit
exit
exit
```

Debug and Troubleshooting

The following configuration enables debugging for DHCP server *dhcp-1* on VPRN 1 on both P-2 and P-3:

```
debug
router "1"
    local-dhcp-server "dhcp-1"
    detail-level medium
    mode egr-ingr-and-dropped
    exit
exit
exit
```

To ensure that the debug output is sent to a session, the following additional configuration is needed:

```
configure
log
    log-id 1
    from debug-trace
    to session
    no shutdown
    exit
exit
exit
```

Operation and Verification

The following command shows all DHCP servers for VPRN 1. The DHCP server names are listed together with their administrative state.

```
*A:P-2# show router 1 dhcp servers all
=====
```

```
Overview of DHCP Servers
=====
Active Leases:      2
Maximum Leases:    159744
Router              Server                               Admin State
-----
Service: 1          dhcp-1                               inService
=====
*A:P-2#
```

The following command shows the DHCP server summary for server *dhcp-1*. The parameters related to failover are shown in bold. The first block applies to the entire DHCP server of *dhcp-1*, the second block is specific to *pool-1*.

```
*A:P-2# show router 1 dhcp local-dhcp-server "dhcp-1" summary
=====
DHCP server dhcp-1  router 1
=====
Admin State           : inService
Operational State     : inService
Persistency State     : shutdown
User Data Base        : N/A
Use gateway IP address : enabled (scope pool)
Use pool from client   : disabled
Send force-renewals    : disabled
Creation Origin        : manual
Lease Hold Time        : 0h0m0s
Lease Hold Time For    : N/A
User-ident             : mac-circuit-id

Failover Admin State : inService
Failover Oper State : normal
Failover Persist Key : N/A
Administrative MCLT : 0h12m0s
Operational MCLT   : 0h12m0s
Startup wait time   : 0h2m0s
Partner down delay : 23h59m59s
Ignore MCLT         : disabled

-----
Pool name : pool-1
-----
Failover Admin State : outOfService
Failover Oper State : shutdown
Failover Persist Key : N/A
Administrative MCLT : 0h10m0s
Operational MCLT   : 0h10m0s
Startup wait time   : 0h2m0s
Partner down delay : 23h59m59s
Ignore MCLT         : disabled

-----
Subnet                Free    %    Stable  Declined  Offered  Rem-pend  Drain
-----
10.10.0.0/24          (L) 100    100%  0        0        0        0        N
                      (R) N/A      0      N/A     N/A     N/A     N/A     N
Totals for pool       100    100%  0        0        0        0
-----

Totals for server     100    100%  0        0        0        0
-----

Interface associations
Interface              Admin
```

```
-----
-----
Local Address Assignment associations
Group interface          Admin
-----
-----
No associations found
=====
*A:P-2#
```

With the DHCP server *dhcp-1* on P-2 in the NORMAL failover state, a user connecting gets an address allocated from a local pool, with the initial lease time set to the MCLT, as follows:

```
1 2017/02/02 15:18:23.20 CET MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp-1
Rx DHCP Discover

ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
siaddr: 0.0.0.0          giaddr: 10.10.0.1
chaddr: 00:00:00:01:01:01  xid: 0x21

DHCP options:
[82] Relay agent information: len = 25
[1] Circuit-id: PE-1|1|int-GRP1|1/1/1:1
[53] Message type: Discover
[255] End
"

2 2017/02/02 15:18:23.20 CET MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp-1
lease added for 10.10.0.12 state=offer
"

3 2017/02/02 15:18:23.20 CET MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp-1
Tx DHCP Offer to relay agent at 10.10.0.1 vrId=2

ciaddr: 0.0.0.0          yiaddr: 10.10.0.12
siaddr: 10.11.11.1        giaddr: 10.10.0.1
chaddr: 00:00:00:01:01:01  xid: 0x21

DHCP options:
[82] Relay agent information: len = 25
[1] Circuit-id: PE-1|1|int-GRP1|1/1/1:1
[53] Message type: Offer
[54] DHCP server addr: 10.11.11.1
[51] Lease time: 720
[1] Subnet mask: 255.255.255.0
[3] Router: 10.10.0.1
[6] Domain name server: length = 8
    1.1.1.1
    1.1.2.2
[255] End
"

4 2017/02/02 15:18:23.22 CET MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp-1
Rx DHCP Request
ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
siaddr: 0.0.0.0          giaddr: 10.10.0.1
chaddr: 00:00:00:01:01:01  xid: 0x21
```

```
DHCP options:
[82] Relay agent information: len = 25
    [1] Circuit-id: PE-1|1|int-GRP1|1/1/1:1
[53] Message type: Request
[50] Requested IP addr: 10.10.0.12
[54] DHCP server addr: 10.11.11.1
[255] End
"

5 2017/02/02 15:18:23.22 CET MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp-1
lease update for 10.10.0.12 state=stable
"

6 2017/02/02 15:18:23.22 CET MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp-1
Tx DHCP Ack to relay agent at 10.10.0.1 vrId=2

    ciaddr: 0.0.0.0          yiaddr: 10.10.0.12
    siaddr: 10.11.11.1       giaddr: 10.10.0.1
    chaddr: 00:00:00:01:01:01  xid: 0x21

DHCP options:
[82] Relay agent information: len = 25
    [1] Circuit-id: PE-1|1|int-GRP1|1/1/1:1
[53] Message type: Ack
[54] DHCP server addr: 10.11.11.1
[51] Lease time: 720
[1] Subnet mask: 255.255.255.0
[3] Router: 10.10.0.1
[6] Domain name server: length = 8
    1.1.1.1
    1.1.2.2
[255] End
"
```

Server *dhcp-1* on P-3 also offers a lease, but the client does not accept that offer so that lease is deleted. Because the client acknowledges the lease allocated by P-2, that lease is synchronized through MCS, as follows:

```
2 2017/02/02 15:18:23.68 CET MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp-1
Rx DHCP Discover

    ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
    siaddr: 0.0.0.0          giaddr: 10.10.0.1
    chaddr: 00:00:00:01:01:01  xid: 0x21

DHCP options:
[82] Relay agent information: len = 25
    [1] Circuit-id: PE-1|1|int-GRP1|1/1/1:1
[53] Message type: Discover
[255] End
"

3 2017/02/02 15:18:23.68 CET MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp-1
lease added for 10.10.0.112 state=offer
"

4 2017/02/02 15:18:23.68 CET MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp-1
Tx DHCP Offer to relay agent at 10.10.0.1 vrId=2
```



```

ciaddr: 0.0.0.0          yiaddr: 10.10.0.112
siaddr: 10.11.12.1       giaddr: 10.10.0.1
chaddr: 00:00:00:01:01:01  xid: 0x21

DHCP options:
[82] Relay agent information: len = 25
      [1] Circuit-id: PE-1|1|int-GRP1|1/1/1:1
[53] Message type: Offer
[54] DHCP server addr: 10.11.12.1
[51] Lease time: 720
[1] Subnet mask: 255.255.255.0
[3] Router: 10.10.0.1
[6] Domain name server: length = 8
      1.1.1.1
      1.1.2.2
[255] End
"

5 2017/02/02 15:18:23.70 CET MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp-1
Rx DHCP Request

ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
siaddr: 0.0.0.0          giaddr: 10.10.0.1
chaddr: 00:00:00:01:01:01  xid: 0x21

DHCP options:
[82] Relay agent information: len = 25
      [1] Circuit-id: PE-1|1|int-GRP1|1/1/1:1
[53] Message type: Request
[50] Requested IP addr: 10.10.0.12
[54] DHCP server addr: 10.11.11.1
[255] End
"

6 2017/02/02 15:18:23.70 CET MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp-1
dropped: Client didn't accept our offer, deleting lease 10.10.0.112
"

7 2017/02/02 15:18:23.70 CET MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp-1
lease added for 10.10.0.12 state=stable
"

```

With one user connected, check the leases on P-2 and P-3, as follows:

```

*A:P-2# show router 1 dhcp local-dhcp-server "dhcp-1" leases

=====
Leases for DHCP server dhcp-1 router 1
=====
IP Address      Lease State      Mac Address      Remaining      Clnt Fail
PPP user name/Opt82 Circuit Id      LifeTime      Type  Ctrl
User-db/Sticky-lease Hostname
-----
10.10.0.12      stable          00:00:00:01:01:01 0h33m31s      dhcp  local
PE-1|1|int-GRP1|1/1/1:1
-----
1 leases found
=====

```

```
*A:P-2#
```

```
*A:P-3# show router 1 dhcp local-dhcp-server "dhcp-1" leases
```

```
=====
Leases for DHCP server dhcp-1 router 1
=====
IP Address      Lease State      Mac Address      Remaining      Clnt  Fail
  PPP user name/Opt82 Circuit Id      LifeTime      Type  Ctrl
  User-db/Sticky-lease Hostname
-----
10.10.0.12      stable          00:00:00:01:01:01 0h33m5s      dhcp  remote
  PE-1|1|int-GRP1|1/1/1:1
-----
1 leases found
=====
*A:P-3#
```

For lease 10.10.0.12, failover control is local on P-2 and remote on P-3, and this matches the pool definitions from the beginning of the configuration section.

The details for the 10.10.0.12 lease can be shown with the following command. The remaining potential expiration time is ahead of the remaining lifetime, as follows:

```
*A:P-2# show router 1 dhcp local-dhcp-server "dhcp-1" leases 10.10.0.12 detail
```

```
=====
Lease for DHCP server dhcp-1 router 1
=====
IP-address      : 10.10.0.12
Lease-state     : stable
Lease started   : 2017/02/02 15:18:23
Last renew     : N/A
Remaining LifeTime : 0h8m24s
Remaining Potential Exp. Time: 0h32m24s
Sticky-lease Host Name : N/A
MAC address     : 00:00:00:01:01:01
Xid            : 0x21
Failover Control : local
Client Type     : dhcp
User-db Host Name : N/A
User-db Address Type : N/A
Persistence Key : N/A
Opt82 Hex Dump  : (length=27)
                  : 52 19 01 17 50 45 2d 31 7c 31 7c 69 6e 74 2d 47
                  : 52 50 31 7c 31 2f 31 2f 31 3a 31
Opt82 Circuit Id : PE-1|1|int-GRP1|1/1/1:1
Opt82 Remote Id  :
Opt82 Subscr Id  :
Opt82 VS System  :
Opt82 VS Clnt MAC :
Opt82 VS Service :
Opt82 VS SAP     :
Opt82 VS String  :
Lease Remaining Hold Time : 0h0m0s
=====
*A:P-2#
```

On renewal of this lease, the offered lease time is increased to the configured lease time, as follows:

```
7 2017/02/02 15:24:24.10 CET MINOR: DEBUG #2001 vprn1 DHCP server
```

```
"DHCP server: dhcp-1
Rx DHCP Request

    ciaddr: 10.10.0.12      yiaddr: 0.0.0.0
    siaddr: 0.0.0.0        giaddr: 0.0.0.0
    chaddr: 00:00:00:01:01:01  xid: 0x21

    DHCP options:
    [82] Relay agent information: len = 25
        [1] Circuit-id: PE-1|1|int-GRP1|1/1/1:1
    [53] Message type: Request
    [255] End
"

8 2017/02/02 15:24:24.10 CET MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp-1
lease update for 10.10.0.12 state=stable
"

9 2017/02/02 15:24:24.10 CET MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp-1
Tx DHCP Ack to client at 10.10.0.12 vrId=2

    ciaddr: 10.10.0.12      yiaddr: 10.10.0.12
    siaddr: 10.11.11.1      giaddr: 0.0.0.0
    chaddr: 00:00:00:01:01:01  xid: 0x21

    DHCP options:
    [82] Relay agent information: len = 25
        [1] Circuit-id: PE-1|1|int-GRP1|1/1/1:1
    [53] Message type: Ack
    [54] DHCP server addr: 10.11.11.1
    [51] Lease time: 1800
    [1] Subnet mask: 255.255.255.0
    [3] Router: 10.10.0.1
    [6] Domain name server: length = 8
        1.1.1.1
        1.1.2.2
    [255] End
"
```

The remaining lifetime and the potential remaining expiration time are adjusted, as follows:

```
*A:P-2# show router 1 dhcp local-dhcp-server "dhcp-1" leases 10.10.0.12 detail

=====
Lease for DHCP server dhcp-1 router 1
=====
IP-address           : 10.10.0.12
Lease-state          : stable
Lease started        : 2017/02/02 15:18:23
Last renew           : 2017/02/02 15:24:24
Remaining LifeTime   : 0h25m1s
Remaining Potential Exp. Time: 0h40m1s
Sticky-lease Host Name : N/A
MAC address          : 00:00:00:01:01:01
Xid                   : 0x21
Failover Control      : local
Client Type           : dhcp
User-db Host Name     : N/A
User-db Address Type  : N/A
Persistence Key       : N/A
Opt82 Hex Dump       : (length=27)
```

```

: 52 19 01 17 50 45 2d 31 7c 31 7c 69 6e 74 2d 47
: 52 50 31 7c 31 2f 31 2f 31 3a 31
Opt82 Circuit Id      : PE-1|1|int-GRP1|1/1/1:1
Opt82 Remote Id      :
Opt82 Subscr Id       :
Opt82 VS System       :
Opt82 VS Clnt MAC     :
Opt82 VS Service      :
Opt82 VS SAP          :
Opt82 VS String       :
Lease Remaining Hold Time : 0h0m0s

=====
*A:P-2#

```

P-3 updates its lease database through MCS as users connect, disconnect, renew, or rebind their leases, as long as the ICL is uninterrupted and the failover state remains NORMAL, as follows:

```

9 2017/02/02 15:39:23.89 CET MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp-1
lease update for 10.10.0.12 state=stable
"

15 2017/02/02 15:39:41.22 CET MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp-1
lease added for 10.10.0.13 state=stable
"

16 2017/02/02 15:40:33.04 CET MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp-1
lease deleted for 10.10.0.13 (delete from peer)

```

A failure of the ICL is emulated by disabling failover for VPRN-1 on P-3, as follows:

```

*A:P-3# configure service vprn 1 dhcp local-dhcp-server "dhcp-1" failover shutdown

```

On P-3, the following debug messages appear:

```

28 2017/02/02 16:23:00.46 CET MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp-1
Failover oper state change from NORMAL to SHUTTING-DOWN
"

29 2017/02/02 16:23:00.46 CET MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp-1
remote lease deleted for 10.10.0.12 (failover shutdown)
"

30 2017/02/02 16:23:00.46 CET MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp-1
Failover oper state change from SHUTTING-DOWN to SHUTDOWN
"

31 2017/02/02 16:23:00.46 CET MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp-1
RX failover peer state COMMUNICATION-INTERRUPTED
"

```

On P-2, the following debug messages appear:

```

32 2017/02/02 16:23:00.45 CET MINOR: DEBUG #2001 vprn1 DHCP server

```

```
"DHCP server: dhcp-1
RX failover peer state SHUTDOWN
"

33 2017/02/02 16:23:00.45 CET MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp-1
Failover oper state change from NORMAL to COMMUNICATION-INTERRUPTED
"
```

Displaying the DHCP server summary for VPRN-1 on P-2 again shows the failover operational state as noCommunication. "Time Left" indicates how much time is left before the failover state changes to the PARTNER-DOWN operational state if no action is taken for resolving the communication issue.

```
*A:P-2# show router 1 dhcp local-dhcp-server "dhcp-1" summary
=====
DHCP server dhcp-1 router 1
=====
Admin State           : inService
Operational State     : inService
Persistency State     : shutdown
User Data Base        : N/A
Use gateway IP address : enabled (scope pool)
Use pool from client   : disabled
Send force-renewals    : disabled
Creation Origin        : manual
Lease Hold Time        : 0h0m0s
Lease Hold Time For    : N/A
User-ident             : mac-circuit-id

Failover Admin State   : inService
Failover Oper State   : noCommunication
Failover Persist Key   : N/A
Time Left            : 23h56m15s before state transition
Administrative MCLT    : 0h12m0s
Operational MCLT       : 0h12m0s
Startup wait time      : 0h2m0s
Partner down delay     : 23h59m59s
Ignore MCLT            : disabled

---snip---

=====
*A:P-2##
```

Check the status for lease 10.10.0.12 again, as follows.

```
*A:P-2# show router 1 dhcp local-dhcp-server "dhcp-1" leases 10.10.0.12 detail
=====
Lease for DHCP server dhcp-1 router 1
=====
IP-address           : 10.10.0.12
Lease-state           : stable
Lease started        : 2017/02/02 15:18:23
Last renew           : 2017/02/02 16:24:24
Remaining LifeTime    : 0h23m35s
Remaining Potential Exp. Time : 0h23m35s
Sticky-lease Host Name : N/A
MAC address          : 00:00:00:01:01:01
Xid                   : 0x21
Failover Control      : local
Client Type           : dhcp
```

```
User-db Host Name      : N/A
User-db Address Type   : N/A
Persistence Key        : N/A
Opt82 Hex Dump         : (length=27)
                       : 52 19 01 17 50 45 2d 31 7c 31 7c 69 6e 74 2d 47
                       : 52 50 31 7c 31 2f 31 2f 31 3a 31
Opt82 Circuit Id       : PE-1|1|int-GRP1|1/1/1:1
Opt82 Remote Id        :
Opt82 Subscr Id        :
Opt82 VS System        :
Opt82 VS Clnt MAC      :
Opt82 VS Service       :
Opt82 VS SAP           :
Opt82 VS String        :
Lease Remaining Hold Time : 0h0m0s

=====
*A:P-2#
```

New clients connecting are allocated addresses from the local address ranges on either P-2 or P-3, even when the server is in the COMM-INT failover state. In this example, both clients are allocated and acknowledged addresses by P-2, as follows:

```
*A:P-2# show router 1 dhcp local-dhcp-server "dhcp-1" leases

=====
Leases for DHCP server dhcp-1 router 1
=====
IP Address      Lease State      Mac Address      Remaining      Clnt  Fail
PPP user name/Opt82 Circuit Id      LifeTime      Type  Ctrl
User-db/Sticky-lease Hostname
-----
10.10.0.12      stable          00:00:00:01:01:01 0h19m2s      dhcp  local
PE-1|1|int-GRP1|1/1/1:1
10.10.0.14      stable          00:00:00:01:01:02 0h10m51s     dhcp  local
PE-1|1|int-GRP1|1/1/1:1
-----
2 leases found
=====
*A:P-2#
```

Solving the communications issue is emulated by enabling failover for VPRN-1 on P-3 again:

```
*A:P-3# configure service vprn 1 dhcp local-dhcp-server "dhcp-1" failover no shutdown
```

The debug log on P-2 shows as follows:

```
49 2017/02/02 16:41:17.60 CET MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp-1
RX failover peer state NORMAL
"

50 2017/02/02 16:41:17.60 CET MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp-1
Failover oper state change from COMMUNICATION-INTERRUPTED to PRE-NORMAL
"
```

The debug log on P-3 shows as follows:

```
37 2017/02/02 16:41:17.60 CET MINOR: DEBUG #2001 vprn1 DHCP server
```

```
"DHCP server:  dhcp-1
Failover oper state change from SHUTDOWN to STARTUP
"

38 2017/02/02 16:41:17.60 CET MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp-1
Failover oper state change from STARTUP to PRE-NORMAL
"

39 2017/02/02 16:41:17.60 CET MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp-1
Failover oper state change from PRE-NORMAL to NORMAL
"

40 2017/02/02 16:41:17.60 CET MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp-1
lease added for 10.10.0.12 state=stable
"

41 2017/02/02 16:41:17.60 CET MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp-1
lease added for 10.10.0.14 state=stable
"

42 2017/02/02 16:41:17.60 CET MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp-1
RX failover peer state NORMAL
"
```

The leases allocated by P-2 are synchronized with peer P-3, so they are marked as remote leases, as follows:

```
*A:P-3# show router 1 dhcp local-dhcp-server "dhcp-1" leases

=====
Leases for DHCP server dhcp-1 router 1
=====
IP Address      Lease State      Mac Address      Remaining      Clnt  Fail
  PPP user name/Opt82 Circuit Id      LifeTime      Type  Ctrl
  User-db/Sticky-lease Hostname
-----
10.10.0.12      stable          00:00:00:01:01:01 0h38m37s      dhcp  remote
  PE-1|1|int-GRP1|1/1/1:1
10.10.0.14      stable          00:00:00:01:01:02 0h31m57s      dhcp  remote
  PE-1|1|int-GRP1|1/1/1:1
-----
2 leases found
=====
*A:P-3#
```

Eventually, P-2's failover state changes to NORMAL again, as follows:

```
57 2017/02/02 16:53:17.85 CET MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp-1
Failover oper state change from PRE-NORMAL to NORMAL
"
```

Conclusion

SR OS supports DHCP server redundancy with failover, providing ISPs the capabilities to offer DHCP service during a partial power loss or partial network outage.

DHCPv4 Server Basics

This chapter describes DHCPv4 server basics.

Topics in this chapter include:

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

Applicability

This chapter is applicable to SR OS routers and is based on SR OS Release 14.0.R4.

Overview

The Dynamic Host Configuration Protocol for IPv4 (DHCPv4) provides a method for assigning addresses to hosts, and conveys additional configuration data to these hosts.

DHCPv4 allows for a flexible mapping of addresses to devices; for example, identified through their MAC address. While the DHCPv4 server owns and manages addresses organized in one or more pools, a DHCPv4 client obtains an address from the DHCPv4 server, which creates a lease for that client. This provides the client the right to use the address, and the server ensures that the address will not be assigned to other clients.

The DHCPv4 server implemented in SR OS has the following features:

- Address management. The DHCPv4 server keeps track of the used and unused addresses. For the used addresses, lease durations are maintained.
- Configuration parameter management. The DHCPv4 server stores parameters that are to be used by clients when they connect.
- Persistency. When enabled, the DHCPv4 server stores the leases on non-volatile storage so that the leases remain across potential node reboots.
- Failover capability. In dual-homed DHCPv4 server scenarios, a primary DHCPv4 server can take over the responsibility of a failing peer.

The DHCPv4 server failover capability is beyond the scope of this chapter.

In this chapter, when DHCP is mentioned, it implies DHCPv4.

Characteristics

IPv4 addresses and parameters are provided by the DHCP server through the Discover – Offer – Request – Acknowledge (DORA) message sequence as explained in the [IPv4 DHCP Hosts](#) chapter.


IPoE clients use DORA messages to communicate with the DHCP server via a relay agent. PPP clients use Link Control Protocol (LCP) and Internet Protocol Control Protocol (IPCP) to communicate with the router, and a local (DHCP) client manages the communication with the DHCP server. IPoE clients and PPP clients are also referred to as native clients and simulated clients, respectively.

When the DHCP server and the relay agent are physically located in the same SR OS node, the server is referred to as a local DHCP server; when they are in different nodes, the server is considered remote. Clients can obtain an address from a local, remote, or (external) third-party DHCP server.

A DHCP server can be used for IPoE users and PPP users simultaneously. A DHCP server must be hosted by a VPRN service or the base router. It can be accessed in either of the following ways; see [Figure 25: Accessing a DHCP server](#):

- When a DHCP user connects, the DORA message sequence running between the DHCP client and the DHCP server also passes through a relay agent, adding and removing user-defined options along the way. The relay agent and the DHCP server can be located in the same or in a different (remote) node.
- When a PPP user connects through LCP and IPCP on a service with an internal DHCP client (local client) configured, the local client manages the DORA communication toward the DHCP server, if the relay agent also has relaying enabled for PPP applications. The local client and the DHCP server can be located in the same or in a different (remote) node.
- When a PPP user connects through LCP and IPCP on a service with local address assignment configured, the DHCP server is accessed through an API-call. See the [ESM SLAAC Prefix Assignment via Local Address Server](#) chapter for an explanation of the local address assignment concept, which also applies to PPP.

Figure 25: Accessing a DHCP server



26095

A DHCP server is supported for the routed CO model as well as for the bridged CO model.

A DHCP server can be hosted by the base router or a VPRN service, for public or private use. Because multiple VPRN services can coexist in a single node, each having its own DHCP server, overlapping address ranges are supported.

DHCP Lease

The DHCP server maintains the following data for every allocation request in a lease:

1. client-type (PPP or DHCP)
2. IP address
3. MAC address
4. lease state
5. option 82, if relevant
6. option 60 (vendor class identifier), if relevant
7. lease timer related data
8. persistence key, if applicable

A lease for a single client is in one of the following states:

1. offered: The IP address was offered to the client. The client still has to acknowledge the offer by sending a DHCP request.
2. stable: The IP address is now in use by the client.
3. force-renew-pending: The IP address is in use by the client, but the server sends a DHCP force-renew message to the client, because an option has changed at pool, subnet, or client (via LUDb) level.
4. remove-pending: The IP address is in use by the client, but the corresponding subnet range is deleted. The server sends a force-renew message to the client to force the client to reinitialize in order to get a new IP address.
5. held: The IP address has been used by the client but the lease has expired. The lease is now in the hold list so that the client can get the same IP address upon the next request for a lease.
6. internal: The IP address has been leased via local address assignment and is in use.
7. internal-orphan: The IP address has been leased via local address assignment and is in use. However, there is no configured subnet to which this lease belongs, because it has been removed or because this lease was installed through dual-homing.
8. internal-offered: The IP address has been offered via local address assignment, but the client has not acknowledged the offer yet.
9. internal-held: The IP address has been offered via local address assignment, but the lease is currently not active. The address is now in the hold list so that the same IP address can be offered to the same client upon request of a lease.
10. sticky: The IP address is reserved for the client and will remain reserved until the reservation for it is cleared. The client will get the same IP address upon the next request for a lease.

User Identification

The key to the leases managed by the DHCP server is configured at server level, and can be set to one of the following values (the default value being **mac-circuit-id**):

```
configure (router | service vprn <service-id>) dhcp local-dhcp-server <server-name>  
    user-ident {client-id|circuit-id|mac|mac-circuit-id|remote-id}
```

The client ID is DHCP option 61; the circuit ID and the remote ID are sub-options 1 and 2 of DHCP option 82, respectively.

Setting **user-ident** to, for example, **circuit-id** can provide a CPE the same IP address regardless of its MAC address; thereby facilitating CPE replacement scenarios.

Lease Hold Time

The usual way for a DHCP client to indicate to the DHCP server it does not need its lease anymore is by sending a release message to the server; this is referred to as a solicited release.

However, when a client gets disconnected, or loses power, no release message is received by the server and the lease ultimately expires; this is referred to as an unsolicited release.

Without a lease hold timer, a lease is immediately deleted when the client sends the release message, or when the lease expires. The corresponding address is returned back to the pool of free addresses, and can be assigned to different clients. There is no guarantee that a client gets the same address again.

With a lease hold timer defined, a lease (entry) is not immediately deleted when the lease timer expires. Instead, the lease is put in the *held* or *internal-held* state. The lease is deleted when the hold timer expires, and the address is returned back to the pool. When the client connects, renews, or rebinds its lease before the hold timer expires, the client gets its previous lease again. There is no guarantee that the client gets the same address.

A lease hold timer can optionally be defined at the DHCP server level using the following command:

```
configure (router | service vprn <service-id>) dhcp local-dhcp-server <server-name>  
    lease-hold-time [days <days>] [hrs <hours>] [min <minutes>] [sec <seconds>]  
  
    <days>           : [0..7305]  
    <hours>           : [0..23]  
    <minutes>         : [0..59]  
    <seconds>         : [0..59]
```

If delayed deletion is also required on reception of a release message (solicited release), use the following command:

```
configure (router | service vprn <service-id>) dhcp local-dhcp-server <server-name>  
    lease-hold-time-for solicited-release
```

The same behavior can be applied to IPSec, but that is beyond the scope of this chapter.

Fixed address allocation using DHCP server

Devices using DHCPv4 but still require the same IPv4 (fixed) can be populated in a LUDB associated with a DHCPv4 server. Sometimes these devices are called static devices because they get a fixed IP address assigned each time they request a DHCP lease. See the [Local User Database for DHCPv4 Server](#) chapter.

The DHCPv4 server also supports lease reservation based on client identifiers provisioned at the time of reservation. After a lease is reserved, it is referred to as a sticky lease in the server.

This sticky lease will subsequently only be assigned to a DHCP client using the same identifiers, though it is not required for such a client to ever exist. Sticky leases are not removed via timeout or DHCP releases but can only be removed via the management interface (SNMP/CLI). Sticky leases are typically not used for devices using DHCPv4 but for applications like vRGW through SNMP.

A sticky lease requires a host name. Identification of the host can be through a MAC address, a circuit ID, a remote ID, or a combination of these. The IP address must be in the range of the parenting pool.

```
tools perform router <router-id> dhcp local-dhcp-server <server-name> pool <pool-name> create-
sticky-lease <hostname>
    [mac <ieee-address>]
    [circuit-id <circuit-id>]
    [client-id <client-id>]
    [requested-ip-address <ip-address>]
    [circuit-id-hex <circuit-id-hex-string>]
    [client-id-hex <client-id-hex-string>]
<hostname> : [32 chars max]
<ieee-address> : xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx
<circuit-id> : [253 chars max]
<client-id> : [255 chars max]
<ip-address> : a.b.c.d
<circuit-id-hex-st*> : [0x0..0xFFFFFFFF...(max 506 hex nibbles)]
<client-id-hex-str*> : [0x0..0xFFFFFFFF...(max 510 hex nibbles)]
```

Address Allocation for Sticky Leases

Sticky leases provide a static mapping between a hardware address and an IP address. This means that a particular device always gets the same IP address.

A sticky lease requires a host name. Identification of the host can be through a MAC address, a circuit ID, a remote ID, or a combination of these. The IP address must be in the range of the parenting pool.

```
tools perform router <router-id> dhcp local-dhcp-server <server-name>
    pool <pool-name> create-sticky-lease <hostname>
        [mac <ieee-address>]
        [circuit-id <circuit-id>]
        [client-id <client-id>]
        [requested-ip-address <ip-address>]
        [circuit-id-hex <circuit-id-hex-string>]
        [client-id-hex <client-id-hex-string>]
<hostname> : [32 chars max]
<ieee-address> : xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx
<circuit-id> : [253 chars max]
<client-id> : [255 chars max]
<ip-address> : a.b.c.d
<circuit-id-hex-st*> : [0x0..0xFFFFFFFF...(max 506 hex nibbles)]
```


```
<client-id-hex-str> : [0x0..0xFFFFFFFF...(max 510 hex nibbles)]
```

As an alternative to sticky leases, an LUDB can also be used to provide a static mapping between a hardware address and an IP address. See the [Local User Database for DHCPv4 Server](#) chapter. However, hosts added to a local user database can only survive a reboot by saving the configuration through the **admin save** command.

Pool and Subnet Management

The DHCPv4 servers manage IPv4 addresses, subnets, and pools. These are hierarchically related to one another; see [Figure 26: Addresses, Subnets, and Pools in a DHCPv4 Server](#).

Figure 26: Addresses, Subnets, and Pools in a DHCPv4 Server



26096

A subnet is identified by an IP address and a netmask, and defines:

- One or more address ranges – The ranges in the subnet that the server can allocate addresses from. Multiple address ranges cannot overlap.
- One or more exclude address ranges (optional) – A sub-range of the preceding range that the server will not allocate addresses from.
- Minimum-free – A notification is generated when the amount of free leases reaches this value (trap and log 99).
- Maximum-declined – Security counter measure, to prevent rogue clients from depleting the subnet. When this maximum value is reached, the oldest declined address will be returned to the pool.
- DHCP options:
 - default-router – up to four addresses can be defined
 - subnet-mask – subnet mask to be used by the clients
 - custom-options – additional options, when required

A pool is identified by name (maximum 32 characters), and defines:

- One or more subnets
- Min-lease-time – requests for a shorter lease time are set to this value; default is 10 min
- Max-lease-time – requests for a longer lease time are set to this value; default is 10 days
- Offer-time – a timer indicating how long an offer remains valid before the address offered is returned to the pool when no Request message is received (default 1 min)

- Minimum-free – a notification is generated when the amount of free leases reaches this value (trap and log 99), with an optional trap when all leases are used
- DHCP options:
 - dns-server – up to four DNS servers can be specified
 - domain-name – the domain to use for DNS resolution when clients provide unqualified host names.
 - lease-renew-time – defines when the client transitions to the renew state (T1)
 - lease-rebind-time – defines when the client transitions to the rebinding state (T2)
 - lease-time – the duration of time that the DHCP server grants to the DHCP client
 - netbios-name-server – defines up to four NetBIOS name servers
 - netbios-node-type – defines the NetBIOS node type (B, P, M, or H)
 - custom-option – additional options, when required

The options added by the DHCP server in response to an allocation request is a combination of the options provided by an LUDB (if applicable), subnet options, and pool options, in this sequence.

Lease Time

A DHCP client can request a specific lease time. The DHCP server checks for this value to be within the bounds as defined at pool level. If the requested lease time is out of bounds, it is set to either the minimum or the maximum value.

If a DHCP client does not request a specific lease time, the DHCP server takes the value from a matching LUDB entry, if available, or from the lease-time parameter defined at pool level, in this sequence. If the pool level lease time is not defined, the maximum lease time is used.

The best practice is to apply the following rule when defining values for the various timers:

```
lease-time > lease-rebind-time > lease-renew-time
```

However, the server does not check consistency of these timers, because the final values offered to the DHCP clients can come from various sources, which are out of the control of the DHCP server.

The local DHCP client always requests a lease time of 1 h to the server for PPP users connecting via the local client.

Address Allocation

When a request arrives at the DHCPv4 server, the server accesses the lease state database using the user ID as a key, checking for an existing lease. If a lease is already available, that lease is used.

Assuming that no lease is present in the lease state database yet, and that the server has a local user database attached, a matching entry is searched for in that local user database; see the [Local User Database for DHCPv4 Server](#) chapter.

In terms of address assignment, an LUDB attached to a DHCP server can return:


- an IP address – This (fixed) address is offered to the requester, where this address must not overlap with the address ranges configured in the local DHCP server.
- a Gi address – This address overrules any Gi address received from the requester.

- a pool name – A free address in one of the subnets in that pool is offered. Optionally, a secondary pool can be defined, which is used in case the primary pool is exhausted.
- **use-gi-address [scope <subnet | pool>]** – When the scope is set to subnet, the server offers an address from the subnet that includes the Gi address. When the scope is set to pool, the server offers an address from the subnet that includes the Gi address, or from the other subnets belonging to the same pool.
- **use-pool-from-client [delimiter <delimiter>]** – The pool name specified in the DHCP client message options and added by the relay agent is used. A free address in one of the subnets in that pool is offered. If two pools are available, the configured delimiting character identifies the splitting-point to find the names of both pools.

If a unique address is found in the LUDB, that address is offered by the server to the requester.

For the general address allocation flow, see [Figure 27: General Address Allocation for DHCP](#). The [Local User Database for DHCPv4 Server](#) chapter applies when an LUDB is attached to the DHCP server.

Figure 27: General Address Allocation for DHCP



26097

Two additional parameters are available at the server level, controlling which pool and subnet an address is taken from, as follows:

```

[no] use-pool-from-client [delimiter <delimiter>]
      <delimiter>          : [1 chars max]

[no] use-gi-address [scope <scope>]
      <scope>             : subnet|pool
  
```


With a requester-provided pool name and **use-pool-from-client** active, the server checks for that pool to exist before selecting a free address from one of the subnets in that pool.

With a requester-provided Gi address and **use-gi-address scope subnet** active, a free address is taken from the subnet that includes the Gi address. With **use-gi-address scope pool**, another subnet in the pool is used if the original subnet is exhausted.

The following rules apply to the DHCP server address allocation flow:

- Assume a DHCP server with an LUDB applied, and **use-gi-address** active:
 - A host lookup failure will not result in the request being dropped. The server sends an offer using an address selected based on the Gi address.
 - A successful host lookup, but returning a non-existent pool name, results in the server dropping the request, so no offer is sent.
- Assume a DHCP server without an LUDB applied, but with **use-pool-from-client** and **use-gi-address** active:
 - A requester not providing a pool name results in the server sending an offer using an address selected based on the Gi address.
 - A requester providing a non-existent pool name results in the server dropping the request, so no offer is sent.

Therefore, **use-pool-from-client** takes precedence over **use-gi-address**. The DHCP server selects an address from a pool if that pool exists. If no pool name is provided to the DHCP server, address selection is based on the Gi address, when allowed through the **use-gi-address** directive.

The pool name provided by a relay agent can be a concatenation of two pool names, where the delimiter character is used to split the string apart in the original pool names.

Subnet Draining

When a subnet is put in the drained state through the drain command, no new leases can be assigned from that subnet. Existing leases are cleaned up upon renewal or rebinding of the client. This is useful in renumbering scenarios; see the [Configuration](#) section for an example.

Force Renew

Parameter force-renews enables DHCP servers to issue DHCP force-renew messages to stable clients, informing them about a configuration change.

With **force-renews** enabled, the server does not need to wait for a client to pass through its renew or rebind sequence to provide the reconfigured options, speeding up the configuration change.

Changes can be applied at the LUDB-level, subnet level, or pool-level.

DHCP Server Persistency

The DHCP protocol does not have a keep-alive mechanism to detect unavailability. Without precaution, a node reboot causes the loss of the DHCP lease state. Because DHCP clients only attempt a reinitialization sequence after expiration of the lease timer, service outages could become unacceptably long.

The DHCP server lease state can be made persistent across reboots. The lease state is then restored from a persistency file stored on the compact flash. Therefore, DHCP clients will only lose connectivity for the duration of the reboot, and no renew or rebind is needed.

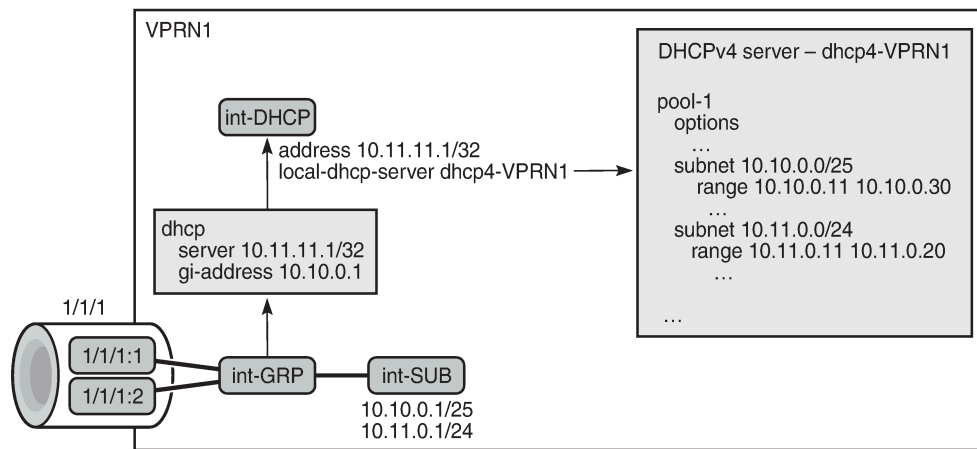
Configuration

Starting a DHCP server up in an SR OS environment requires following steps:

- Configure the DHCP server.
- Configure the interfaces for the DHCP server to listen on.
- Configure one or more relay agents.

The baseline configuration used in this chapter is shown in [Figure 28: Baseline Service Configuration](#).

Figure 28: Baseline Service Configuration



26098

Configure the DHCP Server

One or more DHCP servers can be configured in the base router or in any routed service. VPRN 1 from [Figure 28: Baseline Service Configuration](#) has a single DHCP server named *dhcp4-VPRN1*, with two pools: *pool-1* and *pool-2*. The first pool contains two subnets; the second pool contains a single subnet.

Address selection is primarily based on a pool name (**use-pool-from-client**), and secondarily on the Gi address with the scope set to pool (**use-gi-address scope pool**). This means that address selection will be Gi address-based, if no pool name is provided to the server. Having the scope set to **pool** enables the server to allocate addresses from other subnets within the same pool.

Different options and custom options are defined at different levels. All subnets include an address range. Subnet 10.10.0.0/25 also has an exclusion range, as follows:

```
configure
  service
    vprn 1 customer 1 create
      dhcp
```

```
local-dhcp-server dhcp4-VRPN1 create
  use-gi-address scope pool
  use-pool-from-client
  pool "pool-1" create
    options
      dns-server 1.1.1.1 1.1.2.2
      lease-time hrs 2
      custom-option 150 address 1.1.1.1
    exit
  subnet 10.10.0.0/25 create
    options
      subnet-mask 255.255.255.128
      default-router 10.10.0.1 10.10.0.2
      custom-option 130 string "MyOption1"
    exit
    exclude-addresses 10.10.0.61 10.10.0.70
    address-range 10.10.0.11 10.10.0.126
  exit
  subnet 10.11.0.0/24 create
    options
      subnet-mask 255.255.255.0
      default-router 10.11.0.1
      custom-option 130 string "MyOption2"
    exit
    address-range 10.11.0.11 10.11.0.20
  exit
exit
pool "pool-2" create
  subnet 10.20.0.0/16 create
    options
      subnet-mask 255.255.0.0
      default-router 10.20.0.1
    exit
    address-range 10.20.0.21 10.20.0.120
  exit
exit
no shutdown
exit
exit
exit
exit
exit
```

Configure the DHCP Interface

The DHCP server needs to be listening on one or more interfaces. In the example from [Figure 27: General Address Allocation for DHCP](#), the DHCP server is associated with interface *int-DHCP*, with loopback address 10.11.11.1, as follows. The DHCP server cannot be applied to a group interface.

```
configure
  service
    vprn 1 create
      interface "int-DHCP" create
        address 10.11.11.1/24
        local-dhcp-server "dhcp4-VRPN1"
        loopback
      exit
    exit
  exit
exit
```

Configure Relay Agents

The configuration of the DHCP server must align with the configuration of the relay agents for the server to assign addresses correctly. For example, defining the server to allocate addresses based on a pool name, but not providing a pool name toward the server, might not provide the expected result, because this will not necessarily lead to addresses being assigned and offered to clients.

The DHCP relay agent is configured in the `dhcp` context, as follows:

- **gi-address** – the gateway IPv4 address used by the relay agent
- **server** – up to 8 DHCP servers can be defined by their IPv4 address; only 10.11.11.11 is used in this example
- **client-applications dhcp ppp** – the DHCP server will allocate addresses for DHCP and PPP clients
- **option** – the options added/removed to/from messages toward the server. In the example, the circuit-id, the remote-id, and the pool-name are added.
- **trusted** – this parameter ensures that DHCP messages with option 82 included and the gi-address set to zero are being processed instead of being dropped

```
configure
  service
    vprn 1 customer 1 create
      subscriber-interface "int-SUB" create
        group-interface "int-GRP" create
          dhcp
            option
              action replace
              circuit-id
              remote-id
              vendor-specific-option
              pool-name
            exit
          exit
        server 10.11.11.1
        lease-populate 100
        client-applications dhcp ppp
        gi-address 10.10.0.1
        no shutdown
      exit
    exit
  exit
exit
exit
exit
exit
```

Configure DHCP Server Persistency

The following configuration stores the DHCP server lease-state persistency file on cf1:

```
configure
  system
    persistence
      dhcp-server
        location cf1:
      exit
    exit
  exit
```

```
exit
exit
```

The persistency file is pre-allocated, providing space for the maximum number of allowed leases, which avoids file system space issues during normal operation, as follows:

```
*A:PE1>file cf1:\ # dir

Volume in drive cf1 on slot A has no label.

Volume in drive cf1 on slot A is formatted as FAT32

Directory of cf1:\

09/19/2016  04:29p      <DIR>          .ssh/
09/21/2016  01:58p                248513024  dhcp_serv.006
09/21/2016  01:58p                5825024   dhcp_serv.i06
                2 File(s)                254338048 bytes.
                1 Dir(s)                 7759888384 bytes free.

*A:PE1>file cf1:\ #
```

A message is issued to log-id 99 to indicate that the persistence file is ready for use, as follows:

```
*A:PE1# show log log-id 99

=====
Event Log 99
=====
Description : Default System Log
Memory Log contents  [size=500  next event=10722  (wrapped)]
10721 2016/09/21 12:44:58.24 CEST WARNING: SYSTEM #2037 Base dhcp-server Persistence Report
"Persistency event: dhcp-server persistence file ready for use"
```

The **tools dump persistence summary** command provides persistency information. The following example shows that the file cf1:\dhcp_serv.006 is used for storing persistency records for the DHCP server:

```
*A:PE1# tools dump persistence summary

=====
Persistence Summary on Slot A (active)
=====
Client          Location          #Registrations  File State
                Avg Nr Fragments  #Entries        Subsystem State
                File Fill Level  #Entries Queued
-----
dhcp-server     cf1:\dhcp_serv.006  5               ACTIVE
                1.0               5               OK
                0%               0
-----
Total for cf1:   3% in use
-----

*A:PE1#
```

Persistency records are identified using the persistence key. This key is part of the lease state. The following command shows the persistence key for lease 10.11.0.14:

```
*A:PE1# show router 1 dhcp local-dhcp-server "dhcp4-VPRN1" leases 10.11.0.14 detail
```

```

Lease for DHCP server dhcp4-VRPN1 router 1
=====
IP-address           : 10.11.0.14
Lease-state          : stable
Lease started        : 2016/10/17 15:28:36

---snip---

User-db Address Type : N/A
Persistence Key    : 0x00000004
Lease Remaining Hold Time : 0h0m0s

=====
*A:PE1#

```

The DHCP server lease state records can be shown using the following command. This example shows the record for key 0x00000004:

```

*A:PE1# tools dump persistence dhcp-server record 0x00000004
-----
Persistence Record
-----
Client       : dhcp-server
Persist-Key  : 0x00000004
Filename     : cfl:\dhcp_serv.006
Entries      : Index FedHandle Last Update          Action Valid
                004289 0x00000079 2016/10/17 14:48:46 (UTC) UPDATE Yes
Data         : 151 bytes

    type      : V4 lease
    service Id : 1
    server     : dhcp4-VRPN1
    IP         : 10.11.0.14
    MAC        : 00:00:00:01:01:03
    XID        : 0x00000020
    state      : stable
    lease mode : ET
    start time : 2016/10/17 13:28:36 (UTC)
    last renew : 2016/10/17 14:48:46 (UTC)
    expires    : 2016/10/17 14:58:46 (UTC)
    failctrl   : local
    opt60 len  : 0
    opt61 len  : 0
    opt82 len  : 0
    sticky name:
*A:PE1#

```

DHCP server lease state persistency is typically used together with subscriber management persistency if the DHCP server and subscriber management functions are managed by the same network node; see the [IPv4 DHCP Hosts](#) chapter.

Configure a Sticky Lease

The following command creates a sticky lease with name me-010101, using MAC address 00:00:00:01:02:02 and IP address 10.11.0.20:

```

*A:PE1# tools perform router 1 dhcp local-dhcp-server "dhcp4-VRPN1" pool "pool-1" create-
sticky-lease me-010202 mac 00:00:00:01:02:02 requested-ip-address 10.11.0.20

```

```
=====
Sticky lease creation result
=====
Result           : Success
IP-address       : 10.11.0.20
Lease-state      : sticky
Lease started    : 2016/10/17 17:07:00
Remaining LifeTime : N/A
Sticky-lease Host Name : me-010202
MAC address      : 00:00:00:01:02:02
Persistence Key   : N/A
=====
*A:PE1#
```

No user database may be assigned to the DHCP server to create sticky leases.

A **clear** command can be used to delete a sticky lease, as follows:

```
clear router 1 dhcp local-dhcp-server "dhcp4-VRPN1" sticky-leases hostname "me-010202"
```

Operation and Verification

The following command shows all DHCP servers defined in the system. The maximum and active number of leases are shown. The router and services where the DHCP servers are hosted are listed, together with the server name and an indication whether this server is in- or out-of-service.

```
*A:PE1# show router dhcp servers all

=====
Overview of DHCP Servers
=====
Active Leases:      5
Maximum Leases:     159744

Router              Server                      Admin State
-----
Service: 1          dhcp4-VRPN1                  inService
=====
*A:PE1#
```

The following command shows all leases currently allocated by DHCP server dhcp4-VRPN1 in VRPN-1. In this example, the leases for the DHCP and PPP clients are all "stable". Sticky leases are always shown, even when they are not online.

```
*A:PE1# show router 1 dhcp local-dhcp-server "dhcp4-VRPN1" leases

=====
Leases for DHCP server dhcp4-VRPN1 router 1
=====
IP Address      Lease State      Mac Address      Remaining      Clnt  Fail
  PPP user name/Opt82 Circuit Id      LifeTime      Type  Ctrl
  User-db/Sticky-lease Hostname
-----
10.10.0.11      stable           00:00:00:01:01:01 0h9m16s      dhcp  local
10.10.0.12      stable           00:00:00:01:01:02 0h7m36s      dhcp  local
10.11.0.14      stable           00:00:00:01:01:03 0h9m9s       dhcp  local
```

```
10.11.0.17    stable      00:00:00:00:00:33 0h59m55s    ppp    local
  PE1|1|int-GRP|1/1/1:1
10.11.0.20    sticky      00:00:00:01:02:02 N/A          dhcp   N/A
```

```
me-010202
```

```
-----
5 leases found
```

```
=====
*A:PE1#
```

The following command shows the leases on the same server allocated from the 10.11.0.0/24 subnet:

```
*A:PE1# show router 1 dhcp local-dhcp-server "dhcp4-VRPN1" leases 10.11.0.0/24
```

```
=====
Leases for DHCP server dhcp4-VRPN1 router 1
```

IP Address	Lease State	Mac Address	Remaining LifeTime	Clnt Type	Fail Ctrl
PPP user name/Opt82 Circuit Id					
User-db/Sticky-lease Hostname					
10.11.0.14	stable	00:00:00:01:01:03	0h7m31s	dhcp	local
10.11.0.18	stable	00:00:00:00:00:33	0h59m40s	ppp	local
PE1 1 int-GRP 1/1/1:1					
10.11.0.20	sticky	00:00:00:01:02:02	N/A	dhcp	N/A

```
me-010202
```

```
-----
3 leases found
```

```
=====
*A:PE1#
```

The following command shows the details of a single lease:

```
*A:PE1# show router 1 dhcp local-dhcp-server "dhcp4-VRPN1" leases 10.11.0.18/32 detail
```

```
=====
Lease for DHCP server dhcp4-VRPN1 router 1
```

```
=====
IP-address           : 10.11.0.18
Lease-state          : stable
Lease started        : 2016/10/17 17:15:57
Last renew           : N/A
Remaining LifeTime    : 0h57m55s
Remaining Potential Exp. Time: 0h0m0s
Sticky-lease Host Name : N/A
MAC address          : 00:00:00:00:00:33
Xid                   : 0x8bf01670
Failover Control      : local
Client Type           : ppp
User-db Host Name     : N/A
User-db Address Type  : N/A
Persistence Key       : 0x00000005
Opt82 Hex Dump        : (length=71)
                      : 52 45 01 15 50 45 31 7c 31 7c 69 6e 74 2d 47 52
                      : 50 7c 31 2f 31 2f 31 3a 31 02 06 00 00 00 00 00
                      : 33 09 24 00 00 19 7f 1f 02 06 00 00 00 00 00 33
                      : 06 01 01 01 03 50 45 31 03 04 00 00 00 01 04 07
                      : 31 2f 31 2f 31 3a 31
Opt82 Circuit Id      : PE1|1|int-GRP|1/1/1:1
Opt82 Remote Id       : (hex) 00 00 00 00 00 33
Opt82 VS System       : PE1
```



```
Opt82 VS Clnt MAC      : 00:00:00:00:00:33
Opt82 VS Service       : (hex) 00 00 00 01
Opt82 VS SAP           : 1/1/1:1
Opt82 VS String        :
Opt82 VS PPPoE Session ID :
Opt60 Hex Dump         : (length=10)
                       : 41 4c 55 37 58 58 58 53 42 4d
Lease Remaining Hold Time : 0h0m0s
```

```
=====
*A:PE1#
```

Troubleshooting

The following command shows summary data for the DHCP server:

```
*A:PE1# show router 1 dhcp local-dhcp-server "dhcp4-VRPN1" summary
```

```
=====
DHCP server dhcp4-VRPN1  router 1
=====
Admin State      : inService
Operational State : inService
Persistency State : ok
User Data Base   : N/A
Use gateway IP address : enabled (scope pool)
Use pool from client : enabled
Send force-renewals : disabled
Creation Origin   : manual
Lease Hold Time   : 0h10m0s
Lease Hold Time For : (Not specified)
User-ident        : mac-circuit-id

Failover Admin State : outOfService
Failover Oper State  : shutdown
Failover Persist Key : 0x00000003
Administrative MCLT   : 0h10m0s
Operational MCLT     : 0h10m0s
Startup wait time    : 0h2m0s
Partner down delay   : 23h59m59s
Ignore MCLT          : disabled

-----
Pool name : pool-1
-----
Failover Admin State : outOfService
Failover Oper State  : shutdown
Failover Persist Key : 0x00000001
Administrative MCLT   : 0h10m0s
Operational MCLT     : 0h10m0s
Startup wait time    : 0h2m0s
Partner down delay   : 23h59m59s
Ignore MCLT          : disabled

-----
Subnet           Free    %    Stable  Declined  Offered  Rem-pend  Drain
-----
10.10.0.0/25     0      0%   2       0         0        0        N
10.11.0.0/24     241    98%   3       0         0        0        N
Totals for pool  241    97%   5       0         0        0
-----
```

```

Pool name : pool-2
-----
Failover Admin State : outOfService
Failover Oper State  : shutdown
Failover Persist Key : 0x00000007
Administrative MCLT  : 0h10m0s
Operational MCLT     : 0h10m0s
Startup wait time    : 0h2m0s
Partner down delay   : 23h59m59s
Ignore MCLT          : disabled
-----
Subnet              Free    %    Stable  Declined Offered  Rem-pend Drain
-----
10.20.0.0/16        100    100% 0        0        0        0        N
Totals for pool      100    100% 0        0        0        0
-----

Totals for server    341    98% 5        0        0        0
-----

Interface associations
Interface              Admin
-----
int-VRPN1-DHCPv4      Up
-----

Local Address Assignment associations
Group interface        Admin
-----
=====
*A:PE1#

```

The following command shows DHCP server statistics:

```

*A:PE1# show router 1 dhcp local-dhcp-server "dhcp4-VRPN1" server-stats
=====
Statistics for DHCP Server dhcp4-VRPN1 router 1
=====
Rx Discover Packets      : 2449
Rx Request Packets      : 12752
Rx Release Packets      : 53
Rx Decline Packets      : 0
Rx Inform Packets       : 0

Tx Offer Packets        : 177
Tx Ack Packets          : 1184
Tx Nak Packets          : 63
Tx Forcerenew Packets   : 58

Client Ignored Offers   : 0
Leases Timed Out        : 2

Dropped Bad Packet      : 11205
Dropped Invalid Type    : 0
Dropped No User Database : 0
Dropped Unknown Host    : 0
Dropped User Not Allowed : 0
Dropped Lease Not Ready : 0
Dropped Lease Not Found : 5
Dropped Not Serving Pool : 2297
Dropped Invalid User     : 0
Dropped Overload        : 0

```

```
Dropped Persistence Overload : 0
Dropped Generic Error       : 0
Dropped Destined To Other   : 0
Dropped Address Unavailable : 300
Dropped Max Leases Reached  : 0
Dropped Server Shutdown     : 0
Dropped No Subnet For Fixed IP: 0
Dropped Duplicate From Diff GI: 0
Dropped busy primary audit   : 0
Dropped transmission failed  : 0
```

```
Rx Internal Requests        : 0
Rx Internal Releases        : 0
Dropped Internal w/LUDB     : 0
Dropped Internal w/Failover : 0
Dropped Internal w/Conflicts : 0
```

Failover statistics

```
-----
Dropped Invalid Packets      : 0
Failover Shutdown           : 0
Lease Already Expired        : 0
Maximum Lease Count Reached  : 0
Subnet Not Found             : 0
Range Not Found              : 0
Host Conflict                : 0
Address Conflict             : 0
Peer conflict                : 0
Persistence congestion       : 0
No Lease Hold Time Configured : 0
Invalid Prefix Length        : 0
Lease Not Found              : 0
=====
```

*A:PE1#

The following command shows extended server statistics:

```
*A:PE1# show router 1 dhcp local-dhcp-server "dhcp4-VRPN1" pool-ext-stats
```

Extended pool statistics for server "dhcp4-VRPN1"

```
=====
Current      Peak      Peak Timestamp
-----
Pool pool-1
Local:
  Offered Leases      0      1      10/17/2016 17:15:57
  Stable Leases        5      5      10/17/2016 17:15:57
  Provisioned Addresses 246
  Used Addresses        5      5      10/17/2016 17:21:24
  Free Addresses       241      241      10/17/2016 17:21:24
  Used Pct              3      3      10/17/2016 17:21:24
  Free Pct             97      97      10/17/2016 17:21:24
  Last Reset Time      10/17/2016 15:26:31
-----
Pool pool-2
Local:
  Offered Leases      0      0      10/17/2016 17:22:15
  Stable Leases        0      0      10/17/2016 17:22:15
  Provisioned Addresses 100
  Used Addresses        0      0      10/17/2016 17:22:15
  Free Addresses       100      100      10/17/2016 17:22:15
  Used Pct              0      0      10/17/2016 17:22:15
```

```

Free Pct          100          100          10/17/2016 17:22:15
Last Reset Time   10/17/2016 17:22:15
-----
Number of entries      2
=====
*A:PE1#

```

The following command shows the addresses that are still free in a particular subnet:

```

*A:PE1# show router 1 dhcp local-dhcp-server "dhcp4-VPRN1" free-addresses 10.11.0.0/24
=====
Free addresses
=====
IP Address      Fail Ctrl
-----
10.11.0.11      local
10.11.0.12      local
10.11.0.13      local
10.11.0.15      local
10.11.0.16      local
---snip---
10.11.0.253     local
10.11.0.254     local
-----
No. of free addresses: 241
=====
*A:PE1#

```

The following command shows the DHCP server associations; this is the list of interfaces that the DHCP server is listening on:

```

*A:PE1# show router 1 dhcp local-dhcp-server "dhcp4-VPRN1" associations
=====
DHCP server dhcp4-VPRN1  router 1
=====
Interface associations
Interface                  Admin
-----
int-VPRN1-DHCPv4          Up
-----
Local Address Assignment associations
Group interface            Admin
-----
=====
*A:PE1#

```

The following configuration enables debugging for DHCP server *dhcp4-VPRN1* on VPRN 1:

```

debug
  router "1"
    local-dhcp-server "dhcp4-VPRN1"
    detail-level high
    mode egr-ingr-and-dropped
  exit
exit
exit

```

To ensure that the debug output is sent to a session, the following additional configuration is needed:

```
configure
  log
    log-id 1
      description "Send debug log to the current telnet/ssh session"
      from debug-trace
      to session
      no shutdown
    exit
  exit
exit
```

With this configuration, the following output is shown when the IPE host with MAC address 00:00:00:01:01:01 connects:

```
13 2016/10/17 18:51:12.30 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp4-VRPN1
Rx DHCP Discover

  ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
  siaddr: 0.0.0.0          giaddr: 10.10.0.1
  chaddr: 00:00:00:01:01:01  xid: 0x21

DHCP options:
[82] Relay agent information: len = 23
  [1] Circuit-id: PE1|1|int-GRP|1/1/1:1
[53] Message type: Discover
[255] End

Hex Packet Dump:
01 01 06 00 00 00 00 21 00 00 00 00 00 00 00 00 00 00 00 00 00 0a
---snip---
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
"

14 2016/10/17 18:51:12.30 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp4-VRPN1
lease added for 10.10.0.12 state=offer
"

15 2016/10/17 18:51:12.30 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp4-VRPN1
Tx DHCP Offer to local relay agent 10.10.0.1 vrId=2

  ciaddr: 0.0.0.0          yiaddr: 10.10.0.12
  siaddr: 10.11.11.1        giaddr: 10.10.0.1
  chaddr: 00:00:00:01:01:01  xid: 0x21

DHCP options:
[82] Relay agent information: len = 23
  [1] Circuit-id: PE1|1|int-GRP|1/1/1:1
[53] Message type: Offer
[54] DHCP server addr: 10.11.11.1
[51] Lease time: 600
[1] Subnet mask: 255.255.255.0
[3] Router: 10.10.0.1
[130] Unknown option: len = 9, value = 4d 79 4f 70 74 69 6f 6e 31
[6] Domain name server: length = 8
    1.1.1.1
    1.1.2.2
[150] Unknown option: len = 4, value = 01 01 01 01
[255] End
```

```
Hex Packet Dump:
02 01 06 00 00 00 00 21 00 00 00 00 00 00 00 00 0a 0a 00 0c 0a 0b 0b 01 0a
---snip---
31 7c 31 7c 69 6e 74 2d 47 52 50 7c 31 2f 31 2f 31 3a 31 ff
"

16 2016/10/17 18:51:12.32 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp4-VRPN1
Rx DHCP Request

ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
siaddr: 0.0.0.0          giaddr: 10.10.0.1
chaddr: 00:00:00:01:01:01  xid: 0x21

DHCP options:
[82] Relay agent information: len = 23
      [1] Circuit-id: PE1|1|int-GRP|1/1/1:1
[53] Message type: Request
[50] Requested IP addr: 10.10.0.12
[54] DHCP server addr: 10.11.11.1
[255] End

Hex Packet Dump:
01 01 06 00 00 00 00 21 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 0a
---snip---
31 2f 31 3a 31 ff 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
"

17 2016/10/17 18:51:12.32 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp4-VRPN1
lease update for 10.10.0.12 state=stable
"

18 2016/10/17 18:51:12.52 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp4-VRPN1
Tx DHCP Ack to local relay agent 10.10.0.1 vrId=2

ciaddr: 0.0.0.0          yiaddr: 10.10.0.12
siaddr: 10.11.11.1       giaddr: 10.10.0.1
chaddr: 00:00:00:01:01:01  xid: 0x21

DHCP options:
[82] Relay agent information: len = 23
      [1] Circuit-id: PE1|1|int-GRP|1/1/1:1
[53] Message type: Ack
[54] DHCP server addr: 10.11.11.1
[51] Lease time: 600
[1] Subnet mask: 255.255.255.0
[3] Router: 10.10.0.1
[130] Unknown option: len = 9, value = 4d 79 4f 70 74 69 6f 6e 31
[6] Domain name server: length = 8
      1.1.1.1
      1.1.2.2
[150] Unknown option: len = 4, value = 01 01 01 01
[255] End

Hex Packet Dump:
02 01 06 00 00 00 00 21 00 00 00 00 00 00 00 00 0a 0a 00 0c 0a 0b 0b 01 0a
---snip---
31 7c 31 7c 69 6e 74 2d 47 52 50 7c 31 2f 31 2f 31 3a 31 ff
```

"

When a client terminates its connection, the following output is shown:

```
19 2016/10/17 18:52:05.97 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
Rx DHCP Release

    ciaddr: 10.10.0.12      yiaddr: 0.0.0.0
    siaddr: 0.0.0.0        giaddr: 0.0.0.0
    chaddr: 00:00:00:01:01:01  xid: 0x21

    DHCP options:
    [82] Relay agent information: len = 23
        [1] Circuit-id: PE1|1|int-GRP|1/1/1:1
    [53] Message type: Release
    [54] DHCP server addr: 10.11.11.1
    [255] End

    Hex Packet Dump:
    01 01 06 00 00 00 00 21 00 00 00 00 0a 0a 00 0c 00 00 00 00 00 00 00 00
    ---snip---
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
"
```

```
20 2016/10/17 18:52:05.96 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
lease deleted for 10.10.0.12 (rxd release)
"
```

A PPP user connecting via local address assignment shows the following messages:

```
21 2016/10/17 18:52:15.97 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
Rx internal Request
    primary pool  : pool-2
    ciaddr        : 0.0.0.0
"
```

```
22 2016/10/17 18:52:15.97 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
lease added for 10.20.0.22 state=internal
"
```

When this user terminates the PPP session, the following messages are shown:

```
23 2016/10/17 18:52:26.41 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
Rx internal Release
    ciaddr        : 10.20.0.22
"
```

```
24 2016/10/17 18:52:26.41 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
lease deleted for 10.20.0.22 (rxd internal release)
"
```

Renumbering – Subnet Mask Change

The baseline configuration has the subnet 10.10.0.0/25 defined, providing address space for up to 126 addresses. The range that the server can take free addresses from starts at 10.10.0.11, and ends at 10.10.0.126, excluding the 10.10.0.61 to 10.10.0.70 sub-range.

Assume that subnet 10.10.0.128/25 was removed from a different BNG, and now can be used in this BNG. This subnet can be aggregated with the 10.10.0.0/25 network to become subnet 10.10.0.0/24. At the same time, the requirement is to not disrupt services for already connected users.

The following steps are required at the DHCP server:

- ensure that **no force-renews** is active
- delete the original subnet
- create the new subnet

Preventing the DHCP server from sending force-renew messages is important so that already connected users do not lose their connection, as follows:

```
*A:PE1# configure service vprn 1 dhcp local-dhcp-server dhcp4-VPRN1 no force-renews
```

The following command deletes the original subnet:

```
*A:PE1# configure service vprn 1 dhcp local-dhcp-server "dhcp4-VPRN1"
pool "pool-1" no subnet 10.10.0.0/25
```

Leases are not deleted when the subnet is deleted; their status changes from *stable* to *removePending*, as follows:

```
*A:PE1# show router 1 dhcp local-dhcp-server "dhcp4-VPRN1" leases

=====
Leases for DHCP server dhcp4-VPRN1 router 1
=====
IP Address      Lease State      Mac Address      Remaining      Clnt  Fail
  PPP user name/Opt82 Circuit Id      LifeTime      Type  Ctrl
  User-db/Sticky-lease Hostname
-----
10.10.0.11      removePending    00:00:00:01:01:01 1h57m25s      dhcp  local
10.10.0.12      removePending    00:00:00:03:01:01 0h57m30s      ppp   local
  PE1|1|int-GRP|1/1/1:1
-----
2 leases found
=====
*A:PE1#
```

This status change is also shown in the debug log, as follows:

```
132 2016/10/14 14:10:57.66 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp4-VPRN1
lease 10.10.0.11 scheduled for removal
"

133 2016/10/14 14:10:57.66 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp4-VPRN1
lease 10.10.0.12 scheduled for removal
```



```
"
134 2016/10/14 14:10:57.66 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
lease 10.10.0.11 scheduled for removal
"

135 2016/10/14 14:10:57.66 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
lease 10.10.0.12 scheduled for removal
"
```

Users trying to renew or connect will not get an address as long as no new subnet is defined.

Create the new 10.10.0.0/24 subnet, with the new address range starting at 10.10.0.11 and ending at 10.10.0.254, as follows. The original exclusion range still applies, but a new exclusion address 10.10.0.129 is added, to be described later:

```
configure
  service
    vprn 1 customer 1 create
    dhcp
      local-dhcp-server dhcp4-VPRN1 create
      use-gi-address scope pool
      no force-renews
      pool "pool-1" create
      options
        dns-server 1.1.1.1 1.1.2.2
        lease-time hrs 2
        custom-option 150 address 1.1.1.1
      exit
      subnet 10.10.0.0/24 create
      options
        subnet-mask 255.255.255.0
        default-router 10.10.0.1
      exit
      exclude-addresses 10.10.0.61 10.10.0.70
      exclude-addresses 10.10.0.129 10.10.0.129
      address-range 10.10.0.11 10.10.0.254
    exit
  exit
exit
exit
exit
exit
exit
```

Leases that were in use before return to the *stable* state, if they are not in the exclusion range, as follows:

```
*A:PE1# show router 1 dhcp local-dhcp-server "dhcp4-VPRN1" leases

=====
Leases for DHCP server dhcp4-VPRN1 router 1
=====
```

IP Address	Lease State	Mac Address	Remaining LifeTime	Clnt Type	Fail Ctrl
10.10.0.11	stable	00:00:00:01:01:01	1h49m43s	dhcp	local
10.10.0.12	stable	00:00:00:03:01:01	0h49m48s	ppp	local

```
-----
PE1|1|int-GRP|1/1/1:1
```

```
-----
2 leases found
=====
```

```
*A:PE1#
```

The following command adds the 10.10.0.129/25 address to the *int-SUB* subscriber interface, so that offers in the 10.10.0.128/25 range will not get dropped by the relay agent. Any address in the 10.10.0.128/25 subnet could be used; the lowest one is used in this example. Because this address is in use by the subscriber interface, it must be added to the exclusion list in the DHCP server, as follows:

```
*A:PE1# configure service vprn 1 subscriber-interface int-SUB address 10.10.0.129/25
```

This configuration ensures service continuity for already connected subscribers. They will get their new /24 subnet when they renew or rebind their lease. No change is needed at the relay agent.

Merging the two subnets at the subscriber interface is only possible with a service interruption, because the subscriber interface addresses cannot be deleted when leases are in use. Also the Gi address configured in the dhcp context must be deleted.

```
*A:PE1>config>service>vprn>sub-if# no address 10.10.0.1/25
INFO: PIP #1398 Cannot delete/change address when managed ARPs or leases defined for this
subnet exist - 1 managed-arps or leases exist
*A:PE1>config>service>vprn>sub-if#
```

To also merge the subnets at the subscriber interface, all the leases in these subnets must be deleted. When the address defined at the subscriber interface is also used as the Gi address by the relay agent, the Gi address must be removed first. Then, the subscriber interface address can be deleted and recreated with the correct netmask. Also, the Gi address can be redefined after that. The changes at the DHCP server are similar to the ones defined previously.

Renumbering – Subnet Migration

The following changes to the baseline configuration have to be made to support the migration of DHCP clients from the 10.10.0.0/25 and 10.11.0.0/24 subnets to the 10.12.0.0/20 subnet. For that purpose, the 10.10.0.0/25 and the 10.11.0.0/24 subnets have the keyword **drain** added, so that leases in the corresponding address ranges will not be extended.

This new 10.12.0.0/20 subnet has a new subnet mask, a new default router, and three address ranges. New clients connecting will automatically get addresses from this new subnet. To ensure existing clients will not lose their connection, the **use-gi-address scope** is set to **pool**, so that they get a new lease from the new subnet when renewing or rebinding.

In scenarios where lease times are long (an order of magnitude of months or even years), it can take a considerable time before all clients have a lease in the new subnet. Having DHCP clients supporting force-renew can help speed up the migration process. The following configuration has force-renews enabled.

Address 10.12.0.1 is used as the default router for this subnet, so this address is added to the *int-SUB* subscriber-interface. This address will later be used as the Gi address.

```
configure
  service
    vprn 1
      dhcp
        local-dhcp-server "dhcp4-VRN1" create
          use-gi-address scope pool
          force-renews
```

```
pool "pool-1" create
  options
    dns-server 1.1.1.1 1.1.2.2
    lease-time hrs 2
  exit
  subnet 10.10.0.0/25 create
    drain
    options
      subnet-mask 255.255.255.0
      default-router 10.10.0.1
    exit
    address-range 10.10.0.11 10.10.0.12
  exit
  subnet 10.11.0.0/24 create
    drain
    options
      subnet-mask 255.255.255.0
      default-router 10.10.0.1
    exit
    address-range 10.11.0.11 10.11.0.254
  exit
  subnet 10.12.0.0/20 create
    options
      subnet-mask 255.255.240.0
      default-router 10.12.0.1
    exit
    address-range 10.12.0.10 10.12.12.255
    address-range 10.12.13.1 10.12.14.255
    address-range 10.12.15.10 10.12.15.254
  exit
exit
no shutdown
exit
exit
subscriber-interface "int-SUB"
  address 10.12.0.1/20
exit
exit
exit
exit
```

The following command shows that the original subnets are in the drained state:

```
*A:PE1# show router 1 dhcp local-dhcp-server "dhcp4-VPRN1" summary
=====
DHCP server dhcp4-VPRN1  router 1
=====
Admin State           : inService
Operational State     : inService
Persistency State     : ok

---snip---

-----
Pool name : pool-1
-----
Failover Admin State  : outOfService
Failover Oper State   : shutdown
Failover Persist Key  : 0x00000001
Administrative MCLT   : 0h10m0s
Operational MCLT     : 0h10m0s
Startup wait time     : 0h2m0s
Partner down delay    : 23h59m59s
```

Ignore MCLT	: disabled						
Subnet	Free	%	Stable	Declined	Offered	Rem-pend	Drain
10.10.0.0/25	2	100%	0	0	0	0	Y
10.11.0.0/24	244	100%	0	0	0	0	Y
10.12.0.0/20	4072	99%	3	0	0	0	N
Totals for pool	4318	99%	3	0	0	0	

Totals for server	4318	99%	3	0	0	0	

*A:PE1#							

Because the DHCP server is configured with force-renew, connected clients are sent a force-renew message. In response, the client tries to extend its lease by sending a request message using the current address. The DHCP server sends a negative-acknowledgement (NAK) to the requesting client, because the subnet is in the draining state. This forces the client to go through the DHCP initialization process; a new DORA message sequence is initiated. Therefore, the client gets a free address in the new subnet, with a new netmask, and a new default router, as follows. The same DNS servers are offered, because these pool options were not changed.

```

1 2016/10/15 19:19:36.04 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp4-VRPN1
Tx DHCP ForceRenew to client at 10.10.0.12 vrId=2

    ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
    siaddr: 10.11.11.1       giaddr: 0.0.0.0
    chaddr: 00:00:00:01:01:01  xid: 0x1f

    DHCP options:
    [53] Message type: ForceRenew
    [54] DHCP server addr: 10.11.11.1
    [255] End
"

2 2016/10/15 19:19:36.05 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp4-VRPN1
Rx DHCP Request

    ciaddr: 10.10.0.12       yiaddr: 0.0.0.0
    siaddr: 0.0.0.0          giaddr: 0.0.0.0
    chaddr: 00:00:00:01:01:01  xid: 0x1f

    DHCP options:
    [53] Message type: Request
    [255] End
"

3 2016/10/15 19:19:36.05 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp4-VRPN1
subnet is draining
Tx DHCP Nak to client 10.10.0.12 vrId=2 (via snooping function)

    ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
    siaddr: 0.0.0.0          giaddr: 0.0.0.0
    chaddr: 00:00:00:01:01:01  xid: 0x1f

    DHCP options:

```

```
[53] Message type: Nak
[54] DHCP server addr: 10.11.11.1
[255] End
"

4 2016/10/15 19:19:36.06 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp4-VRPN1
Rx DHCP Discover

    ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
    siaddr: 0.0.0.0          giaddr: 10.10.0.1
    chaddr: 00:00:00:01:01:01  xid: 0x1f

    DHCP options:
    [53] Message type: Discover
    [255] End
"

5 2016/10/15 19:19:36.06 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp4-VRPN1
lease added for 10.12.0.17 state=offer
"

6 2016/10/15 19:19:36.06 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp4-VRPN1
Tx DHCP Offer to local relay agent 10.10.0.1 vrId=2

    ciaddr: 0.0.0.0          yiaddr: 10.12.0.17
    siaddr: 10.11.11.1       giaddr: 10.10.0.1
    chaddr: 00:00:00:01:01:01  xid: 0x1f

    DHCP options:
    [53] Message type: Offer
    [54] DHCP server addr: 10.11.11.1
    [51] Lease time: 7200
    [1] Subnet mask: 255.255.240.0
    [3] Router: 10.12.0.1
    [6] Domain name server: length = 8
        1.1.1.1
        1.1.2.2
    [150] Unknown option: len = 4, value = 01 01 01 01
    [255] End
"

7 2016/10/15 19:19:36.07 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp4-VRPN1
Rx DHCP Request

    ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
    siaddr: 0.0.0.0          giaddr: 10.10.0.1
    chaddr: 00:00:00:01:01:01  xid: 0x1f

    DHCP options:
    [53] Message type: Request
    [50] Requested IP addr: 10.12.0.17
    [54] DHCP server addr: 10.11.11.1
    [255] End
"

8 2016/10/15 19:19:36.07 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp4-VRPN1
lease update for 10.12.0.17 state=stable
"
```

```

9 2016/10/15 19:19:36.24 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp4-VRN1
Tx DHCP Ack to local relay agent 10.10.0.1 vrId=2

ciaddr: 0.0.0.0          yiaddr: 10.12.0.17
siaddr: 10.11.11.1       giaddr: 10.10.0.1
chaddr: 00:00:00:01:01:01 xid: 0x1f

DHCP options:
[53] Message type: Ack
[54] DHCP server addr: 10.11.11.1
[51] Lease time: 7200
[1] Subnet mask: 255.255.240.0
[3] Router: 10.12.0.1
[6] Domain name server: length = 8
    1.1.1.1
    1.1.2.2
[150] Unknown option: len = 4, value = 01 01 01 01
[255] End
"

```

When the original DHCP server subnets are fully drained, they can be safely deleted. However, deleting a subnet from a pool before it is fully drained results in the remaining leases being scheduled for removal, as follows:

```

140 2016/10/10 15:12:42.87 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: dhcp4-VRN1
lease 10.11.0.11 scheduled for removal
"

```

The number of leases pending for removal can be shown using following command:

```

*A:PE1# show router 1 dhcp local-dhcp-server "dhcp4-VRN1" summary
=====
DHCP server dhcp4-VRN1  router 1
=====
Admin State      : inService
Operational State : inService

---snip---

-----
Pool name : pool-1
-----
Failover Admin State : outOfService
Failover Oper State  : shutdown
Failover Persist Key : 0x00000001
Administrative MCLT   : 0h10m0s
Operational MCLT      : 0h10m0s
Startup wait time     : 0h2m0s
Partner down delay    : 23h59m59s
Ignore MCLT           : disabled
-----
Subnet           Free    %    Stable  Declined  Offered  Rem-pend  Drain
-----
10.11.0.0/24     244    100%  0        0         0        0        Y
10.12.0.0/20     4075   100%  0        0         0        0        N
Totals for pool  4319   100%  0        0         0        0
-----
Not subnet related                                Rem-pend
-----
1

```

```
-----  
Totals for server      4319      100% 0          0          0          1  
---snip---  
=====
```

This lease will be deleted when the lease expires.
The relay agent can then have the Gi address updated (10.12.0.1) and the old subnets can be removed from the group interface.

Conclusion

SR OS supports DHCPv4 servers on any routing instance (VPRN or base router), offering pool, subnet, and address management, combined with configuration parameter management and persistency.

Diameter Application NASREQ

This chapter provides information about Diameter Application NASREQ.

Topics in this chapter include:

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

Applicability

The information and configuration in this chapter are based on SR OS Release 19.10.R1.



Note:

This chapter covers the Diameter NASREQ application in combination with the Diameter base protocol implementation that is available from SR OS Release 16.0.R4 onward (configured in the aaa CLI context as diameter node). The legacy Diameter base implementation (configured in the aaa CLI context as diameter-peer-policy) is supported in maintenance mode only, without any further feature enhancement planned. Nokia recommends using or transitioning to the new Diameter base protocol implementation. See also the [Diameter Base Protocol: Establishing a Diameter Peer Connection](#) chapter of this Advanced Configuration Guide

Overview

NASREQ is defined in RFC 7155, *Diameter Network Access Server Application*, and uses the Diameter base protocol defined in RFC 6733, *Diameter Base Protocol*. The purpose of NASREQ in SR OS is to provide subscriber authentication and authorization. NASREQ provides functionality that is also available via RADIUS but uses the Diameter protocol instead.

NASREQ complements the other Diameter applications supported in SR OS:


- Gx provides advanced authorization capabilities for subscribers and usage monitoring, and interfaces with a Policy and Charging Rules Function (PCRF).
- Gy or Diameter Credit Control Application (DCCA) provides on-line charging functionality, and interfaces with an On-line Charging Server (OCS).
- NASREQ provides subscriber authentication and authorization, and interfaces with a AAA server.

These three Diameter applications use the Diameter base protocol, which is described in the [Diameter Base Protocol: Establishing a Diameter Peer Connection](#) chapter of this Advanced Configuration Guide. That chapter also describes the configuration of the Diameter base protocol, which is very similar for all supported Diameter applications, and is not repeated in this chapter.

When a subscriber connects to the BNG, NASREQ is triggered, as would have been the case with RADIUS authentication. [Figure 29: NASREQ trigger](#) shows a sample call flow for NASREQ and Gx

applications when an IPoE IPv4 subscriber session connects. The supported NASREQ messages are Authentication and Authorization Request (AAR) and Authentication and Authorization Answer (AAA), and the BNG assumes that the AAA server does not maintain session state. Therefore, there is no need for the BNG to send a message to the AAA server to indicate when the subscriber has ended the session. The BNG negotiates with the AAA server that it expects a stateless behavior by sending the Auth-Session-State AVP with value NO_STATE_MAINTAINED (1) as defined in RFC 6733, and the AAA server must confirm that by sending back the AVP with the same value.

Figure 29: NASREQ trigger



As shown in [Figure 29: NASREQ trigger](#), NASREQ is triggered when the DHCP discover is received, while Gx, using the Credit Control Request (CCR-I) and Credit Control Answer (CCA-I) messages, is triggered at the end, after the IP address allocation.

Supported NASREQ AVPs in AAA are listed in the *7450 ESS, 7750 SR, and VSR Triple Play Service Delivery Architecture Guide*, and include AVPs for subscriber identification, profiles, service selection, address allocation, and so on.

All the NASREQ AVPs are also supported by RADIUS, and their meaning is described in the *7450 ESS, 7750 SR, and VSR RADIUS Attributes Reference Guide*.

Some AVPs, such as SLA profile and subscriber profile strings can be obtained from both NASREQ in the authentication phase and from Gx in the authorization phase. In this case, the last received value from Gx is used for the subscriber session creation.

The AAA server authenticates a subscriber based on a password. For IPoE subscribers, the password is configured in the BNG because DHCP cannot provide a subscriber password, and for PPP, the authentication is performed through PAP or CHAP. PPPoE PADI authentication via NASREQ is not supported in SR OS.

Configuration

The configuration of NASREQ authentication is performed in four steps:

1. Configure a Diameter node
2. Configure a Diameter application policy.
3. Assign the Diameter application policy.
4. Apply a Python policy (optional).

In the first step, the Diameter node configures the Diameter base protocol and is described in the [Diameter Base Protocol: Establishing a Diameter Peer Connection](#) chapter. As explained in this chapter, a Diameter peer can support multiple applications, such as NASREQ and Gx:

```
# show aaa diameter-node "bng-gx.realm-1.com" routing-table

=====
Routes
=====
Realm-Name
Application    Pref. Id  Server-Identifier
-----
realm-2.com
    nasreq gx    10    1    dra-1.realm-2.com
realm-2.com
    nasreq gx    20    2    dra-2.realm-2.com
-----
No. of routes: 2
=====
```

For the example in this chapter, a direct connection exists between the Diameter NASREQ client (BNG) and two NASREQ servers (aaa-1 and aaa-2):

```
configure
aaa
    diameter
        node "bng-nasreq.nokia.com" create
            peer index 1 "aaa-1.nokia.com" create
                address 2001:db8:2:6::1
                preference 10
                no shutdown
            exit
            peer index 2 "aaa-2.nokia.com" create
                address 172.16.7.2
                preference 20
                no shutdown
            exit
        exit
    exit
exit
```

```
# show aaa diameter-node "bng-nasreq.nokia.com" peers

=====
Peers
=====
Host identity          Status          Default Preference Active
-----
aaa-1.nokia.com        I-Open         No           10         Yes
aaa-2.nokia.com        I-Open         No           20         Yes
-----
No. of peers: 2
=====
```

```
# show aaa diameter-node "bng-nasreq.nokia.com" routing-table
```

```
=====
```

```
Routes
```

```
=====
```

Realm-Name	Application	Pref.	Id	Server-Identifier
nokia.com	nasreq	10	1	aaa-1.nokia.com
nokia.com	nasreq	20	2	aaa-2.nokia.com

```
-----
```

```
No. of routes: 2
```

```
=====
```

When the transport connection for one of the peers is down, the remaining peer is used for authentication. When multiple peers are available, the peer selection is based on the lowest preference value or in case of equal preference, based on the lowest index number. In the example peer aaa-1.nokia.com is used when both peers are up.

In the second step, the Diameter application policy configures the NASREQ application. For example (info detail):

```
configure
  subscriber-mgmt
    diameter-application-policy "diam-nasreq-1"
      description "Diameter application policy - NASREQ"
      application nasreq
      diameter-node "bng-nasreq.nokia.com" destination-realm "nokia.com"
      on-failure failover enabled handling retry-and-terminate
      tx-timer 10
      nasreq
        mac-format "aa:"
        password "YqXTV45qof/g0Y1WFhJbKjjkhg==" hash2
        user-name-format mac
        no user-name-operation
        include-avp
          no called-station-id
          no calling-station-id
          circuit-id
          no imei
          no nas-port
          nas-port-id prefix-type none suffix-type none
          no nas-port-type
          no rat-type
          remote-id
          no user-location-info
        exit
      exit
    exit
  exit
exit
```

A Diameter application policy handles exact one application: NASREQ, Gx, or Gy. In this example, we use **application nasreq** and configure application specific parameters in the **nasreq** context.

A Diameter node is configured in the application policy to select the diameter routing context in which the NASREQ messages should be forwarded to the AAA server: **diameter-node "bng-nasreq.nokia.com" destination-realm "nokia.com"**. The node is identified with its origin host **"bng-nasreq.nokia.com"** and must be configured in the **aaa** context. The destination realm specifies the realm of the NASREQ server

and is used in the Destination-Realm AVP of the application messages. In the example, client and server are in the same realm.

The **on-failure** session failure handling specifies the Diameter application behavior when no answer is received on a request or an answer is received with a protocol error in the result code (except for `DIAMETER_UNABLE_TO_DELIVER` (3002) and `DIAMETER_TOO_BUSY` (3004) that are handled in the Diameter node). Possible configuration options are:

- The subscriber sessions must be authenticated (that is, accepted) in case of failure. NASREQ should not retransmit the AA-Request message.

```
on-failure failover disabled handling continue
```

- The subscriber sessions must be rejected in case of failure. NASREQ should not retransmit the AA-Request message.

```
on-failure failover disabled handling terminate
```

- NASREQ should retransmit the AA-Request and reject the subscriber session when the second attempt also fails (this is the configuration used in the example):

```
on-failure failover enabled handling retry-and-terminate
```

The **tx-timer** is the time that the Diameter application waits to get an answer from the AAA server before applying the configured failure handling. In the example, no **tx-timer** is configured, so the default of 10 seconds applies.

The application specific parameters are configured in the **nasreq** context:

- The format of the MAC address to be used in all NASREQ AVPs with a MAC is configured with the **mac-format** command. The default format is `ab:cd:ef:01:02:03`
- **password** configures the authentication password to be used for IPoE subscriber sessions. For PPP subscribers sessions, the credentials are provided in the PAP or CHAP authentication.
- the AAA server checks the identity of the subscriber based on the User-Name AVP. For PPPoE subscriber sessions, the PAP/CHAP username is copied in the User-Name AVP. For IPoE subscriber sessions, the username can be configured with the command **user-name-format**. Options for the username format are, for example:
 - MAC address (with or without giaddr)
 - circuit-id from the relay agent information (for example, DHCP Option 82 for IPv4)
 - information from DHCP option 60 and 61 (which contain the Client-id and Vendor-Class information)
 - and NAS port Id.
- For both IPoE and PPPoE subscriber sessions, the username format can optionally be manipulated with the **user-name-operation** command to add, remove, or replace a domain name.
- Several optional AVPs can be included in the NASREQ AA-Request message and are configured in the **include-avp** context. In the example, **circuit-id**, **remote-id**, and **nas-port-id** are added.

In the third and last mandatory step, the Diameter application policy must be applied using the **diameter-auth-policy name** command to any of the following:

- a VPRN or IES group interface
- a local user database (LUDB) IPoE or PPP host

- a VPLS capture SAP
- a gtp apn policy

A diameter authentication policy is mutually exclusive with a radius authentication policy:

```
>config>subscr-mgmt>loc-user-db>ipoe>host# auth-policy "auth-policy-1"
MINOR: SVCMGR #1558 cannot configure a nasreq and a radius policy together
```

The fourth optional configuration step enables the manipulation of AA-Request and AA-Answer NASREQ messages with a Python script. A sample script that copies the User-Name AVP into the Subscription-Id AVP (which is not defined for NASREQ) in AA-Request messages is as follows:

```
# show python python-script "nasreq-1" source-in-use

=====
Python script "nasreq-1"
=====
Admin state   : inService
Oper state    : inService
Primary URL   : ftp://*:~@10.1.1.1/./python/diameter/nasreq-1.py
Secondary URL : (Not Specified)
Tertiary URL  : (Not Specified)
Active URL    : primary

-----
Source (dumped from memory)
-----
1 from binascii import *
2 from alc import diameter
3
4 def getint_b2a_hex(val):
5     return int(b2a_hex(val),16)
6
7 def byte_to_binary(n):
8     return ''.join(str((n & (1 << i)) and 1) for i in reversed(range(8)))
9
10 def hex_to_binary(h):
11     return ''.join(byte_to_binary(ord(b)) for b in unhexlify(h))
12
13 def checkbitset(byte,index):
14     return ((byte&(1<<index))!=0)
15
16 def checkRequestOrReply():
17     if checkbitset(int(hex_to_binary(b2a_hex(diameter.flags))),7) is True:
18         return 'R'
19     else :
20         return 'A'
21
22 if getint_b2a_hex(diameter.code) == 265 and checkRequestOrReply() == 'R':
23     try:
24         username = diameter.get_avps(1,0)[0][1]
25         if username != "":
26             diameter.set_grouped_avps(443,0,['@', {(450,0): [('@', '\x00\x00\x00\x04')],
(444,0): [('@',str(username))]}]])
27     except Exception, err: print "Python FAILED to fetch username"
=====
```

Configuration to activate the Python script for the example:

```
configure
python
```

```
python-script "nasreq-1" create
  description "Diameter NASREQ - AAR: copy User-Name in Subscription-Id"
  primary-url "ftp://*:~@10.1.1.1/./python/diameter/nasreq-1.py"
  no shutdown
exit
python-policy "py-nasreq-1" create
  description "Diameter NASREQ - Python"
  diameter aar direction egress script "nasreq-1"
exit
exit
aaa
  diameter
    node "bng-nasreq.nokia.com" create
      python-policy "py-nasreq-1"
      ---snip---
    exit
  exit
exit
```

The Python policy must be configured in the Diameter node and can act on all Diameter base and application messages forwarded or received on any of the peers in the node. The Python policy defines on which messages a script should be activated. In the example, the script "nasreq-1" operates on the AAR messages in the egress direction. In a similar way, Python scripts can also be applied to AAA messages in the ingress direction.

Troubleshooting

The operational state of the Diameter application policy can be shown with the following show command:

```
# show subscriber-mgmt diameter-application-policy "diam-nasreq-1"

=====
DIAMETER application policy "diam-nasreq-1"
=====
Description                : Diameter application policy - NASREQ
Session failover           : enabled
Failover handling          : retry-and-terminate
Peer policy                 : (Not Specified)
Diameter node              : bng-nasreq.nokia.com
Destination-realm          : nokia.com
Application                 : nasreq
Tx timer (s)               : 10
Last management change     : 11/25/2019 13:33:46
-----
NASREQ
-----
Include AVP                 : circuit-id
                           : remote-id
                           : nas-port-id
NAS-Port-Id prefix type    : none
NAS-Port-Id suffix type    : none
User name format           : mac
User name operation        : no-operation
MAC address format         : aa:
Last management change     : 11/25/2019 13:13:04
=====

=====
Associations
=====
```

```
No interfaces found using diameter-auth-policy "diam-nasreq-1".

-----
Local User Database IPOE Hosts using diameter-auth-policy "diam-nasreq-1"
-----
Local User Database                IPOE Host
-----
ludb-1                            sub-02-01_diam
-----
No. of Local User Database IPOE Hosts: 1
-----

No Local User Database PPP Hosts found using diameter-auth-policy "diam-nasreq-1".

No associated SAP's found.

=====
```

The following show command provides Diameter application message statistics for the NASREQ application policy, that includes transmit and receive counters per message type, transmit failures, retransmissions, and statistics per error category:

```
# show subscriber-mgmt diameter-application-policy "diam-nasreq-1" statistics

=====
Diameter node statistics for policy "diam-nasreq-1"
=====
Message                                Requests            Answers
-----
Authorization-Authentication           85                  56
-----

Request message transmission failure* 29
Request message retransmissions        0

Result code                            Sent                Received
-----
(1xxx) Informational                   0                   0
(2xxx) Success                         0                   53
(3xxx) Protocol Errors                 0                   0
(4xxx) Transient Failures              0                   3
(5xxx) Permanent Failures              0                   0
=====
* indicates that the corresponding row element may have been truncated.
```

Diameter debugging is split between node and application level:

```
debug
  diameter
    application
      policy "diam-nasreq-1"
      session-messages
    exit
  exit
  node "bng-nasreq.nokia.com"
  peer "aaa-1.nokia.com"
  exit
  peer "aaa-2.nokia.com"
  exit
exit
```

```
exit
```

In this chapter, the Diameter application level debugging for application messages is explained as well as application message routing errors that are reported in the node level debug. When a Python script is active for the node, the debug messages are logged after Python processing.

The **session-messages** option configured in the diameter application policy debug enables debug output for all Diameter application messages for the specified application policy. For a NASREQ application policy, this includes AA-Request and AA-Answer messages. By default, application error conditions are also logged in the debug output. Debug for application error conditions can be disabled with the debug option **no on-error** at the application or at the policy debug context.

To enable debug for the Python script that manipulates the NASREQ messages, use this debug configuration:

```
debug
  python
    python-script "nasreq-1"
    script-all-info
  exit
exit
exit
```

Following is a debug output example of a new IPoE subscriber setup: a DHCP Discover is received on a capture SAP, followed by the Diameter NASREQ AAR and AAA authentication messages. The Python script debug output shows that AVP 443 (Subscription ID) is added which is also visible in the AAR message debug (logged after the Python script is executed):

```
234756 2019/11/25 16:55:07.994 UTC minor: DEBUG #2001 Base SVCMMGR
SVCMMGR: RX DHCP Packet
  VPLS 10, SAP 1/1/2:*. *

  BootRequest to UDP port 67
  ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
  siaddr: 0.0.0.0          giaddr: 0.0.0.0
  chaddr: 00:02:01:00:00:01  xid: 0xf66ff11b

  DHCP options:
  [82] Relay agent information: len = 21
    [1] Circuit-id: CircuitID
    [2] Remote-id: RemoteID
  [53] Message type: Discover
  [255] End

234757 2019/11/25 16:55:07.994 UTC minor: DEBUG #2001 Base Python Output
Python Output: nasreq-1

234758 2019/11/25 16:55:07.994 UTC minor: DEBUG #2001 Base Python Result
Python Result: nasreq-1
Diameter AVP code 443, SET
  "('@', '\\x00\\x00\\x01\\xc2@\\x00\\x00\\x0c\\x00\\x00\\x00\\x04\\x00\\x00\\x01\\xbc@\\x00\\x00\\x1900:02:01:00:00:01\\x00\\x00\\x00')"
```

```
234759 2019/11/25 16:55:07.995 UTC minor: DEBUG #2001 Base DIAMETER
DIAMETER: Message Transmission
Transmit: "AAR"
Application policy: "diam-nasreq-1"
Node: "bng-nasreq.nokia.com"
```



```
Received peer: N/A
Transmit peer: "aaa-1.nokia.com"
Python policy: "py-nasreq-1"
Header
  ver 1 len 336 flags RP----- code 265
  app-id 1 hbh-id 12691 e2e-id 1731030625
AVPs
  session-Id (263) -M----- [42]
    data [34] (UTF8String) : bng-nasreq.nokia.com;1572949245;91
  auth-appl-id (258) -M----- [12]
    data [4] (Unsigned32) : 1 : Nasreq
  origin-host (264) -M----- [28]
    data [20] (DiameterIdentity) : bng-nasreq.nokia.com
  origin-realm (296) -M----- [17]
    data [9] (DiameterIdentity) : nokia.com
  destination-realm (283) -M----- [17]
    data [9] (DiameterIdentity) : nokia.com
  auth-request-type (274) -M----- [12]
    data [4] (Enumerated) : 3 : AUTHORIZE_AUTHENTICATE
  nas-port-id (87) -M----- [21]
    data [13] (UTF8String) : 1/1/2:2513.20
  origin-state-id (278) -M----- [12]
    data [4] (Unsigned32) : 1572949245
  user-name (1) -M----- [25]
    data [17] (UTF8String) : 00:02:01:00:00:01
  user-password (2) -M----- [11]
    data [3] (OctetString) : 0x50 57 44
  auth-session-state (277) -M----- [12]
    data [4] (Enumerated) : 1 : NO_STATE_MAINTAINED
  agent-circuit-id (1) VM----- [21]
    vendor-id DSL_FORUM
    data [9] (OctetString) : 0x43 69 72 63 75 69 74 49 44
  agent-remote-id (2) VM----- [20]
    vendor-id DSL_FORUM
    data [8] (OctetString) : 0x52 65 6d 6f 74 65 49 44
  subscription-id (443) -M----- [48]
    data [40] (Grouped)
      subscription-id-type (450) -M----- [12]
        data [4] (Enumerated) : 4 : private
      subscription-id-data (444) -M----- [25]
        data [17] (UTF8String) : 00:02:01:00:00:01

234760 2019/11/25 16:55:08.008 UTC minor: DEBUG #2001 Base DIAMETER
DIAMETER: Message Reception
Receive: "AAA"
Application policy: "diam-nasreq-1"
Node: "bng-nasreq.nokia.com"
Received peer: "aaa-1.nokia.com"
Transmit peer: N/A
Python policy: "py-nasreq-1"
Header
  ver 1 len 372 flags -P----- code 265
  app-id 1 hbh-id 12691 e2e-id 1731030625
AVPs
  session-Id (263) -M----- [42]
    data [34] (UTF8String) : bng-nasreq.nokia.com;1572949245;91
  auth-appl-id (258) -M----- [12]
    data [4] (Unsigned32) : 1 : Nasreq
  origin-host (264) -M----- [23]
    data [15] (DiameterIdentity) : aaa-1.nokia.com
  origin-realm (296) -M----- [17]
    data [9] (DiameterIdentity) : nokia.com
  result-code (268) -M----- [12]
```

```

data [4] (Unsigned32) : 2001 : DIAM_RESCODE_SUCCESS
auth-session-state (277) -M----- [12]
data [4] (Enumerated) : 1 : NO_STATE_MAINTAINED
alc-subscriber-id-str (11) V----- [21]
  vendor-id NOKIA
data [9] (UTF8String) : sub-02-01
alc-subscriber-profile-str (12) V----- [25]
  vendor-id NOKIA
data [13] (UTF8String) : sub-profile-1
alc-sla-profile-str (13) V----- [25]
  vendor-id NOKIA
data [13] (UTF8String) : sla-profile-1
alc-msap-service-id (31) V----- [16]
  vendor-id NOKIA
data [4] (Unsigned32) : 1000
alc-msap-interface (33) V----- [25]
  vendor-id NOKIA
data [13] (UTF8String) : group-int-2-1
alc-msap-policy (32) V----- [25]
  vendor-id NOKIA
data [13] (UTF8String) : msap-policy-1
framed-pool (88) ----- [23]
data [15] (UTF8String) : pool-pe2-1000-2
framed-ipv6-pool (100) ----- [23]
data [15] (OctetString) : pool-pe2-1000-2
alc-delegated-ipv6-pool (131) V----- [27]
  vendor-id NOKIA
data [15] (OctetString) : pool-pe2-1000-2

```

Let's consider a failure scenario, for example a configuration error in the diameter application policy where the destination realm for the NASREQ server is incorrect: the *realm realm-1.com* cannot be routed in the node *bng-nasreq.nokia.com*:

```

configure
  subscriber-mgmt
    diameter-application-policy "diam-nasreq-1"
      description "Diameter application policy - NASREQ"
      application nasreq
      diameter-node "bng-nasreq.nokia.com" destination-realm "realm-1.com"
      ---snip---

# show aaa diameter-node "bng-nasreq.nokia.com" routing-table

=====
Routes
=====
Realm-Name
  Application  Pref. Id  Server-Identifier
-----
nokia.com
  nasreq      10    1    aaa-1.nokia.com
nokia.com
  nasreq      20    2    aaa-2.nokia.com
-----
No. of routes: 2
=====

```

In this case, the diameter node level debug displays the routing failure:

```

244903 2019/11/27 14:26:48.705 UTC minor: DEBUG #2001 Base SVCMMGR
SVCMMGR: RX DHCP Packet
VPLS 10, SAP 1/1/2:*.

```

```
BootRequest to UDP port 67
ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
siaddr: 0.0.0.0          giaddr: 0.0.0.0
chaddr: 00:02:01:00:00:01  xid: 0xf66ff11b

DHCP options:
[82] Relay agent information: len = 21
    [1] Circuit-id: CircuitID
    [2] Remote-id: RemoteID
[53] Message type: Discover
[255] End

244904 2019/11/27 14:26:48.705 UTC minor: DEBUG #2001 Base DIAMETER
DIAMETER: Failure
Transmit: "AAR"
Application policy: "diam-nasreq-1"
Node: "bng-nasreq.nokia.com"
Received peer: N/A
Transmit peer: N/A
Python policy: "py-nasreq-1"
Result code: N/A
Error message: "no route to destination"
Failed AVP: N/A

Message:
Header
  ver 1 len 288 flags RP----- code 265
  app-id 1 hbh-id 2384494 e2e-id 1731030625
AVPs
  session-Id (263) -M----- [43]
    data [35] (UTF8String) : bng-nasreq.nokia.com;1572949245;121
  auth-appl-id (258) -M----- [12]
    data [4] (Unsigned32) : 1 : Nasreq
  origin-host (264) -M----- [28]
    data [20] (DiameterIdentity) : bng-nasreq.nokia.com
  origin-realm (296) -M----- [17]
    data [9] (DiameterIdentity) : nokia.com
  destination-realm (283) -M----- [19]
    data [11] (DiameterIdentity) : realm-1.com
  auth-request-type (274) -M----- [12]
    data [4] (Enumerated) : 3 : AUTHORIZE_AUTHENTICATE
  nas-port-id (87) -M----- [21]
    data [13] (UTF8String) : 1/1/2:2513.20
  origin-state-id (278) -M----- [12]
    data [4] (Unsigned32) : 1572949245
  user-name (1) -M----- [25]
    data [17] (UTF8String) : 00:02:01:00:00:01
  user-password (2) -M----- [11]
    data [3] (OctetString) : 0x50 57 44
  auth-session-state (277) -M----- [12]
    data [4] (Enumerated) : 1 : NO_STATE_MAINTAINED
  agent-circuit-id (1) VM----- [21]
    vendor-id DSL_FORUM
    data [9] (OctetString) : 0x43 69 72 63 75 69 74 49 44
  agent-remote-id (2) VM----- [20]
    vendor-id DSL_FORUM
    data [8] (OctetString) : 0x52 65 6d 6f 74 65 49 44
```

Conclusion

In this chapter the configuration and troubleshooting of the Diameter NASREQ application is described. The NASREQ application provides subscriber authentication and authorization.

Diameter Base Protocol: Establishing a Diameter Peer Connection

This chapter provides information about configuring and troubleshooting the Diameter Base protocol to establish a Diameter peer connection.

Topics in this chapter include:

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

Applicability

This information and configuration in this chapter are based on SR OS Release 19.10.R1.



Note:

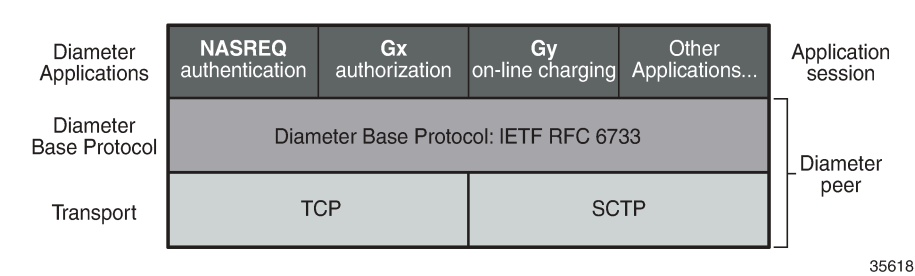
This chapter covers the Diameter base protocol implementation that is available from SR OS Release 16.0.R4 onward (configured in the **aaa** CLI context as **diameter node**). The legacy Diameter base implementation (configured in the **aaa** CLI context as **diameter-peer-policy**) is supported in maintenance mode only, without any further feature enhancement planned. Nokia recommends using or transitioning to the new Diameter base protocol implementation.

Overview

Diameter is an Authentication, Authorization and Accounting (AAA) protocol defined by the IETF in RFC 6733, *Diameter Base Protocol*. While historically wireline access networks were largely based on RADIUS for subscriber authentication, authorization, and accounting, it was decided by 3rd Generation Partnership Project (3GPP) that wireless access networks will be largely based on Diameter. Over time, operators are looking to converge both types of networks, and one of the aspects of this is to replace RADIUS in wireline access networks by Diameter.

Diameter is based on three layers: the transport layer, the Diameter base protocol layer and the Diameter applications as shown in [Figure 30: Diameter protocol stack](#).

Figure 30: Diameter protocol stack




The bottom layer is the transport layer and can be either TCP or SCTP. SR OS supports TCP. The Diameter base protocol implementation in SR OS is based on RFC 6733. The top layer contains the Diameter applications. SR OS supports NASREQ for authentication, Gx for authorization, policy management and usage monitoring and Gy or Diameter Credit Control Application (DCCA) for online charging.

Figure 31: Diameter network topology shows a Diameter network topology that will be used in the configuration examples in this chapter.

A Diameter Client (BNG) is connected via peer-1a and peer-1b to two Diameter Agents (DRA-1 and DRA-2) that provide connectivity to the Diameter Application Servers (PCRF). Via these peers, the BNG can authenticate and perform policy control of subscriber sessions using the NASREQ and Gx applications. The same Diameter Client (BNG) is also directly connected to another Diameter Application Server (OCS) via peer-2. Via this peer, on-line charging can be done for the subscriber sessions using the Gy application.

Figure 31: Diameter network topology



Configuration

The Diameter base protocol and the Diameter applications are configured separately, where the Diameter base protocol must be configured first, and the Diameter applications next. The transport layer configuration is part of the Diameter base protocol layer. This example describes the Diameter base protocol configuration.

The Diameter base protocol and the corresponding transport layer configuration is based on Diameter Nodes. Each Diameter Node represents a Diameter routing instance and contains a list of peers in the routing domain that provide direct or indirect connectivity to application servers.

An example Diameter node configuration that corresponds with the topology in [Figure 31: Diameter network topology](#) is shown below.

```
configure
aaa
  diameter
    node "bng-gx.realm-1.com" create
      description "Authentication and Policy Management"
      ipv6-source-address 2001:db8::2
      source-address 192.0.2.2
      peer index 1 "dra-1.realm-2.com" create
        address 2001:db8:2:6::1
        preference 10
        no shutdown
      exit
      peer index 2 "dra-2.realm-2.com" create
        address 172.16.7.2
        preference 20
        no shutdown
      exit
    exit
  node "bng-gy.realm-1.com" origin-realm "realm-1.com" create
    description "Credit Control"
    router "management"
    peer index 1 "ocs.realm-1.com" create
      address 192.99.3.0
      no shutdown
    exit
  exit
exit
exit
```

A Diameter node configuration requires a unique origin host as key. The origin host is used in Diameter application policies to associate the application with the node. All Diameter base and application messages forwarded via the peers of that node use the configured origin host in the Origin-Host AVP. The value for the Origin-Realm AVP is by default derived from the configured origin host: the realm is the part of the origin host after the first dot (".") as delimiter or equal to the origin host when it does not contain a delimiter. For example, for **node "bng-gx.realm-1.com"**: Origin-Host = "bng-gx.realm-1.com", Origin-Realm = "realm-1.com". It is also possible to explicitly configure an origin realm as shown in the example for the **node "bng-gy.realm-1.com"**.

A node configuration can include a routing context, and an IPv4 and/or IPv6 source address. These parameters are used to establish the TCP transport connection for all peers in the node. The specified source address must be a reachable local interface address in the specified or in the default routing instance. For **node "bng-gx.realm-1.com"** in the example, no routing instance is specified. By default, the TCP connections are established in the Base router using the specified source addresses. For **node "bng-**

gy.realm-1.com" in the example, the out of band routing instance **router "management"** is used to establish the TCP connection of its peer. As no source address is specified, the system will automatically select an interface address, in this case an out of band IP address configured in the Boot Options File (BOF).

Within a Diameter node, up to 5 peers can be configured with an index value between 1 and 5 as key, an IPv4 or IPv6 destination address for the TCP connection, and a mandatory destination host that is used as Destination-Host AVP value for all Diameter base messages on the peer. In a Diameter node, one peer is selected to forward application messages for a specific application session. The other peers provide redundancy when supported by the Diameter application, such as Gy session failover. A Diameter peer for application messages is selected based on following criteria:

1. Forwarding:

If the application message contains a Destination-Host AVP, select the peer in the peer table with a matching configured destination host. This is the forwarding phase.

2. Routing:

When the lookup in the peer table fails, perform a lookup in the realm routing table and select the peer with realm name equal to the Destination-Realm AVP in the application message, and with matching application ID. When multiple peers are matched, select in order of priority until a single peer is found:

- a. the peer with the lowest configured preference (default preference is 50)
- b. the peer with the lowest index

3. Default peer:

When both forwarding lookup in the peer table and routing lookup in the realm routing table were unsuccessful, use the peer configured as **default-peer**. Only a single peer in a node can be configured as a default-peer:

```
>config>aaa>diam>node>peer# default-peer
MINOR: DIAM #2001 Multiple default peer is not allowed
```

For **node "bng-gx.realm-1.com"** in the example, peer-1a with index 1 has a configured preference of 10 and peer-2 with index 2 has a configured preference of 20. Diameter Gx application messages will fail the peer table lookup as the destination host of the PCRF will not be present (no direct connection between Diameter client and Diameter application server):

```
# show aaa diameter-node "bng-gx.realm-1.com" peers

=====
Peers
=====
Host identity          Status      Default Preference Active
-----
dra-1.realm-2.com      I-Open     No          10         Yes
dra-2.realm-2.com      I-Open     No          20         Yes
-----
No. of peers: 2
=====
```

Instead a realm routing table lookup is performed to find the peer for forwarding the application messages. In this case peer-1a (dra-1.realm-2.com) is selected based on the matching destination realm (realm-2.com), application ID (Gx) and the lower preference value:

```
# show aaa diameter-node "bng-gx.realm-1.com" routing-table

=====
```


Routes				
=====				
Realm-Name	Application	Pref.	Id	Server-Identifier

realm-2.com	nasreq gx	10	1	dra-1.realm-2.com
realm-2.com	nasreq gx	20	2	dra-2.realm-2.com

No. of routes: 2				
=====				

The realm routing table is populated based on the Origin-Realm AVP and Application-Id AVP received in the Capability Exchange Answer message together with the configured index and preference values.

Note that Diameter answer messages do not rely on peer or realm routing table lookups. Answers are forwarded over the same route in the reverse direction of the matching requests. This is achieved with a transactional cache in each traversed Diameter node, using the Hop-by-Hop AVP to match requests with answers.

When enabling the peer (no shutdown), the system tries to establish the transport TCP connection. Once the TCP session is up, the system starts a Diameter Capability Exchange using the configured Diameter identity (Origin-Host and Origin-Realm AVPs) and advertising support for all SR OS Diameter applications in Application-Id AVP's (NASREQ, Gx, and Gy). When the Origin-Host AVP in the received CEA message corresponds with the destination host configured for the peer (case insensitive) and at least one application in the CEA is common with the SR OS advertised applications, then the peer moves to the I-Open state (I from Initiator). An example of a Capability Exchange is illustrated in detail in the troubleshooting section.

Optionally, a connection and a watchdog timer can be configured in the Diameter node:

```
configure
aaa
    diameter
        node "bng-gx.realm-1.com" create
        connection-timer 30
        ---snip---
        peer index 1 "dra-1.realm-2.com" create
        connection-timer 30
        watchdog-timer 30
        ---snip---
```

- **connection-timer**

The connection timer or Tc timer controls the frequency at which a transport connection is attempted to be established. The default value is 30 seconds. This timer can be configured per node to be used by all peers or overridden per peer.

- **watchdog-timer**

The watchdog timer controls the frequency at which Device-Watchdog-Request messages are transmitted to the peer, and is called the Tw timer in RFC 3539, *Authentication, Authorization and Accounting (AAA) Transport Profile*. A small timer results in a faster detection of a peer failure at the expense of generating more messages. The timer is configured per peer and its default value is 30 seconds.

A Python policy can be configured in the Diameter node to manipulate Diameter messages transmitted to and/or received on its peers.

```
configure
```

```
aaa
    diameter
        node "bng-gx.realm-1.com" origin-realm "realm-1.com" create
        python-policy "py-diameter-1"
        ---snip---
    exit
exit
exit
```

Manipulating Diameter messages, such as changing the content or format of AVPs using Python is out of the scope of this chapter.

By default, Diameter messages are sent with a DSCP set to AF41. The DSCP value can be changed with the sgt-qos configuration:

```
# configure router sgt-qos application diameter dscp "nc1"
```

SR OS uses TCP as transport and the TCP destination port number is fixed to the standard Diameter base protocol port 3868. The source port is randomly chosen from the ephemeral port range.

Troubleshooting

The status and statistics of the Diameter peers can be verified with following show commands:

```
# show aaa diameter-node "bng-gx.realm-1.com" peers
```

```
=====
Peers
=====
Host identity          Status      Default Preference Active
-----
dra-1.realm-2.com      I-Open     No         10         Yes
dra-2.realm-2.com      I-Open     No         20         Yes
-----
No. of peers: 2
=====
```

```
# show aaa diameter-node "bng-gx.realm-1.com" peer "dra-1.realm-2.com"
```

```
=====
Peer "dra-1.realm-2.com"
=====
Index                  : 1
Status                 : I-Open
Administrative state   : enabled
Active                 : Yes
Active applications    : nasreq gx
Last disconnect cause  : rebooting
Preference             : 10
Default peer           : No
Connection timer (s)   : N/A
Watchdog timer (s)     : 13
Pending messages       : 0
Remote realm           : realm-2.com
Remote IP address       : 2001:db8:2:6::1
Remote TCP port         : 3868
Remote Origin-State-Id : 1574235027
Local host identity     : bng-gx.realm-1.com
Local realm            : realm-1.com
```

```
Local IP address      : 2001:db8::2
Local TCP port       : 53734
Last management change : 11/19/2019 15:05:48
=====
```

```
# show aaa diameter-node "bng-gx.realm-1.com" peer "dra-1.realm-2.com" statistics
```

```
=====
Peer "dra-1.realm-2.com"
=====
```

Message	Sent	Received
Capabilities-Exchange-Request	7	0
Capabilities-Exchange-Answer	0	7
Disconnect-Peer-Request	1	4
Disconnect-Peer-Answer	4	1
Device-Watchdog-Request	1217	778
Device-Watchdog-Answer	778	1217
Application message request	0	0
Application message answer	0	0

```
Last cleared time: N/A
=====
```

To clear the peer statistics, use following command:

```
# clear aaa diameter-node "bng-gx.realm-1.com" peer "dra-1.realm-2.com" statistics
```

Diameter debugging is split between node and application level:

```
debug
  diameter
    application
      policy "diam-nasreq-1"
      session-messages
    exit
  exit
  node "bng-gx.realm-1.com"
  peer "dra-1.realm-2.com"
  peer-to-peer
  exit
exit
exit
exit
```

In this chapter, the Diameter base protocol debugging for peer messages is explained, configured at the node level debug. When a Python script is active for the node, the debug messages are logged after Python processing.

To debug the Diameter base protocol messages for **peer "dra-1.realm-2.com"**, use the following debug commands:

```
debug
  diameter
    node "bng-gx.realm-1.com"
    peer "dra-1.realm-2.com"
    peer-to-peer
  exit
exit
exit
```

```
exit
```

The **peer-to-peer** option enables debug output for all Diameter base messages of the specified peer: Capabilities Exchange, Device Watchdog and Disconnect Peer messages. By default, error conditions are also logged in the debug output. Debug for error conditions can be disabled per Diameter node or per peer with the debug option **no on-error**. Errors reported at the node level include Diameter base errors that are unrelated to a peer, such as a routing problem for a Diameter application message. Errors reported at the peer level include all errors that occur after peer selection and peer connection errors.

Let's start with the peer connection in Closed state (remote end rebooting):

```
*A:pe2# show aaa diameter-node "bng-gx.realm-1.com" peer "dra-1.realm-2.com"

=====
Peer "dra-1.realm-2.com"
=====
Index                : 1
Status               : Closed
Administrative state  : enabled
Active               : No
Active applications   :
Last disconnect cause : rebooting
Preference           : 10
Default peer         : No
Connection timer (s) : 18
Watchdog timer (s)   : N/A
Pending messages     : 0
Remote realm         : (Not Specified)
Remote IP address    : (Not Specified)
Remote TCP port      : (Not Specified)
Remote Origin-State-Id : (Not Specified)
Local host identity   : bng-gx.realm-1.com
Local realm          : realm-1.com
Local IP address     : (Not Specified)
Local TCP port       : (Not Specified)
Last management change : 11/19/2019 15:05:48
=====
```

The **Connection timer (s)** field in above peer details output show that in 18 seconds, a new connection attempt will be made, followed by a Capabilities Exchange when successful. The transmitted Capabilities-Exchange-Request (CER) and received Capabilities-Exchange-Answer (CEA) are shown in the debug output:

```
233997 2019/11/20 19:17:16.271 UTC minor: DEBUG #2001 Base DIAMETER
DIAMETER: Message Transmission
Transmit: "CER"
Application policy: N/A
Node: "bng-gx.realm-1.com"
Received peer: N/A
Transmit peer: "dra-1.realm-2.com"
Python policy: N/A
Header
  ver 1 len 284 flags R----- code 257
  app-id 0 hbh-id 19864 e2e-id 3486524428
AVPs
  origin-host (264) -M----- [26]
    data [18] (DiameterIdentity) : bng-gx.realm-1.com
  origin-realm (296) -M----- [19]
    data [11] (DiameterIdentity) : realm-1.com
  host-ip-addr (257) -M----- [26]
    data [18] (Address) : ipv6 2001:db8::2
```

```
vendor-id (266) -M----- [12]
  data [4] (Unsigned32) : 6527
product-name (269) ----- [13]
  data [5] (UTF8String) : SR-05
auth-appl-id (258) -M----- [12]
  data [4] (Unsigned32) : 1 : Nasreq
auth-appl-id (258) -M----- [12]
  data [4] (Unsigned32) : 4 : Gy
auth-appl-id (258) -M----- [12]
  data [4] (Unsigned32) : 16777238 : Gx
vend-specific-appl-id (260) -M----- [32]
  data [24] (Grouped)
    vendor-id (266) -M----- [12]
      data [4] (Unsigned32) : 10415
    auth-appl-id (258) -M----- [12]
      data [4] (Unsigned32) : 4 : Gy
  vend-specific-appl-id (260) -M----- [32]
    data [24] (Grouped)
      vendor-id (266) -M----- [12]
        data [4] (Unsigned32) : 10415
      auth-appl-id (258) -M----- [12]
        data [4] (Unsigned32) : 16777238 : Gx
supported-vendor-id (265) -M----- [12]
  data [4] (Unsigned32) : 3561
supported-vendor-id (265) -M----- [12]
  data [4] (Unsigned32) : 6527
supported-vendor-id (265) -M----- [12]
  data [4] (Unsigned32) : 10415
supported-vendor-id (265) -M----- [12]
  data [4] (Unsigned32) : 13019
firmware-revision (267) ----- [12]
  data [4] (Unsigned32) : 191001

233998 2019/11/20 19:17:16.275 UTC minor: DEBUG #2001 Base DIAMETER
DIAMETER: Message Reception
Receive: "CEA"
Application policy: N/A
Node: "bng-gx.realm-1.com"
Received peer: "dra-1.realm-2.com"
Transmit peer: N/A
Python policy: N/A
Header
  ver 1 len 240 flags ----- code 257
  app-id 0 hbh-id 19864 e2e-id 3486524428
AVPs
  result-code (268) -M----- [12]
    data [4] (Unsigned32) : 2001 : DIAM_RESCODE_SUCCESS
  origin-host (264) -M----- [25]
    data [17] (DiameterIdentity) : dra-1.realm-2.com
  origin-realm (296) -M----- [19]
    data [11] (DiameterIdentity) : realm-2.com
  host-ip-addr (257) -M----- [26]
    data [18] (Address) : ipv6 2001:db8:2:6::1
  vendor-id (266) -M----- [12]
    data [4] (Unsigned32) : 6527
  product-name (269) ----- [28]
    data [20] (UTF8String) : PythonDiameterAgent1
  origin-state-id (278) -M----- [12]
    data [4] (Unsigned32) : 1574277432
  supported-vendor-id (265) -M----- [12]
    data [4] (Unsigned32) : 10415
  auth-appl-id (258) -M----- [12]
    data [4] (Unsigned32) : 1 : Nasreq
```

```
auth-appl-id (258) -M----- [12]
  data [4] (Unsigned32) : 16777238 : Gx
vend-specific-appl-id (260) -M----- [32]
  data [24] (Grouped)
    vendor-id (266) -M----- [12]
      data [4] (Unsigned32) : 10415
    auth-appl-id (258) -M----- [12]
      data [4] (Unsigned32) : 16777238 : Gx
  firmware-revision (267) ----- [12]
    data [4] (Unsigned32) : 1
```

The result of the successful Capabilities Exchange is that the peer moved to the I-Open state, ready to forward NASREQ and Gx application messages:

```
# show aaa diameter-node "bng-gx.realm-1.com" peer "dra-1.realm-2.com"

=====
Peer "dra-1.realm-2.com"
=====
Index                : 1
Status               : I-Open
Administrative state  : enabled
Active               : Yes
Active applications   : nasreq gx
Last disconnect cause : rebooting
Preference           : 10
Default peer         : No
Connection timer (s) : N/A
Watchdog timer (s)   : 9
Pending messages     : 0
Remote realm         : realm-2.com
Remote IP address     : 2001:db8:2:6::1
Remote TCP port       : 3868
Remote Origin-State-Id : 1574277432
Local host identity   : bng-gx.realm-1.com
Local realm          : realm-1.com
Local IP address      : 2001:db8::2
Local TCP port        : 55199
Last management change : 11/19/2019 15:05:48
=====
```

The **Watchdog timer (s)** field in preceding peer details output shows that in 9 seconds, a Device Watchdog exchange will be initiated. The transmitted Device-Watchdog-Request (DWR) and received Device-Watchdog-Answer (DWA) are shown in the debug output:

```
233999 2019/11/20 19:17:44.268 UTC minor: DEBUG #2001 Base DIAMETER
DIAMETER: Message Transmission
Transmit: "DWR"
Application policy: N/A
Node: "bng-gx.realm-1.com"
Received peer: N/A
Transmit peer: "dra-1.realm-2.com"
Python policy: N/A
Header
  ver 1 len 68 flags R----- code 280
  app-id 0 hbh-id 19865 e2e-id 3486524431
AVPs
  origin-host (264) -M----- [26]
    data [18] (DiameterIdentity) : bng-gx.realm-1.com
  origin-realm (296) -M----- [19]
    data [11] (DiameterIdentity) : realm-1.com
```

```
234000 2019/11/20 19:17:44.271 UTC minor: DEBUG #2001 Base DIAMETER
DIAMETER: Message Reception
Receive: "DWA"
Application policy: N/A
Node: "bng-gx.realm-1.com"
Received peer: "dra-1.realm-2.com"
Transmit peer: N/A
Python policy: N/A
Header
  ver 1 len 92 flags ----- code 280
  app-id 0 hbh-id 19865 e2e-id 3486524431
AVPs
  result-code (268) -M----- [12]
    data [4] (Unsigned32) : 2001 : DIAM_RESCODE_SUCCESS
  origin-host (264) -M----- [25]
    data [17] (DiameterIdentity) : dra-1.realm-2.com
  origin-realm (296) -M----- [19]
    data [11] (DiameterIdentity) : realm-2.com
  origin-state-id (278) -M----- [12]
    data [4] (Unsigned32) : 1574277432
```

Now let's try to bring up the peer in the **node bng-gy.realm-1.com**:

```
# show aaa diameter-node "bng-gy.realm-1.com" peers

=====
Peers
=====
Host identity                               Status      Default Preference Active
-----
ocs.realm-1.com                            Closed      No          50          No
-----
No. of peers: 1
=====
```

Debug is enabled at the peer level for error conditions without the **peer-to-peer** option. Failures are reported, but not all transmitted and received peer messages.

```
debug
  diameter
    node "bng-gy.realm-1.com"
    peer "ocs.realm-1.com"
    exit
  exit
exit
```

The Diameter server is provisioned with an origin host different from the configured destination host for the peer, resulting in a failure and peer reset:

```
234330 2019/11/22 14:57:32.272 UTC minor: DEBUG #2001 management DIAMETER
DIAMETER: Failure
Receive: "CEA"
Application policy: N/A
Node: "bng-gy.realm-1.com"
Received peer: "ocs.realm-1.com"
Transmit peer: N/A
Python policy: N/A
Result code: "DIAM_RESCODE_INVALID_AVP_VALUE"
Error message: "mismatch with locally stored information"
```

```
Failed AVP:
  origin-host (264) -M----- [27]
    data [19] (DiameterIdentity) : ocs.wrong-realm.com

Message:
  Header
    ver 1 len 176 flags ----- code 257
    app-id 0 hbh-id 6050 e2e-id 3486524894
  AVPs
    result-code (268) -M----- [12]
      data [4] (Unsigned32) : 2001 : DIAM_RESCODE_SUCCESS
    origin-host (264) -M----- [27]
      data [19] (DiameterIdentity) : ocs.wrong-realm.com
    origin-realm (296) -M----- [23]
      data [15] (DiameterIdentity) : wrong-realm.com
    host-ip-addr (257) -M----- [14]
      data [6] (Address) : ipv4 192.99.3.0
    vendor-id (266) -M----- [12]
      data [4] (Unsigned32) : 6527
    product-name (269) ----- [28]
      data [20] (UTF8String) : PythonDiameterServer
    origin-state-id (278) -M----- [12]
      data [4] (Unsigned32) : 1574434643
    auth-appl-id (258) -M----- [12]
      data [4] (Unsigned32) : 4 : Gy
    firmware-revision (267) ----- [12]
      data [4] (Unsigned32) : 1

234331 2019/11/22 14:57:32.272 UTC minor: DEBUG #2001 management DIAMETER
DIAMETER: Peer Reset
Node: "bng-gy.realm-1.com"
Peer: "ocs.realm-1.com"
Reason: "failed to parse received CEA"
```

Events

Following events are defined for the Diameter base protocol:

=====					
Application					
ID#	Event Name	P	g/s	Logged	Dropped

2007	tmnxDiamMessageDropped	MI	thr	0	0
2008	tmnxDiamNdPeerStatActiveChanged	MI	thr	46	0
=====					

The *tmnxDiamNdPeerStatActiveChanged* event is generated when the state of a Diameter peer toggles between active / not active:

```
38080 2019/11/22 14:52:02.269 UTC MINOR: DIAMETER #2008 management peer state change
"DIAMETER node bng-gy.realm-1.com, peer ocs.realm-1.com is active"
```

The *tmnxDiamMessageDropped* event is generated when Diameter base drops a malformed message.

Conclusion

As a result of fixed mobile network convergence, Diameter is used in fixed access networks as an alternative for Radius based AAA. Diameter peering provides reliable and secure transport with peer redundancy. Its functionality is defined in a base Diameter protocol specified in RFC 6733. Various applications can be layered on top of base Diameter and they can utilize the robust transport capabilities that Diameter provides.

ESM Basics

This chapter provides information about Enhanced Subscriber Management.

Topics in this chapter include:

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

Applicability

The information and configuration in this chapter are based on SR OS Release 11.0.R4.

Overview

Subscriber management in general includes the following functions:

- subscriber host authentication, identification, addressing, authorization and accounting:
 - authentication – Check whether the subscriber host is allowed access via a Local User Data Base (LUDB) or via a RADIUS server.
 - identification – Fetch the data to use for the subscriber host, including the definition of the subscriber-ID.
 - addressing – Fetch the address information to use, IPv4 and/or IPv6.
 - authorization – Check what the subscriber host is allowed to do.
 - accounting – Both off-line charging (RADIUS and XML) and on-line charging (RADIUS credit control and Diameter credit control application) are supported.
- subscriber host instantiation, based on:
 - a protocol (DHCP, PPPoE/oA/oEoA or ARP) for dynamic hosts, and started through a trigger packet.
 - a static configuration for static hosts.
- subscriber QoS - Ensure per service and per application SLAs, based on:
 - the overall subscriber rate
 - subscriber profiles
 - service level agreement profiles
- subscriber security:
 - Avoid malicious access through anti-spoofing and based on access control lists (ACLs) / IP-filters.
 - DDOS mitigation.

- subscriber persistency – The subscriber state is written to flash-disk (a.k.a. persistent data) and automatically restored on node reboot.
- subscriber resiliency – The subscriber state is retained in case of a node failure in redundant environments, through Multi Chassis Synchronization (MCS).
- subscriber troubleshooting, using:
 - OAM test.
 - mirroring.
 - debugging and event logging.

Enhanced Subscriber Management (ESM) implies subscriber management functions are applied at subscriber level. Subscriber hosts are created, and all of the features listed above apply.

Basic Subscriber Management (BSM) implies subscriber management functions are applied at SAP level. Only a subset of the functions listed above apply.

This example gives an overview of the ESM data required to perform subscriber management functions, and how the ESM data is organized.

The following terminology is used extensively in Triple Play Service Delivery Architecture (TPSDA) and is key to understanding ESM:

- device
- subscriber host
- subscriber

A device is equipment located at the customer premises. Example devices are computers, smartphones, set-top boxes, etc including the Residential Gateway (RGW) providing the connection towards the Internet. The RGW is connected to the access network using an XDSL-connection, a PON-connection, and so on.

A subscriber is a collection of subscriber hosts connected to a single RGW. The subscriber is identified by its subscriber ID, a character string of 32 characters maximum which is used for administrative purposes.

Unlike devices, which are physical entities, subscribers and subscriber hosts are logical entities. These are created dynamically and resources are allocated when a device connects to the network and becomes active.

The following host types are recognized by the Broadband Network Gateway (BNG):

- DHCP hosts
- PPPoE hosts
- ARP hosts

The BNG uses the combination of following parameters to uniquely identify a single subscriber host:

- SAP ID
- IP address
- MAC address
- session ID (PPPoE only)

Multiple subscriber hosts can be associated with a single device.

Figure 32: Bridged RGW Scenario shows a bridged RGW scenario. Subscriber-1 has two devices, device-1 and device-2. Device-1 is a dual stack PPPoE device and is assigned an IPv4 address for host-1 and an

IPv6 SLAAC prefix for host-2, both running over a single PPPoE session. Device-2 is a single stack IPoE device and is assigned an IPv4 address for host-3.

Figure 32: Bridged RGW Scenario

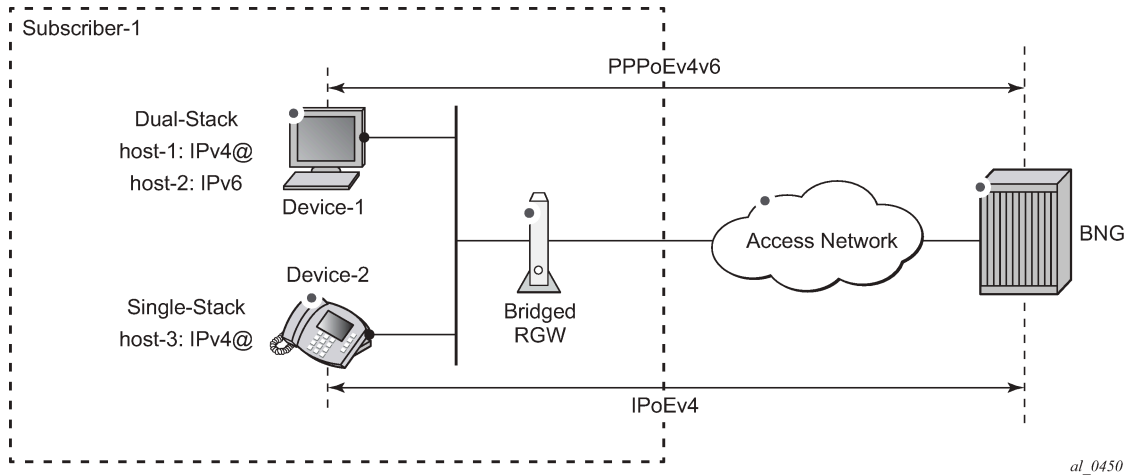
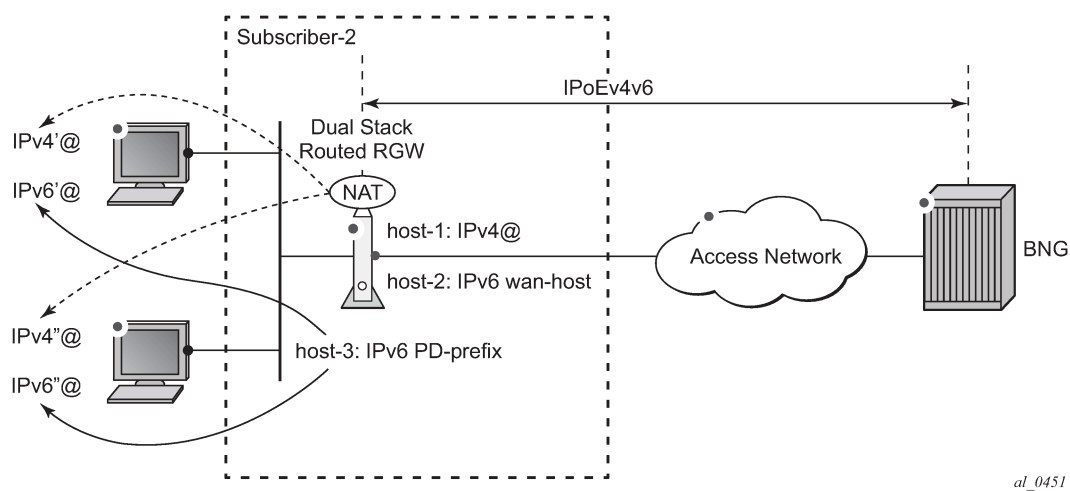


Figure 33: Routed RGW Scenario shows a routed dual stack RGW scenario. Subscriber-2 has two devices but they are hidden behind the RGW. The IP-addresses used by these devices are not known by the BNG. The RGW contains three hosts:

- Host-1 is assigned an outside IPv4 address, which the RGW uses for Network Address Translation (NAT).
- Host-2 is assigned an IPv6 wan host address or SLAAC prefix, which the RGW uses towards the outside network.
- Host-3 is assigned an IPv6 prefix for prefix delegation, which the RGW uses for allocating IPv6 addresses to IPv6 capable devices in the home network.

Figure 33: Routed RGW Scenario



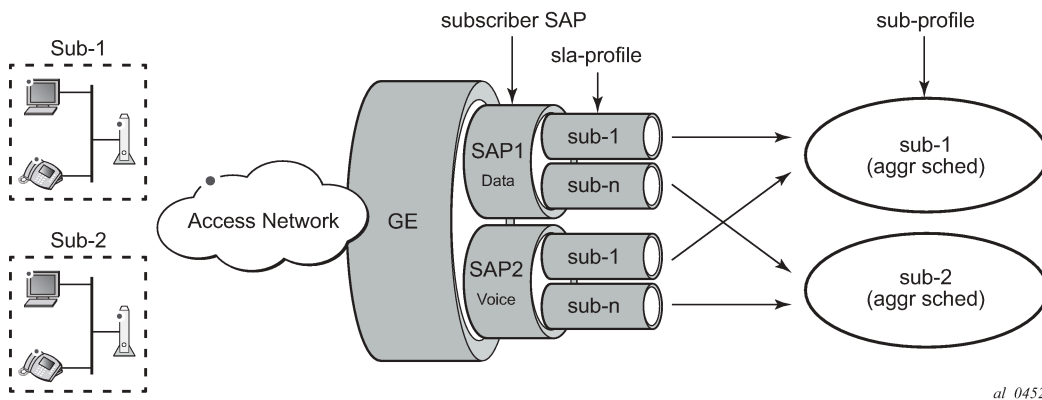
Detailed information on both scenarios can be found in [ESMv6: IPoE Dual Stack Hosts](#) and the [ESMv6: PPPoE Dual Stack Hosts](#) respectively.

Configuration

Figure 34: SLA-Profile and Sub-Profile shows two objects closely related to subscribers and subscriber hosts:

- SLA profile
- Subscriber profile

Figure 34: SLA-Profile and Sub-Profile



al_0452

SLA Profile

A Service Level Agreement profile (SLA profile) is a template identified by name (maximum 32 chars) which defines:

- per service QoS settings as part of the QoS policy (ingress and egress):
 - queue/policer
 - scheduling/priority levels
 - bandwidth limits (CIR/PIR) and queue/policer depths (MBS/CBS)
 - classification
 - (re-)marking
- IP-filters for IPv4 and/or IPv6 (ingress and egress):
 - Access Control Lists (ACL)/Filters
- host-limit
- credit-control

An example SLA profile is shown below.

```
configure
  subscriber-mgmt
    sla-profile "sla-prof-1" create
    ingress
      qos 100
    exit
```

```
        exit
    egress
        qos 100
        exit
    exit
exit
```

An instance of the sla-profile is created at host instantiation time.

The SLA profile enforces traffic control on a per service per subscriber basis. When multiple hosts for the same subscriber use the same SLA profile they share the same set of queues/policers as long as they are on the same SAP. The combined traffic for these hosts is controlled by the settings defined in the SLA profile.

Subscriber Profile

The subscriber profile (sub-profile) is a template identified by name (maximum 32 chars) which defines:

- per subscriber QoS settings:
 - aggregate rate (egress only)
 - scheduler policy (ingress and egress)
 - policer control policy (ingress and egress)
- the accounting profile (RADIUS or XML)
- multicast parameters (igmp-policy, etc.)
- Network Address Translation (NAT) parameters
- the ANCP (Access Node Control Protocol) parameters
- the default sla-profile mappings
- etc.

An example subscriber profile is shown below.

```
configure
  subscriber-mgmt
    sub-profile "sub-prof-1" create
    ingress
      policer-control-policy "pol-ctrl-1"
    exit
  exit
  egress
    scheduler-policy "down-1"
  exit
exit
exit
```

An instance of the sub-profile is created at host instantiation time.

The sub-profile enforces traffic control on a per subscriber basis. The aggregate traffic of all hosts of a particular subscriber is controlled by the settings defined in the sub-profile, as long as the subscriber hosts are terminated on the same card.

QoS details are out of the scope of this example.

Subscriber SAP

A subscriber SAP is a SAP in a VPLS (Bridged CO model, which is not covered), VPRN or IES service (both in the Routed CO model) on which the queues/policers and other resources are allocated and de-allocated on a per subscriber basis.

A static SAP is a subscriber SAP when sub-sla-mgmt is enabled and becomes operational when in the no shutdown state. Multiple subscribers can connect through this SAP simultaneously when multi-sub-sap is enabled.


```
configure
  service
    ies 1
      subscriber-interface "sub-int-1"
        group-interface "grp-int-1-1"
          sap 1/1/1:101 create
            description "data SAP for DLSAM-1"
            sub-sla-mgmt
              multi-sub-sap 400
              ---snip---
              no shutdown
          exit
        exit
      exit
```

Single SAP parameters (profiled-traffic-only and non-sub-traffic) are taken into account only when multi-sub-sap is left at its default of 1.

Subscriber Host Identification and Instantiation

The general subscriber host identification (authentication) process is depicted in [Figure 35: Subscriber Host Identification and Instantiation Process](#) and encompasses fetching the ESM data for a connecting device.

Figure 35: Subscriber Host Identification and Instantiation Process



ESM data is organized as a set of profiles which control the behavior of individual subscribers and subscriber hosts.

Static hosts are instantiated using the manually provisioned ESM data.

For dynamic hosts, the host identification and instantiation process is started by a trigger message (1 in [Figure 35: Subscriber Host Identification and Instantiation Process](#)) which can be one of the following:

- DHCPv4 DISCOVER
- DHCPv6 SOLICIT
- PPPoE PADI
- PPPoA LCP Configuration Request
- ARP REQUEST

The following types of ESM data are distinguished as mandatory or optional ESM data.

Mandatory ESM data is:

- address information (IPv4 and/or IPv6)
- subscriber ID
- sub-profile data
- sla-profile data

Optional ESM data includes:

- app-profile data
- inter-dest-id
- Access Node Control Protocol (ANCP) data

A host is instantiated only when all mandatory ESM data is available for that host.

The ESM data sources (2) are consulted in following predefined sequence:

1. LUDB
2. RADIUS
3. DHCPv4 ACK python
4. Default

The ESM data sources provide:

- the subscriber-id (3)
- the profile strings (5)

Address information (IPv4 and/or IPv6) (4) can be provided by:

- the LUDB or RADIUS ESM data sources
- a DHCP server

The LUDB, RADIUS or DHCPv4 ACK python ESM data sources provide profile strings (5) which have a maximum length of 16 characters. They must be translated into profile names (6) which have a maximum length of 32 characters. The profile name is the key to access the actual profile data. The ESM data source *default* directly returns profile names (7), which do not need any translation.

The ESM string to profile name translation is defined in configurable mapping tables which are part of the subscriber identification policy (sub-ident-policy). Mapping tables can be defined for:

- the sub-profile
- the sla-profile
- the app-profile

In case no mapping is needed because the strings and names are set to the same set of values, then a subscriber identification policy is needed with the attribute **use-direct-map-as-default**. (See section [Subscriber Identification Policy](#).)

Note the subscriber-id does not need to be translated.

The instantiation process (8) ensures the subscriber host is created:

- The subscriber host is added to the active subscriber list.
- Resources (queues, policers, filters, etc) are allocated for SLA enforcement.
- Status information is updated.

As a result the system starts forwarding user data to and from that subscriber host.

A subscriber is instantiated as soon as the first subscriber host for this subscriber is instantiated, and deleted when the last subscriber host for this subscriber is deleted.

Subscriber Identification Policy

The subscriber identification policy is identified by name and defines:

- a description (optional)
- the sub-profile map
- the sla-profile map

- the app-profile map (optional)
- the location of a DHCPv4 ACK python script (optional):
 - primary, secondary and tertiary locations are possible
- the DHCP option used to get the identification strings from (optional)

A subscriber identification policy is needed for dynamic hosts only. In that case one of the ESM data sources from [Figure 35: Subscriber Host Identification and Instantiation Process](#) is used.

Using a python script for subscriber host identification and instantiation is restricted to IPv4 hosts when triggered by the DHCPv4 ACK message.

The first example has no explicit mappings.

```
configure
  subscriber-mgmt
    sub-ident-policy "sub-id-pol-1" create
    description "direct mapping policy"
    sub-profile-map
      use-direct-map-as-default
    exit
    sla-profile-map
      use-direct-map-as-default
    exit
    app-profile-map
      use-direct-map-as-default
    exit
  exit
```

The second example contains explicit mappings.

```
configure
  subscriber-mgmt
    sub-ident-policy "sub-id-pol-2" create
    description "explicit mapping policy"
    sub-profile-map
      entry key "sub-string-1" sub-profile "sub-prof-1"
      entry key "sub-string-2" sub-profile "sub-prof-2"
      entry key "sub-string-3" sub-profile "sub-prof-3"
    exit
    sla-profile-map
      entry key "sla-string-1" sla-profile "sla-prof-1"
      entry key "sla-string-2" sla-profile "sla-prof-2"
      entry key "sla-string-21" sla-profile "sla-prof-20"
      entry key "sla-string-22" sla-profile "sla-prof-20"
    exit
  exit
```

Note that multiple strings can be mapped to the same profile, as the example above shows.

The subscriber identification policy is applied at SAP level in the sub-sla-mgmt context, as shown below.

```
configure
  service
    ies 1
      subscriber-interface "sub-int-1" create
      ---snip---
      group-interface "grp-int-1-1" create
      ---snip---
      sap 1/1/1:111 create
      sub-sla-mgmt
```

```
def-sub-profile "sub-prof-1"
def-sla-profile "sla-prof-1"
sub-ident-policy "sub-id-pol-1"
multi-sub-sap 400
no shutdown
exit
exit
exit
exit
```

Combining Multiple ESM Data Sources

Multiple ESM data sources from [Figure 35: Subscriber Host Identification and Instantiation Process](#) can be consulted for instantiating a subscriber host. When multiple ESM data sources are used, they are accessed in following sequence:

1. LUDB
2. RADIUS
3. DHCPv4 ACK python
4. Default

The outputs from the different data sources are merged. Data is appended as explained in the Flexible Authentication Model in ESM section of this guide.

Default ESM Data Profiles

Default ESM data profiles are used when the other ESM data sources (LUDB, RADIUS, DHCPv4 ACK python) only provide partial data, or no data at all. The default data complements the data provided by the data sources in order to get the full set of ESM data needed for host instantiation.

Default ESM data profiles are defined per SAP in the sub-sla-mgmt context, as shown in the following example.

```
configure
service
  ies 1
    subscriber-interface "sub-int-1-1"
    group-interface "grp-int-1-1"
    sap 1/1/1:111 create
    sub-sla-mgmt
      def-sub-profile "sub-prof-1"
      def-sla-profile "sla-prof-1"
      sub-ident-policy "sub-id-pol-1"
      multi-sub-sap 400
      no shutdown
    exit
```

Address Information

Address information for both IPv4 as well as IPv6 is provided to subscriber hosts using a DHCPv4/v6 server, an LUDB or a RADIUS server.

The IPv4 and IPv6 address information encompasses:

- IPv4 and/or IPv6 address/mask and prefix
- default gateway (in IPv6 announced through RA messages)
- DNS server(s)
- etc

The following cases can be distinguished to obtain an IPv4 or IPv6 address/prefix upon LUDB/RADIUS authentication:

- LUDB/RADIUS returns a unique IPv4/IPv6 address/prefix, directly.
- LUDB/RADIUS returns a pool name, which then is resolved to a unique IP address/prefix through:
 - a DHCP server
 - or Local Address Assignment (LAA)
- LUDB/RADIUS does not return a unique IPv4/IPv6 address/prefix, nor a pool-name. In that case a DHCP server is contacted using a gi-address or link-address to obtain an address/prefix.


Configuring a DHCP server is out of the scope of this example.

Direct Address Assignment using LUDB/RADIUS

In the example shown in [Figure 36: Direct Address Assignment using LUDB/RADIUS](#), a unique IP address is retrieved from LUDB/RADIUS. No interaction with a DHCP server is needed at all.

This scenario applies to IPoE (DHCPv4, DHCPv6 and SLAAC) as well as to PPPoE (IPCP, DHCPv6, and SLAAC).

Figure 36: Direct Address Assignment using LUDB/RADIUS



DHCP clients require a DHCP server address in the DHCP messages, and as such a **server-id** (IPv6) and/or an **emulated-server** (IPv4) must be configured, see the example below. IPCP and SLAAC do not require a DHCP server so these commands can be omitted.

```
configure
  service
    ies 1
      subscriber-interface "sub-int-1" create
      ---snip---
      group-interface "grp-int-1-1" create
      ipv6
```

```

dhcp6
  ---snip---
  proxy-server
    server-id duid-ll
    client-applications dhcp ppp
    no shutdown
  exit
exit
arp-populate
dhcp
  proxy-server
    emulated-server 10.1.1.254
    no shutdown
  exit
  ---snip---
exit

```

The emulated-server address must be a unique address in one of the subnets allowed on this subscriber interface.


Indirect Address Assignment using a DHCP server

In the example shown in [Figure 37: Indirect Address Assignment using a DHCP Server](#), the IPv4/IPv6 address/prefix is obtained from a DHCP server, using a pool-name as returned by LUDB/RADIUS.

The DHCP server (DHCPv4/DHCPv6) can be co-located with the BNG node (internal DHCP server) or can be external to the BNG node.

This scenario applies to both DHCPv4/DHCPv6 as well as to PPP. PPP relies on an internal DHCP client and communicates with the PPP client through IPCP.

Figure 37: Indirect Address Assignment using a DHCP Server



Relaying is configured for IPv4 and for IPv6 separately, at group-interface level.

The DHCPv4 relay agent is configured in the dhcp context:

- **gi-address** – The gateway IPv4-address used by the relay-agent.
- **server** - The DHCP server IPv4-address, 10.11.11.11 in the example.
- **client-applications dhcp ppp**
 - The DHCP server accepts requests for DHCP and for PPP.
- **option** – The options added/removed to/from messages towards the server. In the example the circuit-id, the remote-id and the pool-name are added.
- **trusted** – This parameter ensures that DHCP messages with option 82 included and the gi-address set to zero are being processed instead of being dropped.

The DHCPv6 relay agent is configured in the ipv6 dhcp6 relay context:

- **server** - The DHCP server IPv6 address, 2001:DB8::11 in the example.
- **client-applications dhcp ppp**

The DHCP server accepts requests for DHCP and for PPP.

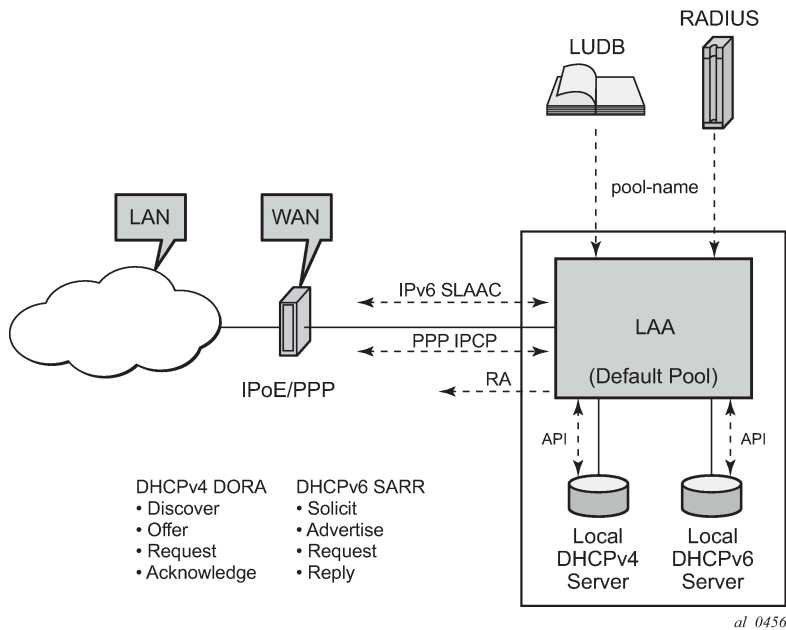
```
configure
  service
    ies 1
      subscriber-interface "sub-int-1"
        ---snip---
        group-interface "grp-int-1-1"
          ipv6
            ---snip---
            dhcp6
              relay
                server 2001:DB8::11
                client-applications dhcp ppp
                no shutdown
              exit
            exit
          exit
        dhcp
          option
            action replace
            circuit-id
            remote-id
            vendor-specific-option
            pool-name
          exit
        exit
      server 10.11.11.1
      trusted
      lease-populate 100
      client-applications dhcp ppp
      gi-address 10.1.1.254
      no shutdown
    exit
```

Indirect Address Assignment using LAA

In the example shown in [Figure 38: Indirect Address Assignment using LAA](#) LUDB/RADIUS returns a pool-name which is resolved into an IP address using Local Address Assignment (LAA). LAA is implemented

using a procedure call to an internal DHCP server, and not through the typical DHCPv4 DORA or DHCPv6 SARR sequence.

Figure 38: Indirect Address Assignment using LAA



This scenario applies to PPPv4 hosts as well as to IPv6 SLAAC hosts. LAA can also be used to provide the IA-NA address in DHCPv6 proxy scenarios where an LUDB or a RADIUS server provides an IPv6 Delegated Prefix (Delegated-IPv6-Prefix) and an IPv6 WAN Address Pool (Framed-IPv6-Pool).

LAA is configured at group-interface level and overrules the relay scenario for the IPCP DHCP client (PPP) when both are configured.

The IPv4 parameters are configured in the local-address-assignment context:

- **server** – The internal DHCPv4 server to contact, identified by name.
- **default-pool** – Defines the pool name to use in case neither the LUDB nor the RADIUS server provides a pool-name.
- **client-application** – ppp-v4.

The IPv6 parameters are configured in the local-address-assignment ipv6 context:

- **server** – the DHCPv6 server to contact, identified by name.
- **client-application** – ppp-slaac and/or ipoe-wan.

An example is show below.

```
configure
service
  ies 1
    subscriber-interface "sub-int-1" create
    ---snip---
    group-interface "grp-int-1-1" create
      local-address-assignment
        server "int-dhcp-v4"
        client-application ppp-v4
```

```
default-pool "pool4-2"
ipv6
    server "int-dhcp-v6"
    client-application ppp-slaac ipoe-wan
exit
no shutdown
exit
```

LAA details are out of the scope of this example.

ESM Data Retrieval Examples

The following ESM data source scenarios are examined below:

- static host ESM data retrieval
- dynamic host ESM data retrieval:
 - RADIUS
 - LUDB + RADIUS access

DHCPv4 ACK python details are out of the scope of this example.

ESM Data Retrieval for Static Hosts

The example below shows the creation of a static host. The IP address is mandatory, a MAC address is optional.

```
configure
service
    ies 1
        subscriber-interface "sub-int-1"
        group-interface "grp-int-1-1"
        sap 1/1/1:111
        description "sap for customer 1"
        static-host ip 10.1.1.101 create
            sla-profile "sla-prof-2"
            sub-profile "sub-prof-2"
            subscriber 03-7654321
            no shutdown
        exit
    exit
exit
```

Identification is not needed for static hosts: the subscriber-id, sub-profile and sla-profile used, are configured explicitly. Instantiation takes place at host creation time (**no shutdown**).

The following command shows the details for a single static host. The forwarding state is *Not Fwding*, meaning that no traffic can be forwarded to and from that host in this state.

```
*A:BNG# show service id 1 static-host ip-address 10.1.1.101 detail
=====
Static Hosts for service 1
=====
```

Sap	IP Address	Configured MAC	Dynamic MAC
Subscriber		Admin State	Fwding State
1/1/1:111	10.1.1.101	N/A	N/A
03-7654321		Up	Not Fwding


```
-----
Subscriber-interface : sub-int-1
Group-interface      : grp-int-1-1
Sub Profile          : sub-prof-2
SLA Profile          : sla-prof-2
App Profile          : N/A
-----
Number of static hosts : 1
=====
*A:BNG# #
```

Once the MAC address of the host is learned through the ARP protocol, the forwarding state is set to *Fwding* and traffic now can be forwarded to and from that host.

```
*A:BNG# show service id 1 static-host
=====
Static Hosts for service 1
=====
Sap      IP Address      Configured MAC      Dynamic MAC
Subscriber                               Admin State         Fwding State
-----
1/1/1:111      10.1.1.101      N/A                 00:00:00:99:99:99
03-7654321      Up               Fwding
-----
Number of static hosts : 1
=====
*A:BNG# #
```

ESM Data Retrieval through RADIUS

A RADIUS server can provide ESM data, including:

- subscriber-id
- address information
- sub-profile-string
- sla-profile-string
- etc.

The BNG sends a RADIUS Access-Request including the User-Name attribute identifying the host for authentication purposes. The User-Name format is configurable on the BNG:

- MAC-address
- circuit-ID, remote-ID
- user@domain
- etc.

An excerpt from the Free-RADIUS user file used for this example is shown below. The first line of a block contains the User-Name with the credentials.

```
00:00:00:22:22:22  Auth-Type := Local, Cleartext-Password := ""
    Alc-Subsc-ID-Str = "sub-22",
    Alc-Subsc-Prof-Str = "sub-string-1",
    Alc-SLA-Prof-Str = "sla-string-1"

BSAN64|40|1/1/2:300 Auth-Type := Local, User-Password == "LetMeIn"
```

```
Alc-Subsc-ID-Str = "subscriber-300",
Alc-Subsc-Prof-Str = "subpro1-string",
Alc-SLA-Prof-Str = "slapro1-string",
Session-Timeout = 600

sub202@provider Cleartext-Password := "sub202"
Alc-Subsc-ID-Str = "sub-44",
Alc-Subsc-Prof-Str = "sub-string-3",
Alc-SLA-Prof-Str = "sla-string-22",
Framed-IP-Address = 10.2.1.202,
Framed-IP-Netmask = 255.255.255.0
```

The lines following the User-Name with the credentials lists the attributes the RADIUS server returns after successfully authenticating the user, and include:

- Alc-Subsc-ID-Str
- Alc-Subsc-Prof-Str
- Alc-SLA-Prof-Str
- Framed-IP-Address
- Framed-IP-Netmask
- etc.

Authentication Policy

An authentication policy is needed for retrieval of the data from a RADIUS server, indirectly indicating the RADIUS server(s) to contact.

The following steps are needed to create an authentication policy:

1. Define one or more RADIUS servers.
 2. Define a RADIUS server policy.
 3. Define an authentication policy.
1. The example defines a single RADIUS server, indicating the name, the address and the secret to use.

```
configure
router
radius-server
server "server-1" address 192.168.202.84 secret secret-1 create
exit
```

2. The radius-server-policy refers to the server defined above. In this example the Base router is used to reach the RADIUS server.

```
configure
aaa
radius-server-policy "rad-serv-pol-1" create
description "Radius AAA server policy"
servers
router "Base"
server 1 name "server-1"
exit
exit
```

3. The authentication policy specifies to include the circuit-id and the remote-id attributes towards the RADIUS server as well as the radius-server-policy to use. Multiple authentication policies can be defined.

```
configure
  subscriber-mgmt
    authentication-policy "auth-pol-1" create
      pppoe-access-method pap-chap
      include-radius-attribute
        circuit-id
        remote-id
    exit
  radius-server-policy "rad-serv-pol-1"
```

The authentication policy is applied at a group-interface. The example below shows both group interfaces using the same authentication policy.

```
configure
  service
    ies 1
      subscriber-interface "sub-int-1"
        group-interface "grp-int-1-1"
          ---snip---
          authentication-policy "auth-pol-1"
          ---snip---
        exit
      group-interface "grp-int-1-2"
        ---snip---
        authentication-policy "auth-pol-1"
        ---snip---
      exit
```

DHCP Host

The following command shows the strings returned by the RADIUS server for the DHCPv4 host with MAC-address 00:00:00:22:22:22.

```
A:BNG# show service id 1 dhcp lease-state mac 00:00:00:22:22:22 detail
=====
DHCP lease states for service 1
=====
Service ID           : 1
IP Address           : 10.1.1.9
Client HW Address    : 00:00:00:22:22:22
Subscriber-interface : sub-int-1
Group-interface      : grp-int-1-1
SAP                  : 1/1/1:112
Up Time              : 0d 00:03:30
Remaining Lease Time : 9d 23:56:30
Remaining SessionTime: N/A
Persistence Key      : N/A

Sub-Ident            : "sub-22"
Sub-Profile-String   : "sub-string-2"      # profile string before translation
SLA-Profile-String   : "sla-string-2"      # profile string before translation
App-Profile-String   : ""
Lease ANCP-String    : ""
```

```

Lease Int Dest Id      : ""
Category-Map-Name     : ""

---snip---

Lease-Time             : 10d 00:00:00
DHCP Server Addr      : 10.11.11.1
Radius User-Name       : "00:00:00:22:22:22"
-----
Number of lease states : 1
=====
A:BNG#

```

The strings returned by the RADIUS server are translated according the subscriber identification profile for retrieval of the actual profile data.

The actual profiles being used can be found using following command.

```

*A:BNG# show service id 1 subscriber-hosts mac 00:00:00:22:22:22 detail
=====
Subscriber Host table
=====
Sap      Subscriber
IP Address
MAC Address      PPPoE-SID Origin      Fwding State
-----
1/1/1:112      sub-22
10.1.1.6
00:00:00:22:22:22  N/A      DHCP      Fwding
-----
Subscriber-interface : sub-int-1
Group-interface      : grp-int-1-1
Sub Profile          : sub-prof-2      # profile name after translation
SLA Profile          : sla-prof-2      # profile name after translation
App Profile          : N/A
Egress Q-Group       : N/A

---snip---

-----
Number of subscriber hosts : 1
=====
*A:BNG#

```

PPP Host

The following command shows the strings returned by the RADIUS server for the PPP host with user-name sub202@provider.

```

*A:BNG# show service id 1 ppp session user-name sub202@provider detail
=====
PPP sessions for service 1
=====

User-Name          : sub202@provider

Description         : svc:1 sap:1/1/1:212 mac:00:00:00:44:44:44 sid:1
Up Time            : 0d 00:17:13
Type               : oE
Termination        : local

```

```

IP/L2TP-Id/If-Id      : 10.2.1.202
MC-Standby            : No
Session Time Left     : N/A

---snip---

Subscriber            : "sub-44"
Sub-Profile-String    : "sub-string-3" # profile string before translation
SLA-Profile-String    : "sla-string-22" # profile string before translation
ANCP-String           : ""
Int-Dest-Id           : ""
App-Profile-String    : ""
Category-Map-Name     : ""

---snip---

Radius Class          :
Radius User-Name      : sub202@provider
---snip---
-----
No. of sessions: 1
=====
*A:BN# #

```

The actual profiles being used can be found using following command.

```

*A:BN# show service id 1 subscriber-hosts mac 00:00:00:44:44:44 detail
=====
Subscriber Host table
=====
Sap      Subscriber
IP Address
MAC Address      PPPoE-SID Origin      Fwding State
-----
1/1/1:212      sub-44
10.2.1.202
00:00:00:44:44:44      1          IPCP          Fwding
-----
Subscriber-interface : sub-int-2
Group-interface      : grp-int-2-1
Sub Profile          : sub-prof-3 # profile name after translation
SLA Profile          : sla-prof-20 # profile name after translation
App Profile          : N/A

---snip---

-----
Number of subscriber hosts : 1
=====
*A:BN#

```

For both the DHCPv4 host as well as the PPPv4 host, the output aligns with the data provided by the RADIUS server and the defined profile maps.

ESM Data Retrieval through a Local User Database

A Local User Database (LUDB) can provide ESM data for DHCP and PPP hosts, including:

- subscriber-id
- address information

- sub-profile string
- sla-profile string
- MSAP service-id
- retail service-id
- etc.

Retrieval of data in an LUDB requires matching criteria, which can be one of the following items, or a combination of the following items, including:

- MAC-address
- circuit-id/remote-id
- username
- SAP-id
- etc.

The example below defines an LUDB named *ludb-1* which matches DHCP hosts by means of MAC address, and PPP hosts by means of username.

```
configure
  subscriber-mgmt
    local-user-db "ludb-1" create
      dhcp
        match-list mac
        host "host-121" create
          host-identification
            mac 00:00:00:aa:aa:aa
          exit
          address 10.1.1.121
          identification-strings 254 create
            subscriber-id "sub-121"
            sla-profile-string "sla-string-3"
            sub-profile-string "sub-string-2"
          exit
          options
            subnet-mask 255.255.255.0
            default-router 10.1.1.254
          exit
          ipv6-wan-address-pool "pool6-2"
          ipv6-delegated-prefix-pool "pool6-2"
          no shutdown
        exit
        host "host-122" create
          host-identification
            mac 00:00:00:bb:bb:bb
          exit
          auth-policy "auth-pol-1"
          no shutdown
        exit
      exit
    exit
  ppp
    match-list username
    host "host-123" create
      host-identification
        username "user@domain"
      exit
      address 10.1.1.123/32
      password pap user
      identification-strings 254 create
```

```
        subscriber-id "sub-123"
        sla-profile-string "sla-string-3"
        sub-profile-string "sub-string-1"
    exit
    no shutdown
exit
exit
no shutdown
exit
```



Note:

- The entire LUDB can be disabled.
- Individual host entries can be disabled.

An LUDB can be applied at group interface level in different contexts:

- dhcp
- ipv6 dhcp6
- ppp
- pppoe

For the LUDB to provide ESM data, no authentication policy may be applied at group interface level, and the LUDB itself must be in the no shutdown state.

The following example shows the LUDB named *ludb-1* applied to group-interface *grp-int-1-2* in the dhcp and the pppoe context.

```
configure
service
  ies 1
    subscriber-interface "sub-int-1"
      address 10.1.1.254/24
      address 10.1.2.254/24
      group-interface "grp-int-1-2"
        dhcp
          ---snip---
          gi-address 10.1.1.254
          user-db "ludb-1"
          no shutdown
        exit
      no authentication-policy
      sap 1/1/1:121 create
        sub-sla-mgmt
          def-sub-profile "sub-prof-1"
          def-sla-profile "sla-prof-1"
          sub-ident-policy "sub-id-pol-1"
          multi-sub-sap
          no shutdown
        exit
      exit
    sap 1/1/1:122 create
      sub-sla-mgmt
        def-sub-profile "sub-prof-1"
        def-sla-profile "sla-prof-1"
        sub-ident-policy "sub-id-pol-1"
        multi-sub-sap
        no shutdown
      exit
    exit
```

```

pppoe
  session-limit 100
  user-db "ludb-1"
  no shutdown
exit
exit

```

For *host-121*, with MAC-address 00:00:00:aa:aa:aa, all ESM data is provided by the LUDB. The detailed DHCP lease state shows the actual profile strings, the subscriber-ID and the IP address used. The Lease Info origin is set to UserDb.

```

A:BNG# show service id 1 dhcp lease-state mac 00:00:00:aa:aa:aa detail
=====
DHCP lease states for service 1
=====
Service ID           : 1
IP Address           : 10.1.1.121
Client HW Address    : 00:00:00:aa:aa:aa
Subscriber-interface : sub-int-1
Group-interface      : grp-int-1-2
SAP                  : 1/1/1:121
---snip---
Sub-Ident            : "sub-121"
Sub-Profile-String   : "sub-string-2"
SLA-Profile-String   : "sla-string-3"
---snip---
Lease Info origin    : UserDb
---snip---
Ip-Netmask           : 255.255.255.0
Broadcast-Ip-Addr    : 10.1.1.255
Default-Router       : 10.1.1.254
Primary-Dns          : N/A
Secondary-Dns        : N/A
---snip---
Relay Agent Information
  Circuit Id         : BNG|1|grp-int-1-2|1/1/1:121
Radius User-Name     : ""
-----
Number of lease states : 1
=====
*A:BNG#

```

The detailed DHCPv6 lease state for the same host is shown below. The Lease State info now is DHCP for both IA-NA as well as for IA-PD where the IA-NA and the IA-PD are allocated by the DHCP server based on the ipv6-wan-address-pool and the ipv6-delegated-prefix-pool respectively, as indicated by the LUDB.

```

*A:BNG# show service id 1 dhcp6 lease-state mac 00:00:00:aa:aa:aa detail
=====
DHCP lease states for service 1
=====
Service ID           : 1
IP Address           : 2001:DB8:201::1/128
Client HW Address    : 00:00:00:aa:aa:aa
Subscriber-interface : sub-int-1
Group-interface      : grp-int-1-2
SAP                  : 1/1/1:121
---snip---
Sub-Ident            : "sub-121"
Sub-Profile-String   : "sub-string-2"
SLA-Profile-String   : "sla-string-3"
---snip---
Pool Name            : "pool6-2"

```



```
Dhcp6 Server Addr : 2001:DB8::11
---snip---
Lease Info origin : DHCP
---snip---
Radius User-Name : ""
-----
Service ID : 1
IP Address : 2001:DB8:202::/56
Client HW Address : 00:00:00:aa:aa:aa
Subscriber-interface : sub-int-1
Group-interface : grp-int-1-2
SAP : 1/1/1:121
---snip---
Sub-Ident : "sub-121"
Sub-Profile-String : "sub-string-2"
SLA-Profile-String : "sla-string-3"
---snip---
Pool Name : "pool6-2"
Dhcp6 Server Addr : 2001:DB8::11
---snip---
Lease Info origin : DHCP
---snip---
Radius User-Name : ""
-----
Number of lease states : 2
=====
*A:BNG#
```

Because only an authentication policy is defined for *host-122*, with MAC address 00:00:00:bb:bb:bb, the profile strings and the subscriber-ID are provided by the RADIUS server. The IP address is provided by the DHCP server. The Lease Info origin is set to DHCP.

```
A:BNG# show service id 1 dhcp lease-state mac 00:00:00:bb:bb:bb detail
=====
DHCP lease states for service 1
=====
Service ID : 1
IP Address : 10.1.1.5
Client HW Address : 00:00:00:bb:bb:bb
Subscriber-interface : sub-int-1
Group-interface : grp-int-1-2
SAP : 1/1/1:122
---snip---
Sub-Ident : "sub-122"
Sub-Profile-String : "sub-string-3"
SLA-Profile-String : "sla-string-1"
---snip---
Lease Info origin : DHCP

Ip-Netmask : 255.255.255.0
Broadcast-Ip-Addr : N/A
Default-Router : 10.1.1.254
---snip---
Lease-Time : 10d 00:00:00
DHCP Server Addr : 10.11.11.1

Relay Agent Information
Circuit Id : BNG|1|grp-int-1-2|1/1/1:122
Radius User-Name : "00:00:00:bb:bb:bb"
-----
Number of lease states : 1
=====
```

```
A:BNG#
```

For *host-123*, the subscriber-ID, the profile strings and the IP-address are provided by the LUDB. Similar to *host-121*, no interaction with a RADIUS server is needed at all. The IP origin is set to local-user-db.

```
A:BNG# show service id 1 ppp session user-name "user@domain" detail
=====
PPP sessions for service 1
=====

User-Name           : user@domain
Description          : svc:1 sap:1/1/1:122 mac:00:00:00:cc:cc:cc sid:1
Up Time              : 0d 00:46:10
Type                 : oE
---snip---
PPP MTU              : 1492
PPP Auth-Protocol    : PAP
PPP User-Name        : user@domain

Subscriber-interface : sub-int-1
Group-interface      : grp-int-1-2

IP Origin           : local-user-db
DNS Origin           : none
NBNS Origin          : none

Subscriber         : "sub-123"
Sub-Profile-String : "sub-string-1"
SLA-Profile-String : "sla-string-3"
---snip---
IP Address         : 10.1.1.123/32
---snip---
Circuit-Id           :
Remote-Id            :

Radius Session-T0    : N/A
Radius Class          :
Radius User-Name     :
Logical-Line-Id      :
-----
No. of sessions: 1
=====
A:BNG#
```

LUDB details are out of the scope of this example.

Optional ESM Data

Limits

The maximum number of hosts, subscribers, sessions and leases are checked during the host instantiation process, and can be defined at following levels:

- sla-profile level: host-limit
- sap-level: multi-sub-sap
- group-interface level: lease-populate, session-limit, sap-session-limit

The sla-profile optionally defines a maximum number of hosts allowed by this profile. An example is shown below.

```
configure
  subscriber-mgmt
    sla-profile "sla-prof-3" create
      host-limit 4
    ingress
      qos 100
    exit
  egress
    qos 100
  exit
exit
```

By default only hosts from a single subscriber can connect through a SAP. This condition can be lifted as follows.

```
configure
  service
    ies 1
      subscriber-interface "sub-int-1"
        group-interface "grp-int-1-1" sub-sla-mgmt
          sap 1/1/1:112 create
            sub-sla-mgmt
              def-sub-profile "sub-prof-1"
              def-sla-profile "sla-prof-1"
              sub-ident-policy "sub-id-pol-2"
              multi-sub-sap
              no shutdown
            exit
          exit
        exit
      exit
    exit
```

The maximum number of leases assigned to subscriber hosts by the DHCP server can be defined at group interface level as follows.

```
configure
  service
    ies 1
      subscriber-interface "sub-int-1"
        address 10.1.1.254/24
        group-interface "grp-int-1-1"
          dhcp
            proxy-server
              emulated-server 10.1.1.254
            no shutdown
          exit
        server 10.11.11.1
        trusted
        lease-populate 100
        client-applications dhcp ppp
        gi-address 10.1.1.254
        no shutdown
      exit
    exit
```

The maximum number of sessions for PPP as well as for PPPoE, including the sap-session-limit, is also defined at group interface level, as the example below shows.

```
configure
service
  ies 1
    subscriber-interface "sub-int-1"
    group-interface "grp-int-1-1"
      ppp
        session-limit 50
      exit
      pppoe
        session-limit 40
        sap-session-limit 30
      no shutdown
    exit
```

Filtering

Filter policies allow for selectively dropping, forwarding or redirecting ingress/egress traffic. They are also known as access control lists (ACLs) and can optionally be included in the SLA-profile.

An example sla-profile with IP-filters is show below.

```
configure
subscriber-mgmt
  sla-profile "sla-prof-3" create
  host-limit 4
  ingress
    qos 100
    exit
    ip-filter 1
    ipv6-filter 1
  exit
  egress
    qos 100
    exit
    ip-filter 2
    ipv6-filter 2
  exit
```

Filter policies must be defined before they can be used in an SLA profile.

Accounting

Accounting policies define how to count the traffic for billing purposes on a per service basis. Two types of accounting are available:

- RADIUS accounting
- XML accounting

For RADIUS accounting, the accounting data is stored on the RADIUS accounting server. For XML accounting, the accounting data is stored locally on a flash-disk on the node.

The example below shows the definition of a radius-accounting-policy and reuses the radius-server-policy defined before.

```
configure
  subscriber-mgmt
    radius-accounting-policy "rad-acct-pol-1" create
      host-accounting interim-update
      update-interval 5
      include-radius-attribute
        framed-ip-addr
        sla-profile
        sub-profile
      exit
      radius-accounting-server
        router "Base"
        server 1 address 192.168.202.84 secret secret-1
      exit
    radius-server-policy "rad-serv-pol-1"
  exit
```

The radius-accounting policy is referred to from the SUB profile as shown in the example below.

```
configure
  subscriber-mgmt
    sub-profile "sub-prof-1" create
      radius-accounting-policy "rad-acct-pol-1"
      ingress
        policer-control-policy "pol-ctrl-1"
      exit
    exit
    egress
      scheduler-policy "down-1"
    exit
  exit
exit
```

For XML accounting, an accounting policy is referred to from the SUB profile.

```
configure
  subscriber-mgmt
    sub-profile "sub-prof-2" create
      accounting-policy 10
      collect-stats
      ingress
        scheduler-policy "sched-up-1"
      exit
      egress
        scheduler-policy "sched-down-1"
      exit
    exit
  exit
```

The **collect-stats** command activates the generation of accounting files.

Accounting policies must be defined before they can be used in a SUB profile.

RADIUS accounting and XML accounting details are out of the scope of this example.

Application Profile

An application profile is needed for supporting Application Assurance.

Traffic is diverted to an MS-ISA MDA which processes the traffic according the application profile, which for that purpose is assigned to:

- an ESM subscriber, or an group of ESM subscribers
- a SAP
- a spoke SDP

ESM subscribers are assigned an application profile either statically or dynamically.

The example below shows the assignment of an application profile to a static host.

```
configure
  service
    ies 1
      subscriber-interface "sub-int-1"
        group-interface "grp-int-1-1"
          sap 1/1/1:111
            description "sap for customer 1"
            static-host ip 10.1.1.101 create
            app-profile "app-prof-3"
            sla-profile "sla-prof-2"
            sub-profile "sub-prof-1"
            subscriber 03-7654321
            no shutdown
          exit
        exit
      exit
```

For dynamic hosts, the same rules apply as for the sub-profile scriber and the sla-profile, meaning that the app-profile can be found through:

1. LUDB
2. RADIUS
3. DHCPv4 ACK python
4. Default

The example below shows the assignment of a default application profile at the SAP level.

```
configure
  service
    ies 1
      subscriber-interface "sub-int-1"
        group-interface "grp-int-1-1" sub-sla-mgmt
          sap 1/1/1:112 create
            sub-sla-mgmt
              def-sub-profile "sub-prof-1"
              def-sla-profile "sla-prof-1"
              def-app-profile "app-prof-1"
              sub-ident-policy "sub-id-pol-2"
              multi-sub-sap
              no shutdown
            exit
          exit
        exit
      exit
```

Application profiles must be defined before they can be referenced.

Application Assurance and Application Assurance subscribers details are out of the scope of this example.

Intermediate Destinations

An intermediate destination is an aggregation point in the network and identified through the intermediate destination identity (inter-dest-id).

Most typically Access Nodes (ANs) are considered intermediate destinations.

The inter-dest-id is an optional per subscriber attribute.

The intermediate destination is used in the BNG for supporting QoS, as an example:

- By shaping the aggregate rate of all egress traffic of subscribers connected to an AN to prevent congestion of the link towards that AN.
- To ensure fairness amongst the subscriber hosts connected to that AN. For that purpose the inter-dest-id is linked to an hierarchical scheduler via a virtual port (Vport).

The example below shows the assignment of the inter-dest-id to a static host.

```
configure
  service
    ies 1
      subscriber-interface "sub-int-1"
        group-interface "grp-int-1-1"
          sap 1/1/1:111
            description "sap for customer 1"
            static-host ip 10.1.1.101 create
              inter-dest-id "bsan-1"
              sla-profile "sla-prof-2"
              sub-profile "sub-prof-1"
              subscriber 03-7654321
              no shutdown
            exit
          exit
        exit
      exit
    exit
```

The inter-dest-id does not to be translated.

Intermediate destination details are is out of the scope of this example.

Conclusion

This chapter explains basic ESM concepts including the definition of subscribers and subscriber hosts. The host identification and instantiation process was explained, indicating the mandatory and optional ESM data fetched during the process (address information, sub-profile, sla-profile, app-profile, inter-dest-id, etc.). The address information retrieval process was explained. The different sources from where the ESM data can originate were mentioned, including some examples. Also the sub-ident-policy and the authentication-policy, needed during the identification and instantiation process, are explained in detail.

ESM 128-bit Mode for DHCPv6 IA_NA WAN Hosts

This chapter describes ESM 128-bit Mode for DHCPv6 IA_NA WAN Hosts.

Topics in this chapter include:

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

Applicability

The information and configuration in this chapter are based on SR OS Release 16.0.R4 and cover both IPoE and PPPoE subscribers.

Basic knowledge of ESM is a prerequisite.

Overview

Nokia Triple Play Service Delivery Architecture (TPSDA) supports both 64-bit and 128-bit WAN mode. This chapter describes 64-bit and 128-bit WAN mode for DHCPv6 IA_NA WAN hosts and provides an example 128-bit WAN mode subscriber interface configuration for IPoE IPv6 hosts.

The 64-bit and 128-bit WAN mode is a creation time configurable parameter on the subscriber interface. The parameter specifies the host addressing schema for a DHCPv6 WAN host (DHCPv6 IA_NA). By default, all subscriber interfaces use the 64-bit WAN mode. This means that each WAN host created on the system would consume an entire /64 prefix. More specifically, the WAN host would construct only one 128-bit address out of the entire /64 prefix. By comparison, the 128-bit WAN mode allows a single 128-bit address for each DHCPv6 WAN host. Therefore, when using DHCPv6 IA_NA, the 128-bit WAN mode reduces IPv6 address consumption. All ESM features are supported regardless of the configured WAN mode.

In 64-bit mode, the subscriber interface expects all WAN hosts to have /64 unique prefixes. The system will only distinguish individual IPv6 subscribers by the first 64 bits of an IPv6 address. When IPv6 WAN addresses are /64 unique, a WAN host consumes an entire /64 prefix (with the exception of IPoE-bridged mode, described later). For example, if the first IPv6 WAN host is assigned an DHCPv6 IA_NA address of 2001::1, the next address must be a 64-bit increment. Therefore, the second IPv6 WAN host must be assigned a DHCPv6 IA_NA address of 2001:0:0:1::1. In this mode, all WAN hosts are assigned a /64 prefix regardless of whether the host is using SLAAC or DHCPv6. This provides the flexibility for allowing the subscriber to choose between using SLAAC or DHCPv6 when assigned a /64 prefix.

The 128-bit mode allows the system to distinguish individual IPv6 WAN hosts using the full 128-bit address. For example, if the first IPv6 WAN host is assigned an address of 2001::1, the next host address can be a 128-bit increment, which is 2001::2. Compared to 64-bit mode, a single WAN host no longer consumes an entire /64 prefix. When the system is operating in 128-bit mode, a /96 prefix is automatically

created in the FIB to assist in subscriber host lookup. Therefore, in 128-bit mode, subnets using a prefix length of 96 or longer are recommended. Within a /96 subnet, about 4 billion IPv6 addresses (WAN hosts) can be supported.

Extra considerations are required for subscriber interfaces operating in 128-bit mode.

- All subscribers within a subscriber interface should use a single /96 prefix (or longer; for example, /97).
- In the case where multiple prefixes are required (for example, service differentiation), Nokia recommends keeping the number of /96 prefixes to a minimum.
- The 128-bit WAN mode can support SLAAC hosts. The system will not derive a /96 prefix from the SLAAC host. However, the system does not allow a SLAAC host and a DHCPv6 IA_NA host to share the same subscriber interface prefix. SLAAC hosts must use a different prefix than a DHCPv6 IA_NA host.
- The routing instance (VPRN or IES) supports a mixture of 64-bit and 128-bit WAN mode subscriber interfaces.

The local DHCPv6 server supports assigning both incremental DHCP IA_NA /64 and /128 addresses. The local DHCP server will determine which subscriber interface the DHCP IA_NA request is from. With this information, the server will automatically determine whether the next incremental /64 or /128 address should be assigned. There is no configuration required on the local DHCP server; this is performed automatically by the server.

In addition, it is possible to use the local DHCP server for DHCPv6 IA_NA address assignment while using AAA for DHCPv6 IA_PD prefix assignment. For this, the local address assignment feature is required (configured under group-interface). The client application is IPoE WAN and the DHCPv6 IA_NA address is assigned via a pool name retrieved from AAA or LUDB. The local address assignment also supports incremental /64 and /128 address assignment. This is also determined automatically by the local DHCP server without requiring any additional configuration.

For IPoE-bridged mode, Nokia recommends using only the 64-bit WAN mode. IPoE-bridged mode allows a single subscriber up to 128 bridge hosts sharing the same /64 prefix. Nokia does not recommend using 128-bit WAN mode when IPoE bridging is required. This is due to the 128-bit WAN mode generating at least one /96 prefix per subscriber. There is a limit of prefixes supported per subscriber interface. In 64-bit WAN mode, the subscriber is assigned an entire /64 prefix, but the mode allows each host to have an incremental 128-bit address.



Note:

In 64-bit mode, all bridge hosts must use the same SLA profile, while in 128-bit mode, each DHCPv6 IA_NA host can have a unique SLA profile.

Table 1 compares the 64-bit WAN mode with the 128-bit WAN mode.

Table 6: 64-bit WAN Mode Versus 128-bit WAN Mode

	64-bit WAN mode	128-bit WAN mode	Comments
IPv6 WAN hosts	IPv6 WAN hosts must be assigned 64-bit incremental addresses.	IPv6 WAN hosts can be assigned 128-bit incremental addresses.	
DHCPv6 local-dhcp-server	Automatically assigns incremental /64	Automatically assigns incremental /128 addresses to DHCP IA_NA requests.	Requires no configuration.

	64-bit WAN mode	128-bit WAN mode	Comments
	addresses to DHCP IA_NA requests.		
Local-address-assignment for DHCPv6 IA_NA	Automatically assigns incremental /64 addresses to DHCP IA_NA requests.	Automatically assigns incremental /128 addresses to DHCP IA_NA requests.	Local address assignment is used when the local DHCP server is used to assign the DHCPv6 IA_NA addresses while AAA assigns the DHCPv6 IA_PD addresses.
IPoE-bridged-mode	Supported. Recommended mode. Up to 128 bridge hosts can share the same /64 prefix. All bridge hosts must share the same SLA profile.	Supported. Not recommended mode. Up to 128 bridge hosts can share the same /64 prefix. DHCPv6 IA_NA hosts can have unique SLA profiles SLAAC hosts must share the same SLA profile.	
FIB	No FIB impact.	A host creation will automatically install a /96 FIB entry. Hosts that share the same /96 will not create new FIB entries. If all hosts within the same /96 entry are removed from the system, the /96 FIB entry will automatically be removed.	128-bit WAN mode can impact FIB scaling if used incorrectly. 128-bit WAN mode is intended for deployments where all WAN hosts under a single subscriber interface share a /96 prefix (or smaller). A /96 prefix can serve up to 4 billion WAN subscribers.

Configuration

Create a 128-bit WAN mode subscriber interface for IPoE IPv6 hosts.

As already indicated, the WAN mode is a creation-time parameter. It is not possible to toggle between 64-bit and 128-bit mode once the subscriber interface is created.

1. Create a subscriber interface as a 128-bit WAN mode interface, as follows:

```
*A:BNG-1>config>service>ies# info
-----
description "BNG-1"
subscriber-interface "sub-int-1" wan-mode model28 create
```

2. Following is the rest of a subscriber management configuration on the BNG. For 128-bit WAN mode, the prefix length for the WAN host is 96. Nokia recommends using prefix length 96 or longer when using 128-bit WAN mode.

```
*A:BNG-1>config>service>ies# info
-----
description "BNG-1"
subscriber-interface "sub-int-1" wan-mode model28 create
address 10.255.255.253/8
ipv6
  subscriber-prefixes
    prefix 2001:db8::/96 wan-host
  exit
exit
group-interface "group-int-1" create
dhcp
  server 192.168.0.1
  lease-populate 10
  client-applications dhcp
  gi-address 10.255.255.253
  no shutdown
exit
ipv6
  dhcp6
    proxy-server
      no shutdown
    exit
    relay
      link-address 2001:db8::
      server 2001:db9::1
      client-applications dhcp
      no shutdown
    exit
  exit
exit
sap 1/1/5:4 create
  sub-sla-mgmt
    def-sub-id use-sap-id
    def-sub-profile "sub-profile-1"
    def-sla-profile "sla-profile-1"
    sub-ident-policy "sub-ident-policy-1"
    multi-sub-sap 10
    no shutdown
  exit
exit
exit
exit
```

With the preceding configuration, the subscriber interface is ready to support 128-bit incremental IPv6 address assignment, as follows:

```
*A:BNG-1> show service active-subscribers hierarchy

=====
Active Subscribers Hierarchy
=====
-- sub_1 (sub-profile-1)
  |
  +- sap:1/1/5:4 - sla:sla-profile-1
    |
```

```

+-- IP0E-session - mac:00:00:10:10:13:13 - svc:1
|
+-- 2001:db8::1/128 - DHCP6

-- sub_2 (sub-profile-1)
|
+-- sap:1/1/5:4 - sla:sla-profile-1
|
+-- IP0E-session - mac:00:00:10:10:13:14 - svc:1
|
+-- 2001:db8::2/128 - DHCP6

-----
Number of active subscribers : 2
Flags: (N) = the host or the managed route is in non-forwarding state
=====

```

If the subscriber interface is in 64-bit WAN mode, the system will reject a subsequent subscriber host that requests addresses that overlap with the /64 prefix of existing hosts, as follows:

```

*A:eng-BNG-2>config>service>vprn>sub-if>grp-if>ipv6# show log log-id 99

=====
Event Log 99
=====
Description : Default System Log
Memory Log contents [size=500 next event=367512 (wrapped)]

367507 2018/09/19 20:18:37.656 UTC WARNING: DHCP #2005 IES1 Lease State Population Error
"Lease state table population error on SAP 1/1/5:4 in service 1 - Host with IP 2001:db8::2/
128 and MAC 00:00:10:10:13:14 conflicts with existing host in service 1"
=====

```

Conclusion

Different operators have unique IPv6 WAN addressing requirements. The 64-bit WAN mode is most suitable for assigning a /64 prefix to a subscriber (WAN host). This mode gives the subscriber flexibility to either use the /64 prefix for SLAAC or one address out of the /64 prefix for a DHCPv6 IA_NA address. The 128-bit WAN mode is for operators who prefer to use DHCPv6 IA_NA addressing. With 128-bit mode, each subscriber is assigned a single 128-bit address, minimizing IPv6 WAN address consumption. By using different WAN modes, addressing scaling and efficiency can be improved.

ESM IPv4: Multicast in a Wholesale/Retail Scenario

This chapter describes ESM IPv4 multicast configurations in a wholesale/retail scenario.

Topics in this chapter include:

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

Applicability

This chapter applies to SR OS routers, covers multicast in a wholesale and retail scenario for both IPoE and PPPoE subscribers, and was originally written for SR OS Release 11.0.R1. The CLI is updated to Release 15.0.R3.

Overview

Triple Play Service Delivery Architecture (TPSDA) allows operators to integrate High Speed Internet (HSI), voice and video services within a single network. The goal of this configuration example is to provide a walkthrough of a wholesale/retail multicast setup.

There are two wholesale/retail models in TPSDA. In the first model, a retail service instance is co-located with the wholesale service on the wholesale router whereas in the second model, for PPP services only, the retail service is on a separate BNG. [Figure 39: Wholesale/Retail Model 1](#) shows the first model, and consists of two 7750s. The first is a wholesaler Broadband Network Gateway (BNG) with a co-located retail service and the second is a retailer router. [Figure 40: Wholesale/Retail Model 2](#) shows the second model where the retail service is hosted by a separate router and the connection between the wholesale and retail utilizes Layer 2 Tunneling Protocol (L2TP).

Figure 39: Wholesale/Retail Model 1





Figure 40: Wholesale/Retail Model 2



In both models, the 7450 is used as an aggregation switch to aggregate the traffic of the PPPoE and IPoE subscribers. The wholesale router and the retail router are 7750s. The wholesale router is connected to the aggregation switch and to the retail router. The retail router is connected directly to the multicast source. The 7450 can only be an L2TP Access Concentrator (LAC), whereas the 7750 can be a LAC or an L2TP Network Server (LNS).

There are two basic requirements for a subscriber to receive multicast streams. First, the group interface for the subscribers must have IGMP enabled. Second, the Enhanced Subscriber Management (ESM) subscriber must be allowed to receive multicast streams by having IGMP enabled. When both requirements are met, the BNG will process the subscribers' IGMP messages, otherwise, IGMP messages are dropped. All customer premise equipment (CPE) originated IGMP messages are aggregated via the 7450 and passed onto the wholesale BNG. It is always the retail VPRN that processes the IGMP messages. The wholesale VPRN SAPs forward the actual multicast streams.

Configuration

Note that basic knowledge of multicast and ESM is assumed.

ESM Wholesale-Retail Multicast

There are various ways to provide wholesale and retail multicast function.

- For the IPoE and PPPoE Layer 3 wholesale/retail model, the wholesale and the retail services reside on separate VPRNs.
- For the PPPoE Layer 2 wholesale/retail model, L2TP is used.

ESM Layer 3 Wholesale-Retail Multicast

[#unique_137/unique_137_Connect_42_d2e4](#)Figure 41: Layer 3 Wholesale/Retail depicts a Layer 3 wholesale/retail scenario for both IPoE and PPPoE. The first BNG contains both the wholesale and retail configuration. There are two options for the retail BNG to deliver the multicast streams to the wholesale BNG:


1. MVPN between the BNGs

or

2. If using a routed interface between the BNGs, multicast routing is required.

This example uses the second option for delivery of the multicast streams in order to keep the configuration simple.

Figure 41: Layer 3 Wholesale/Retail



Step 1 - Configure wholesale service on BNG-1

The configuration for the wholesale service on BNG-1 with the group interface added to IGMP is as follows. This configuration applies to both IPoE and PPPoE.

```
# on BNG-1
configure
service
  vprn 1 customer 1 create
    description "wholesale"
    route-distinguisher 65536:1
    interface "system" create
      address 192.0.2.20/32
      loopback
    exit
    subscriber-interface "sub-vprn-1-ws" create
      unnumbered "system"
      group-interface "grp-vprn-1-ws" create
        arp-populate
        dhcp
        client-applications dhcp ppp
        no shutdown
      exit
      authentication-policy "auth-radius-1"
      sap 1/1/2:1 create
        sub-sla-mgmt
          def-sub-id use-sap-id
          def-sub-profile "sub-profile-1"
          def-sla-profile "sla-profile-1"
          sub-ident-policy "sub-ident-1"
          multi-sub-sap 10
          no shutdown
        exit
      exit
    pppoe
      session-limit 10
      sap-session-limit 10
      no shutdown
    exit
  exit
exit
igmp
```

```

        group-interface "grp-vprn-1-ws"
            no shutdown
        exit
        no shutdown
    exit
    no shutdown
exit
exit
exit
exit

```

Step 2 - Configure retail service on BNG-1

Also on BNG-1, a separate VPRN is configured for the retail provider. The retail configuration is a little different from the wholesale configuration. The configuration for the retail VPRN with IGMP and PIM enabled is as follows. This configuration is applicable to both IPoE and PPPoE. The multicast streams received in the retail VPRN are forwarded to the wholesale VPRN. Other retail VPRNs can offer multicast streams as well, and the same multicast addresses can be re-used as long as the address is assigned to a different retail VPRN.

```

# on BNG-1
configure
    service
        vprn 2 customer 2 create
            description "retail"
            route-distinguisher 65536:2
            interface "system" create
                address 172.16.1.1/32
                loopback
            exit
            interface "int-BNG-1-BNG-2" create
                address 192.168.12.1/30
                sap 1/1/1 create
                exit
            exit
            subscriber-interface "sub-vprn-2-rt" \
                fwd-service 1 fwd-subscriber-interface "sub-vprn-1-ws" create
                address 10.255.255.254/8
                dhcp
                    server 172.16.1.2
                    lease-populate 200
                    client-applications dhcp ppp
                    gi-address 10.255.255.254
                    no shutdown
                exit
            exit
        exit
    igmp
        group-interface fwd-service 1 "grp-vprn-1-ws"
            no shutdown
        exit
        no shutdown
    exit
    pim
        interface "int-BNG-1-BNG-2"
            exit
            no shutdown
        exit
    ospf
        area 0.0.0.0
            interface "sub-vprn-2-rt"
                interface-type point-to-point
                no shutdown
            exit
        exit
    exit

```



```
        exit
        interface "system"
            no shutdown
        exit

        interface "int-BNG-1-BNG-2"
            interface-type point-to-point
            no shutdown
        exit
    exit
    no shutdown
exit
no shutdown
exit
exit
exit
exit
```

Step 3 - Configure an IGMP policy

Per host replication is mandatory in a wholesale/retail scenario. A single SAP in a wholesale service might be shared among different retailers. A wholesale host that has requested a multicast group will always have the multicast delivered directly. Other hosts on the SAPs might belong to a different retailer and therefore 1) retailers might not have the same multicast group and sources and 2) their bandwidth should not be impacted by other hosts' multicast. Per-host replication is configured in the **igmp-policy igmp-policy-1**. This is mandatory for both IPoE and PPPoE subscribers.

```
# on BNG-1
configure
    subscriber-mgmt
        igmp-policy "igmp-policy-1" create
            per-host-replication
        exit
        sla-profile "sla-profile-1" create
        exit
        sub-profile "sub-profile-1" create
            igmp-policy "igmp-policy-1"
        exit
    exit
exit
```

Step 4 - Configure retail service on BNG-2

The interfaces are added to OSPF and to PIM on the retail BNG that is connected to the multicast source.

```
# on BNG-2
configure
    service
        vprn 2 customer 2 create
            dhcp
                local-dhcp-server "dhcp-RETAIL" create
                    use-gi-address
                    pool "pool-RETAIL-1" create
                        subnet 10.0.0.0/8 create
                            options
                                default-router 10.255.255.254
                            exit
                        address-range 10.0.0.1 10.0.0.100
                    exit
                exit
            exit
            no shutdown
        exit
    exit
```

```

route-distinguisher 65536:2
interface "system" create
    address 172.16.1.2/32
    local-dhcp-server "dhcp-RETAIL"
    loopback
exit
interface "int-BNG-2-BNG-1" create
    address 192.168.12.2/30
    sap 2/1/1 create
    exit
exit
interface "int-BNG-2-MCAST" create
    address 192.168.4.1/30
    sap 1/1/6 create
    exit
exit
pim
    interface "int-BNG-2-BNG-1"
    exit
    interface "int-BNG-2-MCAST"
    exit
    no shutdown
exit
ospf
    area 0.0.0.0
        interface "system"
            no shutdown
        exit
        interface "int-BNG-2-BNG-1"
            interface-type point-to-point
            no shutdown
        exit
        interface "int-BNG-2-MCAST"
            passive
            no shutdown
        exit
    exit
    no shutdown
exit
no shutdown
exit
exit
exit

```

With the configuration from all previous steps, the wholesale/retail setup is ready to process IGMP messages. Now an IGMPv3 request is sent to the wholesale SAP. The (S,G) is (192.168.4.2, 239.255.1.1) and the subscriber IP address is 10.0.0.1. The following output shows that the (S,G) is not registered in the wholesale VPRN, but in the retail VPRN.

```

*A:BNG-1# show router 1 igmp group
=====
IGMP Interface Groups
=====
No Matching Entries
=====
IGMP Host Groups
=====
No Matching Entries
=====
IGMP SAP Groups
=====
No Matching Entries
=====

```

```
*A:BNG-1#
```

```
*A:BNG-1# show router 2 igmp group
```

```
=====
```

```
IGMP Interface Groups
```

```
=====
```

```
No Matching Entries
```

```
=====
```

```
IGMP Host Groups
```

```
=====
```

```
(192.168.4.2,239.255.1.1)
```

```
  Fwd List   : 10.0.0.1
```

```
UpTime: 0d 00:01:41
```

```
-----
```

```
Entries : 1
```

```
=====
```

```
IGMP SAP Groups
```

```
=====
```

```
No Matching Entries
```

```
=====
```

```
*A:BNG-1#
```

The following command shows all subscribers' (S,G) pairs.

```
*A:BNG-1# show service active-subscribers igmp detail
```

```
=====
```

```
Active Subscribers Detail
```

```
=====
```

Subscriber HostAddr GrpAddr SrcAddr	IGMP-Policy GrpItf Type Type	Up-Time	NumGroups Mode Blk/Fwd
sub-video-1 10.0.0.1 239.255.1.1 192.168.4.2	igmp-policy-1 grp-vprn-1-ws Dynamic Dynamic	0d 00:02:16	1 Include Fwd

```
-----
```

```
Number of Subscribers : 1
```

```
=====
```

```
*A:BNG-1#
```

Only the retail VPRN is responsible for processing the IGMP messages. Therefore to troubleshoot a wholesale/retail setup, debug is only relevant on the retail router instance.

```
debug
  router "2"
    igmp
      group-interface fwd-service "1" "grp-vprn-1-ws"
      host "10.0.0.1"
      packet mode egr-ingr-and-dropped
    exit
  exit
exit
```

```
1 2017/07/06 13:58:14.57 CEST MINOR: DEBUG #2001 vprn2 IGMP[3]
"IGMP[3]: RX-PKT
[000 01:09:17.920] IGMP host 10.0.0.1 V3 PDU: 10.0.0.1 -> 224.0.0.22 pduLen 20
  Type: V3 REPORT maxrespCode 0x0 checksum 0x2352
  Num Group Records: 1
  Group Record 0
```

```
Type: ALW_NEW_SRCS, AuxDataLen 0, Num Sources 1
Mcast Addr: 239.255.1.1
Source Address List
    192.168.4.2
"

2 2017/07/06 13:58:14.57 CEST MINOR: DEBUG #2001 vprn2 IGMP[vprn2 inst 3]
"IGMP[vprn2 inst 3]: igmpIfGroupAdd
Adding 239.255.1.1 to IGMP host 10.0.0.1 database"

3 2017/07/06 13:58:14.57 CEST MINOR: DEBUG #2001 vprn2 IGMP[vprn2 inst 3]
"IGMP[vprn2 inst 3]: igmpProcessGroupRec
Process group rec ALW_NEW_SRCS received on host 10.0.0.1 for group 239.255.1.1 i
n mode INCLUDE. Num srcls 1"

4 2017/07/06 13:58:14.57 CEST MINOR: DEBUG #2001 vprn2 IGMP[vprn2 inst 3]
"IGMP[vprn2 inst 3]: igmpIfSrcAdd
Adding i/f source entry for host 10.0.0.1 (192.168.4.2,239.255.1.1) to IGMP fwdL
ist Database, redir if N/A"
```

The same debug configuration can be used for troubleshooting IGMP leave messages, as follows.

```
16 2017/07/06 14:01:34.78 CEST MINOR: DEBUG #2001 vprn2 IGMP[3]
"IGMP[3]: RX-PKT
[000 01:12:38.120] IGMP host 10.0.0.1 V3 PDU: 10.0.0.1 -> 224.0.0.22 pduLen 20
    Type: V3 REPORT maxrespCode 0x0 checksum 0x2252
    Num Group Records: 1
    Group Record 0
    Type: BLK_OLD_SRCS, AuxDataLen 0, Num Sources 1
    Mcast Addr: 239.255.1.1
    Source Address List
        192.168.4.2
"

17 2017/07/06 14:01:34.78 CEST MINOR: DEBUG #2001 vprn2 IGMP[vprn2 inst 3]
"IGMP[vprn2 inst 3]: igmpProcessGroupRec
Process group rec BLK_OLD_SRCS received on host 10.0.0.1 for group 239.255.1.1 i
n mode INCLUDE. Num srcls 1"

18 2017/07/06 14:01:34.78 CEST MINOR: DEBUG #2001 vprn2 IGMP[vprn2 inst 3]
"IGMP[vprn2 inst 3]: igmpProcessIfSrcTimerExp
Source Timer expired for IGMP host 10.0.0.1 (192.168.4.2,239.255.1.1)"

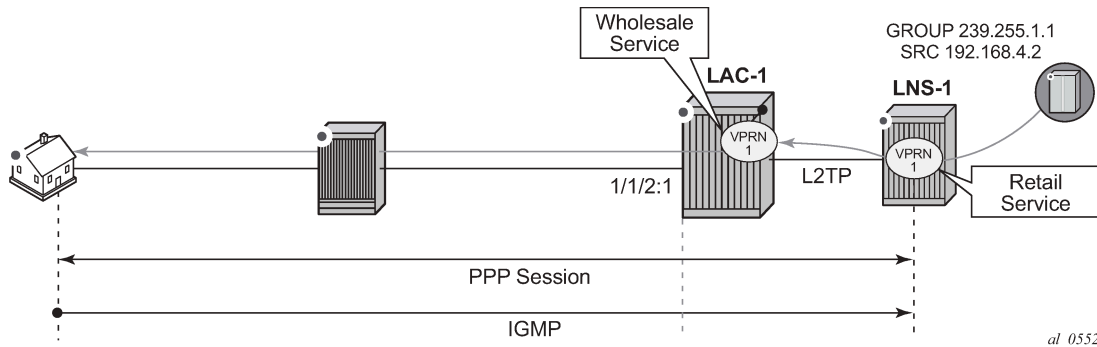
19 2017/07/06 14:01:34.78 CEST MINOR: DEBUG #2001 vprn2 IGMP[vprn2 inst 3]
"IGMP[vprn2 inst 3]: igmpIfSrcDel
Deleting i/f source entry for host 10.0.0.1 (192.168.4.2,239.255.1.1) from IGMP
Database. DeleteFromAvl: 1 !Redir 0"

20 2017/07/06 14:01:34.78 CEST MINOR: DEBUG #2001 vprn2 IGMP[vprn2 inst 3]
"IGMP[vprn2 inst 3]: igmpIfGroupDel
Deleting 239.255.1.1 from IGMP host 10.0.0.1 database"
```

ESM L2TP Wholesale/Retail Multicast

As previously mentioned, the other option for PPPoE wholesale/retail is to use an L2TP connection as shown in [Figure 42: L2TP Wholesale-Retail Multicast](#). LAC-1 contains the wholesale configuration while LNS-1 contains the retail configuration.

Figure 42: L2TP Wholesale-Retail Multicast



A partial configuration for the wholesale LAC service is as follows. It is using the local database named *ppp-retail*, defined with the **pppoe user-db** command, to authenticate subscribers. The wholesale LAC does not process any IGMP messages so it passes all messages to the retailer LNS.

```
# on LAC-1
configure
service
  vprn 1 customer 2 create
  route-distinguisher 65536:1
  ---snip---
  interface "system" create
  address 192.168.1.1/32
  loopback
  exit
  subscriber-interface "int-SUB-LAC-1" create
  unnumbered "system"
  group-interface "int-GRP-LAC-1" create
  sap 1/1/2:1 create
  sub-sla-mgmt
  def-sub-id use-sap-id
  def-sub-profile "sub-profile-1"
  def-sla-profile "sla-profile-1"
  sub-ident-policy "sub-ident-1"
  multi-sub-sap 1000
  no shutdown
  exit
  exit
  pppoe
  session-limit 10
  sap-session-limit 10
  user-db "ppp-retail"
  no shutdown
  exit
  exit
  exit
  ---snip---
  l2tp
  group "grp-retail-1" create
  tunnel "tnl-retail-1" create
  local-address 192.168.1.1
  local-name "lac-1"
  peer 192.168.1.2
  no shutdown
  exit
  no shutdown
  exit
```

```

        no shutdown
    exit
    no shutdown
exit
exit
exit
exit

```

The local user database *ppp-retail* is defined as follows:

```

# on LAC-1
configure
  subscriber-mgmt
    local-user-db "ppp-retail" create
    ppp
      match-list username
      host "retail1.com" create
      host-identification
        username "retail1.com" domain-only
      exit
      identification-strings 254 create
        sla-profile-string "sla-profile-1"
        sub-profile-string "sub-profile-1"
      exit
      l2tp
        group "grp-LAC-RETAIL-1" service-id 1
      exit
      no shutdown
    exit
  exit
  no shutdown
exit
exit
exit

```

The retailer BNG hosts the L2TP Network Server (LNS). The following is a configuration extract for the LNS. To support multicast, IGMP is enabled on the ESM group-interface in the retail service, and PIM is enabled on the interface to the multicast source.

```

# on LNS-1
configure
  service
    vprn 1 customer 2 create
      route-distinguisher 65536:1
      ---snip---
      interface "system" create
        address 192.168.1.2/32
        loopback
      exit
      interface "int-VPRN-1-MCAST" create
        address 192.168.4.1/30
        sap 1/1/6 create
        exit
      exit
      subscriber-interface "int-SUB-LNS-1" create
        address 10.1.1.254/24
        group-interface "int-GRP-LNS-1" lns create
          sap-parameters
            sub-sla-mgmt
              def-sla-profile "sla-profile-1"
              def-sub-profile "sub-profile-1"
              sub-ident-policy "sub-ident-1"
            exit
          exit
        exit
      exit
    exit
  exit

```

```

        exit
    exit
exit
igmp
    group-interface "int-GRP-LNS-1"
        no shutdown
    exit
    no shutdown
exit
l2tp
    group "lns" protocol v2 create
        lns-group 1
        ppp
            default-group-interface "int-GRP-LNS-1" service-id 1
            keepalive 90 hold-up-multiplier 5
            proxy-authentication
            proxy-lcp
            user-db "ppp-retail"
        exit
        tunnel "tunnel" create
            remote-name "lac-1"
            no shutdown
        exit
        no shutdown
    exit
    no shutdown
exit
pim
    interface "int-VRN-1-MCAST"
    exit
    no shutdown
exit
    no shutdown
exit
exit
exit

```

With the previous configuration applied, the wholesale/retail multicast setup can be verified. Firstly, send an IGMP message from the subscriber; the following example uses IGMPv3. The (S,G) sent is (192.168.4.2, 239.255.1.1) from the subscriber with IP address 10.1.1.1.

```
*A:LNS-1# show service active-subscribers igmp detail
```

```
=====
Active Subscribers Detail
=====
```

Subscriber HostAddr GrpAddr SrcAddr	IGMP-Policy GrpItf Type	Up-Time	NumGroups Mode Blk/Fwd
user1@retail1.com 10.1.1.1 239.255.1.1 192.168.4.2	igmp-policy-1 int-GRP-1 Dynamic Dynamic	0d 00:01:34	1 Include Fwd

```
-----
Number of Subscribers : 1
=====
```

```
*A:LNS-1#
```

The IGMP group is not seen in the wholesale router instance (as shown by the following output on LAC-1), however, it is seen in the retail router instance (as shown by the second output below on LNS-1).

```
*A:LAC-1# show router 1 igmp group
=====
IGMP Interface Groups
=====
No Matching Entries
=====
IGMP Host Groups
=====
No Matching Entries
=====
IGMP SAP Groups
=====
No Matching Entries
=====
*A:LAC-1#
```

```
*A:LNS-1# show router 1 igmp group
=====
IGMP Interface Groups
=====
No Matching Entries
=====
IGMP Host Groups
=====
(192.168.4.2,239.255.1.1)
  Fwd List   : 10.1.1.1                                UpTime: 0d 00:02:58
-----
Entries : 1
=====
IGMP SAP Groups
=====
No Matching Entries
=====
*A:LNS-1#
```

Only the retail BNG (LNS-1) is responsible for processing the IGMP messages. Therefore, to troubleshoot ESM multicast for an L2TP service, the following debug commands are used on the LNS.

```
debug
  router "1"
    igmp
      group-interface "int-GRP-1"
      host "10.1.1.1"
      packet mode egr-ingr-and-dropped
    exit
  exit
exit
```

```
56 2017/07/05 15:21:09.51 CEST MINOR: DEBUG #2001 vprn1 IGMP[2]
"IGMP[2]: RX-PKT
[000 00:26:16.790] IGMP host 10.1.1.1 V3 PDU: 10.1.1.1 -> 224.0.0.22 pduLen 20
  Type: V3 REPORT maxrespCode 0x0 checksum 0x2352
  Num Group Records: 1
    Group Record 0
      Type: ALW_NEW_SRCS, AuxDataLen 0, Num Sources 1
      Mcast Addr: 239.255.1.1
      Source Address List
        192.168.4.2
```



```
"  
  
57 2017/07/05 15:21:09.51 CEST MINOR: DEBUG #2001 vprn1 IGMP[vprn1 inst 2]  
"IGMP[vprn1 inst 2]: igmpIfGroupAdd  
Adding 239.255.1.1 to IGMP host 10.1.1.1 database"  
  
58 2017/07/05 15:21:09.51 CEST MINOR: DEBUG #2001 vprn1 IGMP[vprn1 inst 2]  
"IGMP[vprn1 inst 2]: igmpProcessGroupRec  
Process group rec ALW_NEW_SRCS received on host 10.1.1.1 for group 239.255.1.1 i  
n mode INCLUDE. Num srcs 1"  
  
59 2017/07/05 15:21:09.51 CEST MINOR: DEBUG #2001 vprn1 IGMP[vprn1 inst 2]  
"IGMP[vprn1 inst 2]: igmpIfSrcAdd  
Adding i/f source entry for host 10.1.1.1 (192.168.4.2,239.255.1.1) to IGMP fwdL  
ist Database, redir if N/A"
```

The IGMP leave messages can also be seen in the debug, as follows.

```
1 2017/07/05 16:41:03.05 CEST MINOR: DEBUG #2001 vprn1 IGMP[2]  
"IGMP[2]: RX-PKT  
[000 00:24:30.330] IGMP host 10.1.1.1 V3 PDU: 10.1.1.1 -> 224.0.0.22 pduLen 20  
Type: V3 REPORT maxrespCode 0x0 checksum 0x2252  
Num Group Records: 1  
Group Record 0  
Type: BLK_OLD_SRCS, AuxDataLen 0, Num Sources 1  
Mcast Addr: 239.255.1.1  
Source Address List  
192.168.4.2  
"  
  
2 2017/07/05 16:41:03.06 CEST MINOR: DEBUG #2001 vprn1 IGMP[vprn1 inst 2]  
"IGMP[vprn1 inst 2]: igmpProcessGroupRec  
Process group rec BLK_OLD_SRCS received on host 10.1.1.1 for group 239.255.1.1 i  
n mode INCLUDE. Num srcs 1"  
  
3 2017/07/05 16:41:03.06 CEST MINOR: DEBUG #2001 vprn1 IGMP[vprn1 inst 2]  
"IGMP[vprn1 inst 2]: igmpProcessIfSrcTimerExp  
Source Timer expired for IGMP host 10.1.1.1 (192.168.4.2,239.255.1.1)"  
  
4 2017/07/05 16:41:03.06 CEST MINOR: DEBUG #2001 vprn1 IGMP[vprn1 inst 2]  
"IGMP[vprn1 inst 2]: igmpIfSrcDel  
Deleting i/f source entry for host 10.1.1.1 (192.168.4.2,239.255.1.1) from IGMP  
Database. DeleteFromAvl: 1 !Redir 0"
```

Conclusion

Multicast is an essential part of Triple Play Services. The SR/ESS TPSDA solution is much more than a baseline multicast delivery. It includes individual subscriber awareness and provides each retailer a separate routing context for managing their own multicast content. Subscriber awareness allows for the fine-tuning of each subscriber multicast experience and also for troubleshooting on a per subscriber basis. This example provides a complete configuration walkthrough for multicast delivery for both IPoE and PPPoE in a wholesale/retail model.

ESM IPv4: Multicast with Redirection

This chapter describes ESM IPv4 multicast with redirection configurations.

Topics in this chapter include:

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

Applicability

The configuration in this chapter was initially based on SR OS Release 11.0.R1 and covers both IPoE and PPPoE subscribers. The CLI is updated to Release 15.0.R3.

Overview


Triple Play Service Delivery Architecture (TPSDA) allows operators to integrate High Speed Internet (HSI), voice, and video services within a single network. The goal of this chapter is to walk through a TPSDA multicast architecture with redirection. The topics are divided into the following sections:

- Single BNG with redirection
- SRRP BNG configuration with Multi Chassis Link Aggregation (MC-LAG) and redirection
- IPoE ESM multicast walkthrough
- PPPoE ESM multicast walkthrough
- Subscriber Routed Redundancy Protocol (SRRP)
 - Multi-Chassis Synchronization (MCS) walkthrough
- IGMP Debugging

In [ESM IPv4: Multicast with SRRP](#), multicast is directly distributed to a subscriber through a subscriber SAP. This chapter walks through another popular model which redirects all multicast streams to a common routed interface for all subscribers. When multicast is put on the common routed interface, one single copy of a multicast stream is delivered to multiple subscribers. In this model, per-subscriber replication of multicast streams is done on an access node or on the aggregation network in order to minimize the bandwidth consumed by the multicast traffic in access/aggregation.

[Figure 43: Network Topology Overview](#) shows two BNGs configured with SRRP to provide redundancy. The P router is connected to the multicast source and is connected to both BNGs. The connections between the BNGs and the P router, and the multicast source and the P router, are running PIM to provide multicast delivery. On the access side, the two BNGs are connected to an aggregation switch via MC-LAG aggregating the traffic for both PPPoE and IPoE subscribers. The BNGs facing the subscriber side are IGMP aware and will respond to any subscribers' IGMP requests.

Figure 43: Network Topology Overview



There are two requirements for a subscriber to receive multicast streams. First, the ESM group-interface must have IGMP enabled. Second, each subscriber's subscriber profile must be customized to allow them to receive multicast streams. When both requirements are met, the BNG will process the subscribers' IGMP messages, otherwise, IGMP messages are dropped. All customer premise equipment (CPE) IGMP messages are aggregated via the 7450 and passed onto the BNGs. Because the BNGs are running SRRP, the SRRP master is the only BNG processing and answering the IGMP messages. Protocol Independent Multicast (PIM) is used between the BNG and the P router to request the multicast streams. If PIM is successful in retrieving the multicast group, the multicast stream is forwarded toward the individual subscribers. This is the typical multicast delivery for TPSDA.

Configuration

This example builds on the ESM multicast foundation described in chapter [ESM IPv4: Multicast with SRRP](#). It starts with a single BNG setup with redirection.


ESM Multicast Interface Redirection

[Figure 44: Single BNG Setup with Multicast Redirection](#) shows a popular ESM multicast model that redirects all multicast streams to a dedicated router interface. When configuring a redirected interface be aware that:

1. Redirection between Global Routing Table (GRT) interfaces and VPRN interfaces is not supported
2. GRT interfaces are interfaces that reside in the base router or in an IES.
3. Redirection can be performed between interfaces in the GRT or between the interfaces in any VPRN (even different VPRNs).

The following scenario start with a simple ESM multicast configuration for BNG, without redirection.

Figure 44: Single BNG Setup with Multicast Redirection



al_0271

The BNG-1 baseline configuration excluding multicast redirection is as follows. The local DHCP server is hosted on BNG-1, and is bound to interface *int-DHCP-lb1*.

```
configure
router Base
  dhcp
    local-dhcp-server "dhcp-local-server" create
    use-gi-address scope pool
    pool "pool-1" create
    subnet 10.0.0.0/8 create
    options
      subnet-mask 255.0.0.0
      default-router 10.255.255.254
    exit
    address-range 10.0.0.10 10.0.0.254
  exit
  exit
  no shutdown
exit
interface "int-DHCP-lb1"
  address 192.168.0.1/32
  loopback
  local-dhcp-server "dhcp-local-server"
  no shutdown
exit
interface "system"
  address 192.0.2.1/32
  no shutdown
exit
exit
exit
```

Subscribers are located in the 10.0.0.0/8 subnet. The multicast stream (S,G) is (192.168.4.2, 239.255.1.1).

```
configure
service
  ies 1 customer 1 create
  description "BNG-1"
  interface "int-BNG-1-MCAST-S1" create
    address 192.168.4.1/30
    sap 1/1/5 create
  exit
exit
subscriber-interface "sub-int-1" create
```

```
        address 10.255.255.254/8
        group-interface "grp-int-1" create
        srrp-enabled-routing
        dhcp
            server 192.168.0.1
            trusted
            lease-populate 10
            client-applications dhcp ppp
            gi-address 10.255.255.254
            no shutdown
        exit
        authentication-policy "auth-1"
        sap 1/1/4:1 create
            sub-sla-mgmt
                def-sub-profile "sub-profile-1"
                def-sla-profile "sla-profile-1"
                sub-ident-policy "sub-ident-1"
                multi-sub-sap 10
                no shutdown
            exit
        exit
        pppoe
            no shutdown
        exit
    exit
    no shutdown
exit
exit
exit
```

To support multicast, group-interface *grp-int-1* is IGMP enabled, and interface *int-BNG-1-MCAST-S1* is PIM enabled.

```
configure
router Base
    igmp
        group-interface "grp-int-1"
        no shutdown
    exit
    no shutdown
exit
    pim
        interface "int-BNG-1-MCAST-S1"
        exit
        no shutdown
    exit
exit
exit
```

The subscriber management profiles enabling multicast for ESM subscribers are defined as follows:

```
configure
subscriber-mgmt
    igmp-policy "igmp-policy-1" create
    exit
    sub-profile "sub-profile-1" create
        igmp-policy "igmp-policy-1"
    exit
exit
exit
```

A router interface is defined to redirect all multicast streams to, and added to the igmp context, as follows.

```
configure
  service
    ies 1
      interface "int-REDIRECT" create
        address 192.168.10.1/30
        sap 1/1/2:1 create
        exit
      exit
    exit
  exit
exit
```

```
configure
  router
    igmp
      interface "int-REDIRECT"
        no shutdown
      exit
      no shutdown
    exit
  exit
exit
```

A router redirection policy is defined, redirecting every (S,G) toward the redirected interface, as follows:

```
configure
  router
    policy-options
      begin
      policy-statement "mcast-REDIR-IF"
        default-action accept
        multicast-redirection fwd-service 1 "int-REDIRECT"
      exit
    exit
  commit
  exit
exit
exit
```

The redirection policy then is applied to the igmp policy, which in turn is applied to the subscriber profile, as follows:

```
configure
  subscriber-mgmt
    igmp-policy "igmp-policy-1"
      redirection-policy "mcast-REDIR-IF"
    exit
  sub-profile "sub-profile-1" create
    igmp-policy "igmp-policy-1"
  exit
exit
exit
```

From this point on all multicast streams will be redirected to the *int-REDIRECT* interface.

Now send an IGMPv3 join message and then use the **show router igmp group** command to verify that all multicast streams are redirected. In this example, IGMPv3 is used with an (S,G) of (192.168.4.2, 239.255.1.1).

For PPPoE subscribers, the multicast (S,G) shows up on both the redirected interface and the host, as follows:

```
*A:BNG-1# show router igmp group
=====
IGMP Interface Groups
=====
(192.168.4.2,239.255.1.1)                UpTime: 0d 00:00:30
  Fwd List  : int-REDIRECT
-----
Entries : 1
=====
IGMP Host Groups
=====
(192.168.4.2,239.255.1.1)                UpTime: 0d 00:00:30
  Fwd List  : 10.0.0.10
-----
Entries : 1
=====
IGMP SAP Groups
=====
No Matching Entries
=====
*A:BNG-1#
```

For IPoE subscribers, the multicast (S,G) shows up on both the redirected interface and the SAP.

```
*A:BNG-1# show router igmp group
=====
IGMP Interface Groups
=====
(192.168.4.2,239.255.1.1)                UpTime: 0d 00:01:40
  Fwd List  : int-REDIRECT
-----
Entries : 1
=====
IGMP Host Groups
=====
(192.168.4.2,239.255.1.1)                UpTime: 0d 00:01:40
  Fwd List  : 10.0.0.12
-----
Entries : 1
=====
IGMP SAP Groups
=====
No Matching Entries
=====
*A:BNG-1#
```

Now the *int-REDIRECT* interface is the only interface sending out multicast streams. The first command shows that the group interface does not register any multicast group (Num-Groups=0). The second command shows that all multicast groups are registered against the redirected interface (Num-Groups=1).

```
*A:BNG-1# show router igmp group-interface
=====
IGMP Group-Interfaces
=====
FwdSvc Group-Interface                Adm/Opr-State                Import-Policy
```

```

      SAP                      Adm/Opr-Version      Num-Groups
-----
1      grp-int-1              Up/Up                 none
      1/1/4:1                 3/3                 0
-----
Group-Interfaces = 1, SAPs = 1
=====
*A:BNG-1#

*A:BNG-1# show router igmp interface

=====
IGMP Interfaces
=====
Interface          Adm  Oper Querier          Cfg/Opr Num      Policy
                   Version Groups
-----
int-REDIRECT        Up   Up   192.168.10.1       3/3     1       none
-----
Interfaces : 1
=====
*A:BNG-1#

```

Debug facilities can be used to troubleshoot multicast redirection issues. The following output shows the multicast stream is redirected to a regular routed interface after an IGMP join.

```

74 2017/06/29 09:52:21.07 CEST MINOR: DEBUG #2001 Base IGMP[1]
"IGMP[1]: RX-PKT
[000 19:43:29.700] IGMP host 10.0.0.12 V3 PDU: 10.0.0.12 -> 224.0.0.22 pduLen 20
  Type: V3 REPORT maxrespCode 0x0 checksum 0x2352
  Num Group Records: 1
    Group Record 0
      Type: ALW_NEW_SRCS, AuxDataLen 0, Num Sources 1
      Mcast Addr: 239.255.1.1
      Source Address List
        192.168.4.2
"

75 2017/06/29 09:52:21.07 CEST MINOR: DEBUG #2001 Base IGMP[Base inst 1]
"IGMP[Base inst 1]: igmpIfGroupAdd
Adding 239.255.1.1 to IGMP host 10.0.0.12 database"

76 2017/06/29 09:52:21.07 CEST MINOR: DEBUG #2001 Base IGMP[Base inst 1]
"IGMP[Base inst 1]: igmpProcessGroupRec
Process group rec ALW_NEW_SRCS received on host 10.0.0.12 for group 239.255.1.1
in mode INCLUDE. Num srcs 1"

77 2017/06/29 09:52:21.07 CEST MINOR: DEBUG #2001 Base IGMP[Base inst 1]
"IGMP[Base inst 1]: igmpIfSrcAdd
Adding i/f source entry for host 10.0.0.12 (192.168.4.2,239.255.1.1) to IGMP fwd
List Database, redir if interface int-REDIRECT [ifIndex 6]"

```

The following output shows what happens when an IGMP leave message is sent so that the multicast stream is no longer being forwarded.

```

170 2017/06/29 10:05:22.07 CEST MINOR: DEBUG #2001 Base IGMP[1]
"IGMP[1]: RX-PKT
[000 19:56:30.700] IGMP host 10.0.0.12 V3 PDU: 10.0.0.12 -> 224.0.0.22 pduLen 20
  Type: V3 REPORT maxrespCode 0x0 checksum 0x2252
  Num Group Records: 1
    Group Record 0

```



```
Type: BLK_OLD_SRCS, AuxDataLen 0, Num Sources 1
Mcast Addr: 239.255.1.1
Source Address List
    192.168.4.2
"

171 2017/06/29 10:05:22.07 CEST MINOR: DEBUG #2001 Base IGMP[Base inst 1]
"IGMP[Base inst 1]: igmpProcessGroupRec
Process group rec BLK_OLD_SRCS received on host 10.0.0.12 for group 239.255.1.1
in mode INCLUDE. Num srcs 1"

172 2017/06/29 10:05:22.07 CEST MINOR: DEBUG #2001 Base IGMP[Base inst 1]
"IGMP[Base inst 1]: igmpProcessIfSrcTimerExp
Source Timer expired for IGMP host 10.0.0.12 (192.168.4.2,239.255.1.1)"

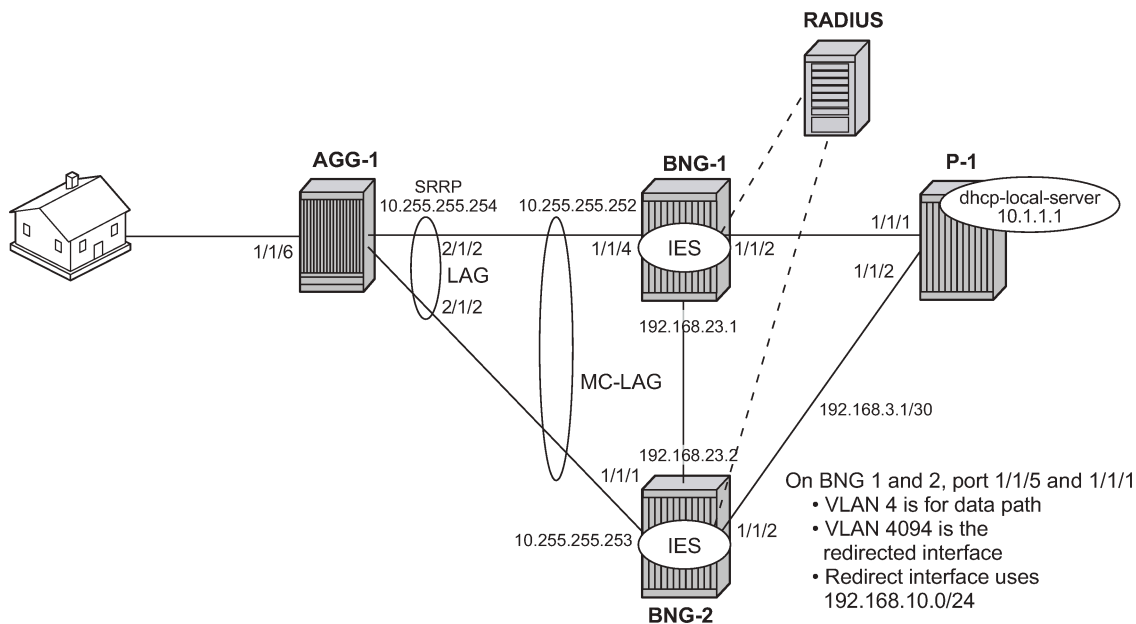
173 2017/06/29 10:05:22.07 CEST MINOR: DEBUG #2001 Base IGMP[Base inst 1]
"IGMP[Base inst 1]: igmpIfSrcDel
Deleting i/f source entry for host 10.0.0.12 (192.168.4.2,239.255.1.1) from IGMP
Database. DeleteFromAvl: 1 Redir 0"

174 2017/06/29 10:05:22.07 CEST MINOR: DEBUG #2001 Base IGMP[Base inst 1]
"IGMP[Base inst 1]: igmpIfGroupDel
Deleting 239.255.1.1 from IGMP host 10.0.0.12 database"
```

ESM Multicast Redirection with SRRP and MC-LAG

[Figure 45: Network Topology with MC-LAG](#) shows a numbered SRRP setup with MC-LAG SAPs serving both IPoE and PPPoE subscribers. [ESM IPv4: Multicast with SRRP](#) covers the configuration of regular SRRP SAPs, consequently this example provides configuration guidelines to use a different type of SAP: SRRP MC-LAG SAPs. Note that redirection on SRRP SAPs without MC-LAG is also supported. The configuration of the RADIUS server is out of the scope of this example.

Figure 45: Network Topology with MC-LAG



The baseline configuration for BNG-1 excluding the IGMP configuration is as follows. The configuration begins with the MC-LAG configuration. ESM is configured in an IES service, but it is also possible to configure ESM in a VPRN. The redirection interface must be in the same routing instance as the group-interface, this applies to both regular SRRP SAPs and MC-LAG SAPs. In the following example, the MC-LAG is **lag-1**, customer data traffic is using VLAN 4, MC-LAG control traffic is using VLAN 5, and the redirected multicast streams are using VLAN 4094.

```
# Baseline for BNG-1
configure
  lag 1
    mode access
    encap-type dot1q
    port 1/1/4 priority 1
    lacp active administrative-key 32768
    no shutdown
  exit
exit
```

```
configure
  redundancy
    multi-chassis
      peer 192.0.2.3 create
      mc-lag
        lag 1 lacp-key 1 system-id 00:00:00:00:00:01 system-priority 100
        no shutdown
      exit
    sync
    igmp
    srrp
    sub-mgmt ipoe pppoe
    port lag-1 create
```

```
        range 4-4 sync-tag "mclagdata"
        range 5-5 sync-tag "mclagcontrol"
    exit
    no shutdown
exit
no shutdown
exit
exit
exit
exit
exit

configure
service
  ies 1 customer 1 create
  description "BNG-1"
  redundant-interface "mclink-BNG-1-BNG-2" create
  address 192.168.11.0/31
  ip-mtu 1500
  spoke-sdp 23:1 create
  no shutdown
  exit
exit
interface "int-LAG-REDIRECTED" create
  address 192.168.10.252/24
  vrrp 1
  backup 192.168.10.254
  exit
  sap lag-1:4094 create
  exit
exit
subscriber-interface "sub-int-1" create
  address 10.255.255.252/8 gw-ip-address 10.255.255.254 track-srrp 1
  group-interface "grp-int-1" create
  srrp-enabled-routing
  dhcp
  server 192.168.0.1
  lease-populate 10
  client-applications dhcp ppp
  gi-address 10.255.255.252
  no shutdown
  exit
  authentication-policy "auth-1"
  redundant-interface "mclink-BNG-1-BNG-2"
  sap lag-1:4 create
  sub-sla-mgmt
  def-sub-id use-sap-id
  def-sub-profile "sub-profile-1"
  def-sla-profile "sla-profile-1"
  sub-ident-policy "sub-ident-1"
  multi-sub-sap 10
  no shutdown
  exit
exit
sap lag-1:5 create
exit
srrp 1 create
  message-path lag-1:5
  priority 200
  no preempt
  no shutdown
exit
pppoe
  no shutdown
```

```
        exit
    exit
    exit
    no shutdown
    exit
    exit
    exit
exit

configure
router Base
    interface "int-BNG-1-BNG-2"
        address 192.168.23.1/30
        port 1/1/3:4040
        no shutdown
    exit
    interface "int-BNG-1-P-1"
        address 192.168.24.1/30
        port 1/1/2:4040
        no shutdown
    exit
    interface "system"
        address 192.0.2.2/32
        bfd 100 receive 100 multiplier 3
        no shutdown
    exit
    ospf 0
        traffic-engineering
        area 0.0.0.0
            interface "system"
                no shutdown
            exit
            interface "int-BNG-1-BNG-2"
                interface-type point-to-point
                metric 10000
                no shutdown
            exit
            interface "int-BNG-1-P-1"
                interface-type point-to-point
                no shutdown
            exit
            interface "sub-int-1"
                no shutdown
            exit
            interface "int-LAG-REDIRECTED"
                no shutdown
            exit
        exit
        no shutdown
    exit
    pim
        interface "int-BNG-1-P-1"
        exit
    exit
    exit
exit
exit
```

The baseline configuration for BNG-2 excluding the IGMP configuration is as follows. The default SRRP priority for BNG-2 is lower than the SRRP priority for BNG-1 and therefore BNG-2 will be in standby mode.

```
# Baseline for BNG-2
configure
    lag 1
```

```
mode access
encap-type dot1q
port 1/1/1 priority 1
lacp active administrative-key 32768
no shutdown
exit
exit
```

```
configure
  redundancy
    multi-chassis
      peer 192.0.2.2 create
      mc-lag
        lag 1 lacp-key 1 system-id 00:00:00:00:00:01 system-priority 100
        no shutdown
      exit
      sync
        igmp
        srrp
        sub-mgmt ipoe pppoe
        port lag-1 create
          range 4-4 sync-tag "mclagdata"
          range 5-5 sync-tag "mclagcontrol"
        exit
        no shutdown
      exit
      no shutdown
    exit
  exit
exit
exit
exit
```

```
configure
  service
    ies 1 customer 1 create
      description "BNG-2, SRRP1"
      redundant-interface "mclink-BNG-2-BNG-1" create
        address 192.168.11.1/31
        ip-mtu 1500
        spoke-sdp 32:1 create
          no shutdown
        exit
      exit
      interface "int-LAG-REDIRECTED" create
        address 192.168.10.253/24
        vrrp 2
          backup 192.168.10.254
        exit
        sap lag-1:4094 create
        exit
      exit
      subscriber-interface "sub-int-1" create
        address 10.255.255.253/8 gw-ip-address 10.255.255.254 track-srrp 1
        group-interface "grp-int-1" create
          srrp-enabled-routing
          dhcp
            server 192.168.0.1
            lease-populate 10
            client-applications dhcp ppp
            gi-address 10.255.255.253
            no shutdown
          exit
        exit
      exit
```

```
        authentication-policy "auth-1"
        redundant-interface "mclink-BNG-2-BNG-1"
        sap lag-1:4 create
            sub-sla-mgmt
                def-sub-id use-sap-id
                def-sub-profile "sub-profile-1"
                def-sla-profile "sla-profile-1"
                sub-ident-policy "sub-ident-1"
                multi-sub-sap 10
                no shutdown
            exit
        exit
        sap lag-1:5 create
        exit
        srrp 1 create
            message-path lag-1:5
            priority 150
            no shutdown
        exit
        pppoe
            no shutdown
        exit
    exit
exit
exit
exit
exit
```

```
configure
router Base
    interface "int-BNG-2-BNG-1"
        address 192.168.23.2/30
        port 1/1/3:4040
        no shutdown
    exit
    interface "int-BNG-2-P-1"
        address 192.168.34.1/30
        port 1/1/2:4040
        no shutdown
    exit
    interface "system"
        address 192.0.2.3/32
        bfd 100 receive 100 multiplier 3
        no shutdown
    exit
    ospf 0
        traffic-engineering
        area 0.0.0.0
            interface "system"
                no shutdown
            exit
            interface "int-BNG-2-BNG-1"
                interface-type point-to-point
                metric 10000
                no shutdown
            exit
            interface "int-BNG-2-P-1"
                interface-type point-to-point
                no shutdown
            exit
            interface "sub-int-1"
                no shutdown
```

```
        exit
        interface "int-LAG-REDIRECTED"
            no shutdown
        exit
    exit
    no shutdown
exit
igmp
    no shutdown
exit
pim
    interface "int-BNG-2-P-1"
        exit
        no shutdown
    exit
exit
exit
exit
```

The baseline configuration for the 7450 aggregation switch is as follows. It has a LAG interface configured. There are two VPLSs. The first is VPLS 1 which is used to receive all redirected multicast traffic on VLAN 4094. The second is VPLS 2 which is responsible for passing all subscriber traffic on VLAN 4.

```
# Baseline for AGG-1
configure
    lag 1
        mode access
        encap-type dot1q
        port 2/1/1
        port 2/1/2
        lacp active administrative-key 32768
        no shutdown
    exit
exit
```

```
configure
    service
        vpls 4 customer 1 create
            description "for user traffic"
            stp
                shutdown
            exit
            sap 1/1/6:4 create
                description "to spirent"
                no shutdown
            exit
            sap lag-1:4 create
                no shutdown
            exit
            no shutdown
        exit
    exit
exit
```

```
configure
    service
        vpls 4094 customer 1 create
            description "for the multicast streams"
            stp
                shutdown
            exit
            sap lag-1:4094 create
```

```

        no shutdown
    exit
    sap 1/1/5:4094 create
        no shutdown
    exit
    no shutdown
exit
exit
exit

```

The baseline configuration for the P router is as follows. It is now responsible for DHCP address assignment (moved from BNG-1 in the previous configuration to allow for redundant operations in case of failure of either BNG-1 or BNG-2) and is also attached to the multicast source.

```

# Baseline for P-1
configure
router
dhcp
    local-dhcp-server "dhcp-local-server" create
        use-gi-address scope pool
        pool "pool-1" create
            subnet 10.0.0.0/8 create
                options
                    subnet-mask 255.0.0.0
                    default-router 10.255.255.254
                exit
            address-range 10.0.0.10 10.0.0.254
        exit
    exit
    no shutdown
exit
interface "int-DHCP-lb1"
    address 192.168.0.1/32
    loopback
    local-dhcp-server "dhcp-local-server"
    no shutdown
exit
interface "int-P-1-BNG-1"
    address 192.168.24.2/30
    port 1/1/1:4040
    no shutdown
exit
interface "int-P-1-BNG-2"
    address 192.168.34.2/30
    port 1/1/2:4040
    no shutdown
exit
interface "int-P-1-mcast-source"
    address 192.168.4.1/30
    port 1/1/5:4092
    no shutdown
exit
interface "system"
    address 192.0.2.4/32
    no shutdown
exit
ospf 0
    traffic-engineering
    area 0.0.0.0
        interface "system"
            no shutdown
        exit

```



```
        interface "int-DHCP-lb1"
            no shutdown
        exit
        interface "int-P-1-BNG-1"
            interface-type point-to-point
            no shutdown
        exit
        interface "int-P-1-BNG-2"
            interface-type point-to-point
            no shutdown
        exit
        interface "int-P-1-mcast-source"
            passive
            no shutdown
        exit
    exit
    no shutdown
exit
pim
    interface "int-P-1-BNG-1"
    exit
    interface "int-P-1-BNG-2"
    exit
    interface "int-P-1-mcast-source"
    exit
    ---snip---
    no shutdown
    exit
exit
exit
```

Enable IGMP on Group Interface and Redirect Interface

The following configuration shows how to add the group-interface and redirect interface to IGMP. If ESM is configured in a VPRN, each VPRN will have its own IGMP instance. Remember to apply the following configuration to both BNG-1 and BNG-2.

```
# on BNG-1 and BNG-2
configure
    router
        igmp
            group-interface "grp-int-1"
            no shutdown
        exit
        interface "int-LAG-REDIRECTED"
            no shutdown
        exit
        no shutdown
    exit
exit
exit
```

Next, the IGMP policy is configured to redirect all multicast streams to a dedicated interface. The following configuration outlines the steps necessary to enable multicast redirection.

Define a router redirection policy. This will redirect every (S,G) toward the redirected interface.

```
configure
    router
        policy-options
```

```
begin
  policy-statement "mcast-REDIR-IF"
    default-action accept
    multicast-redirection fwd-service 1 "int-LAG-REDIRECTED"
  exit
exit
commit
exit
exit
exit
```

Apply the redirection policy to the IGMP policy, as follows:

```
configure
  subscriber-mgmt
    igmp-policy "igmp-policy-1" create
    redirection-policy "mcast-REDIR-IF"
  exit
exit
exit
```

Add multi-chassis synchronization for the redirected interface on BNG-1. This will synchronize the IGMP state on this MC-LAG interface. A similar configuration is required on BNG-2.

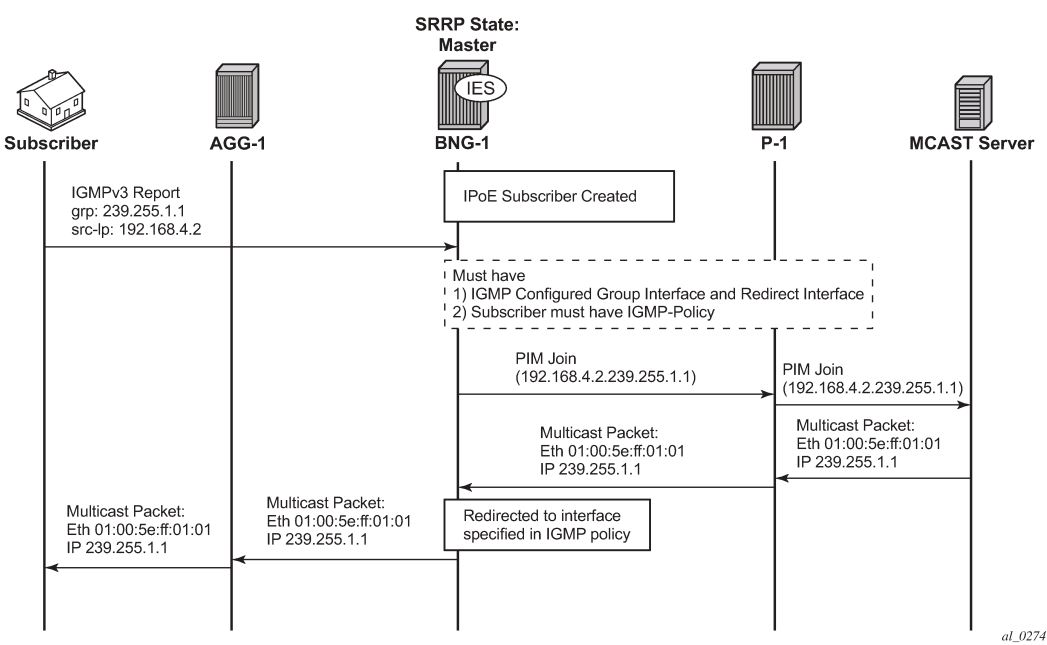
```
# on BNG-1
configure
  redundancy
    multi-chassis
      peer 192.0.2.3 create
      sync
        igmp
        srrp
        sub-mgmt ipoe pppoe
        port lag-1 create
          range 4-4 sync-tag "mclagdata"
          range 5-5 sync-tag "mclagcontrol"
          range 4094-4094 sync-tag "mclagmulticast"
        exit
      no shutdown
    exit
  no shutdown
exit
exit
exit
exit
exit
```

ESM IGMP IPoE walkthrough

With the baseline configuration applied, the BNG is ready to process IGMP messages and deliver multicast streams to the subscribers through the redirected interface. Figure 4 shows the message flow for IPoE subscribers requesting and receiving multicast traffic.

- The group-interface and redirect interface must have IGMP enabled.
- The subscriber must be associated with an IGMP-policy via a sub-profile.

Figure 46: IPoE Multicast Message Flow



To verify the (ESM enabled) group-interface and the redirect interface are ready for multicast, use the following show commands. Remember the IES service ID is 1, the group-interface name is *grp-int-1* and the interface name is *int-LAG-REDIRECTED*.

Verify if the group interface and redirected interface have IGMP enabled.

```
*A:BNG-1# show router igmp group-interface
=====
IGMP Group-Interfaces
=====
FwdSvc Group-Interface      Adm/Opr-State      Import-Policy
  SAP      Adm/Opr-Version      Num-Groups
-----
1      grp-int-1      Up/Up      none
      lag-1:4      3/3      0
      lag-1:5      3/3      0
-----
Group-Interfaces = 1, SAPs = 2
=====
*A:BNG-1#

*A:BNG-1# show router igmp interface
=====
IGMP Interfaces
=====
Interface      Adm  Oper  Querier      Cfg/Opr Num      Policy
              Version Groups
-----
int-LAG-REDIRECTED  Up   Up    192.168.10.253  3/3    0      none
-----
Interfaces : 1
```

```
=====
*A:BNG-1#
```

Ensure the subscriber is associated with an IGMP-policy. Because the IGMP-policy is associated with a subscriber-profile, verification of an IGMP-policy is performed via the sub-profile.

```
*A:BNG-1# show subscriber-mgmt sub-profile "sub-profile-1"

=====
Subscriber Profile sub-profile-1
=====
Description      : (Not Specified)
I. Sched. Policy : N/A
E. Sched. Policy : N/A                      E. Agg Rate Limit: Max
I. Policer Ctrl. : N/A
E. Policer Ctrl. : N/A
I. vport-hashing : Disabled
I. sec-sh-hashing: Disabled
Q Frame-Based Ac*: Disabled
Acct. Policy     : N/A                      Collect Stats    : Disabled
ANCP Pol.       : N/A
Accu-stats-pol  : (Not Specified)
HostTrk Pol.    : N/A
IGMP Policy     : igmp-policy-1
MLD Policy      : N/A
PIM Policy      : N/A
Sub. MCAC Policy: N/A
NAT Policy      : N/A
Firewall Policy : N/A
UPnP Policy     : N/A
NAT Prefix List : N/A
Def. Encap Offset: none                    Encap Offset Mode: none
Avg Frame Size  : N/A
Vol stats type  : full
Preference      : 5
LAG hash class  : 1
LAG hash weight : 1
-----
Radius Accounting
-----
Policy          : N/A
Session Opti.Stop: False
-----
HSMMDA-2
-----
I. Qos Policy   : 1                      E. Qos Policy   : 1
E. WRR Policy   : N/A                    E. Agg Rate Limit: Max
Pkt Byte Offset : add 0*
-----
Last Mgmt Change : 06/27/2017 12:01:27
=====
* indicates that the corresponding row element may have been truncated.
---snip---
*A:BNG-1#
```

After the verification, the BNGs are ready to deliver multicast streams. Next, initiate an IGMP report from a subscriber requesting a multicast channel. In this example, IGMPv3 with SSM is used. If the IPoE subscriber is receiving multicast through the subscriber SAP then the IGMP group will be associated with the SAP. Because redirection is used, the IGMP group is associated with the redirected interface instead. The following output shows that when an IGMP message is received and processed, an (S,G) binding

is associated with the redirected interface. The example uses an IGMPv3 SSM message requesting (192.168.4.2, 239.255.1.1). The subscriber IP address is 10.0.0.11.

```
*A:BNG-1# show router igmp group
=====
IGMP Interface Groups
=====
(192.168.4.2,239.255.1.1)                               UpTime: 0d 00:01:16
  Fwd List  : int-LAG-REDIRECTED
-----
Entries : 1
=====
IGMP Host Groups
=====
(192.168.4.2,239.255.1.1)                               UpTime: 0d 00:01:16
  Fwd List  : 10.0.0.11
-----
Entries : 1
=====
IGMP SAP Groups
=====
No Matching Entries
=====
*A:BNG-1#
```

Next, verify the individual subscribers and their IGMP information. First verify the IGMP policy related to the subscriber.

```
*A:BNG-1# show service active-subscribers igmp detail
=====
Active Subscribers Detail
=====
Subscriber
HostAddr      IGMP-Policy
GrpAddr      GrpItf
SrcAddr      Type          Up-Time          NumGroups
              Type          Mode
              Type          Blk/Fwd
-----
ipoe-sub-1    igmp-policy-1
10.0.0.11     grp-int-1
239.255.1.1   Dynamic        0d 00:03:30      1
192.168.4.2   Dynamic        Include
              Fwd
-----
Number of Subscribers : 1
=====
*A:BNG-1#
```

Because the IGMP-policy controls bandwidth, interoperability, and restricts multicast groups, it is useful to view what is defined in the IGMP-policy if the subscriber fails to receive multicast streams.

```
*A:BNG-1# show subscriber-mgmt igmp-policy "igmp-policy-1"
=====
IGMP Policy igmp-policy-1
=====
Import Policy      :
Admin Version      : 3
Num Subscribers    : 1
Host Max Group     : No Limit
Host Max Sources   : No Limit
Host Max Group Sources : No Limit
```

```
Query Interval           : None
Query Last Member Interval : None
Query Response Interval  : None
Router Alert Check       : Enabled
Fast Leave               : yes
Redirection Policy        : mcast-REDIR-IF
Per Host Replication      : no
Egress Rate Modify       : no
Mcast Reporting Destination Name :
Mcast Reporting Admin State : Disabled
=====
```

*A:BNG-1#

The following command lists the (S,G)s that all subscribers are requesting. The operational status for the host is Up. The multicast stream is not delivered directly over the subscriber SAP, but over the *int-LAG-REDIRECTED* interface instead.

```
*A:BNG-1# show router igmp hosts detail
```

```
=====
IGMP Host 10.0.0.11
=====
```

```
Oper Status      : Up      MacAddress       : 00:10:94:00:00:02
Oper version     : 3        Subscriber          : ipoe-sub-1
Num Groups       : 1        GrpItf             : grp-int-1
Max Grps Till Now: 1        IGMP-Policy        : igmp-policy-1
PPPoE SessionId  : N/A      Next query time: 0d 00:00:32
FwdSvcId         : 1        Max Srcs Allow*: No Limit
Max Grps Allowed : No Limit Max Grp Srcs A*: No Limit
Qry Interval     : 125      Qry Last Mbr I*: 1
Qry Resp Interval: 10      Router Alert C*: Enabled
```

```
-----
IGMP Group
-----
```

```
Group Address    : 239.255.1.1   Up Time       : 0d 00:06:33
Expires          : Not running   Mode          : Include
V1 Host Timer    : Not running   Type          : Dynamic
V2 Host Timer    : Not running   Compat Mode: IGMP Version 3
Redir.SvcId      : N/A          Redir.Intf   : int-LAG-REDIRECTED
```

```
-----
Source Address   Expires      Type      Fwd/Blk
-----
192.168.4.2      0d 00:02:55 Dynamic    Fwd
-----
```

```
Hosts : 1
=====
```


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

*A:BNG-1#

ESM IGMP PPPoE Walkthrough

The same baseline configuration is used for PPPoE subscriber. [Figure 47: PPPoE Multicast Flow](#) shows the message flow for delivery of multicast streams to PPPoE subscribers.

Figure 47: PPPoE Multicast Flow



By default, PPPoE subscribers receive multicast streams via Ethernet unicast over subscriber SAPs. PPPoE does not have a multicast mechanism and requires all data traffic to be unicasted. However, because multicast streams are redirected, the streams are sent as multicast at both Layers 2 and 3 (the Layer 2 header will have a multicast destination MAC address and the Layer 3 header will have a multicast destination IP address).

Verify the IGMP on the group-interface. It shows very little difference from the IPoE group interface. No multicast streams are delivered directly over the subscriber SAP group interface.

```
*A:BNG-1# show router igmp group-interface detail
```

```
=====
IGMP Group-Interfaces
=====
FwdSvc/Grp-Intf   : 1/grp-int-1
Admin-Status      : Up
Import-Policy     : none
Router-Alert-Check : Enabled
MCAC Policy Name  :
MCAC Max Unconst BW: no limit
MCAC In use Mand BW: 0
MCAC In use Opnl BW: 0
Qry Interval      : None
Qry Resp Interval : None
MCAC If-Policy Name:

Oper-Status       : Up
Subnet-Check      : Enabled
Sub-Hosts-Only    : Enabled
MCAC Const Adm St : Enable
MCAC Max Mand BW  : no limit
MCAC Avail Mand BW : unlimited
MCAC Avail Opnl BW : unlimited
Qry Last Mbr Inter*: None

-----
SAP                : lag-1:4
Admin/Oper version: 3/3
Max Groups Allowed: No Limit
Max Sources Allow*: No Limit
Max Grp Srcs Allo*: No Limit
Qry Interval       : 125
Qry Resp Interval  : 10
Num Groups         : 0
Max Groups Till Now: 1
Qry Last Memb Inte*: 1
-----
```

```
SAP                : lag-1:5
Admin/Oper version: 3/3      Num Groups          : 0
Max Groups Allowed: No Limit Max Groups Till Now: 0
Max Sources Allow*: No Limit
Max Grp Srcs Allo*: No Limit
Qry Interval       : 125      Qry Last Memb Inte*: 1
Qry Resp Interval  : 10
```

```
-----
Group-Interfaces = 1, SAPs = 2
=====
```

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

*A:BNG-1#

All multicast streams should be delivered over the redirected interface. The following output shows the IGMP group for a PPPoE subscriber and also that the multicast stream is associated with the *int-LAG-REDIRECTED* interface. The (S,G) is (192.168.4.2, 239.255.1.1) and the subscriber IP address is 10.0.0.12.

```
*A:BNG-1# show router igmp group
=====
IGMP Interface Groups
=====
(192.168.4.2,239.255.1.1)          UpTime: 0d 00:02:03
  Fwd List  : int-LAG-REDIRECTED
-----
Entries : 1
=====
IGMP Host Groups
=====
(192.168.4.2,239.255.1.1)          UpTime: 0d 00:02:03
  Fwd List  : 10.0.0.12
-----
Entries : 1
=====
IGMP SAP Groups
=====
No Matching Entries
=====
*A:BNG-1#
```

The following output shows all the subscribers and the (S,G)s they have joined. Note that there is only one PPPoE subscriber and the multicast stream is redirected.

```
*A:BNG-1# show router igmp hosts detail
=====
IGMP Host 10.0.0.12
=====
Oper Status       : Up      MacAddress        : 00:10:94:00:00:03
Oper version      : 3        Subscriber         : pppoe-sub-1
Num Groups        : 1        GrpItf            : grp-int-1
Max Grps Till Now: 1        IGMP-Policy       : igmp-policy-1
PPPoE SessionId   : 1        Next query time: 0d 00:00:23
FwdSvcId          : 1        Max Srcs Allow*:  No Limit
Max Grps Allowed  : No Limit Max Grp Srcs A*:  No Limit
Qry Interval      : 125      Qry Last Mbr I*:  1
Qry Resp Interval: 10        Router Alert C*:  Enabled
-----
IGMP Group
-----
```



```

Group Address      : 239.255.1.1    Up Time   : 0d 00:02:50
Expires           : Not running    Mode      : Include
V1 Host Timer     : Not running    Type      : Dynamic
V2 Host Timer     : Not running    Compat Mode: IGMP Version 3
Redir.SvcId       : N/A            Redir.Intf : int-LAG-REDIRECTED

```

```

-----
Source Address    Expires      Type      Fwd/Blk
-----
192.168.4.2      0d 00:02:44  Dynamic   Fwd
-----

```

Hosts : 1

=====

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

*A:BNG-1#

To view the (S,G)s of a single subscriber, use the following command.

```
*A:BNG-1# show service active-subscribers igmp subscriber "pppoe-sub-1" detail
```

=====

Active Subscribers Detail

=====

Subscriber HostAddr GrpAddr SrcAddr	IGMP-Policy GrpItf Type Type	Up-Time	NumGroups Mode Blk/Fwd
pppoe-sub-1 10.0.0.12 239.255.1.1 192.168.4.2	igmp-policy-1 grp-int-1 Dynamic Dynamic	0d 00:04:42	1 Include Fwd

Number of Subscribers : 1

=====

*A:BNG-1#

ESM IGMP MCS

The BNGs are configured with SRRP for both IPoE and PPPoE subscribers. This provides stateful redundancy when the master BNG fails. The SRRP master BNG will be the only BNG processing and answering IGMP messages, while the standby BNG synchronizes the state information of all subscribers via MCS in real time. In the event of a failure, the standby takes over and starts processing all IGMP messages. As the standby BNG has the full state information of all subscribers, including the (S,G)s they have joined, PIM starts sending joins for those (S,G)s immediately after failover. Restoration of all multicast streams happens quickly and relies on the PIM configuration and the underlying routing infrastructure. Note that the PIM command **non-dr-attract-traffic** can be used to reduce the failover outage by attracting multicast to the non designated PIM router.

The following output shows the items that are synchronized between the BNGs. To reduce the ESM multicast restoration time, it is important that all subscriber related data (IPoE, PPPoE, SRRP and IGMP) are kept in sync. BNG-1 has system IP address 192.0.2.2 and BNG-2 has system IP address 192.0.2.3.

```
*A:BNG-1# show redundancy multi-chassis sync peer 192.0.2.3 detail
```

=====

Multi-chassis Peer Table

=====

Peer

```

-----
Peer IP Address      : 192.0.2.3
Description          : (Not Specified)
Authentication       : Disabled
Source IP Address    : 192.0.2.2
Admin State          : Enabled
Warm standby         : No
Remote warm standby  : No
-----
Sync-status
-----
Client Applications  : IGMP SUBMGMT-IPOE SUBMGMT-PPPOE SRRP
Sync Admin State     : Up
Sync Oper State      : Up
Sync Oper Flags      :
DB Sync State        : inSync
Num Entries          : 31
Lcl Deleted Entries  : 0
Alarm Entries        : 0
OMCR Standby Entries : 0
OMCR Alarm Entries   : 0
Rem Num Entries      : 31
Rem Lcl Deleted Entries : 2
Rem Alarm Entries    : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries : 0

=====
MCS Application Stats
=====
Application          : igmp
Num Entries          : 3
Lcl Deleted Entries  : 0
Alarm Entries        : 0
OMCR Standby Entries : 0
OMCR Alarm Entries   : 0
-----
Rem Num Entries      : 3
Rem Lcl Deleted Entries : 2
Rem Alarm Entries    : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries : 0
-----

---snip---

-----
Application          : subMgmtIpoE
Num Entries          : 1
Lcl Deleted Entries  : 0
Alarm Entries        : 0
OMCR Standby Entries : 0
OMCR Alarm Entries   : 0
-----
Rem Num Entries      : 1
Rem Lcl Deleted Entries : 0
Rem Alarm Entries    : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries : 0
-----
Application          : srrp
Num Entries          : 26
Lcl Deleted Entries  : 0
Alarm Entries        : 0

```

```
OMCR Standby Entries      : 0
OMCR Alarm Entries       : 0
-----
Rem Num Entries          : 26
Rem Lcl Deleted Entries  : 0
Rem Alarm Entries        : 0
Rem OMCR Standby Entries : 0
Rem OMCR Alarm Entries   : 0
-----

---snip---

-----
Application              : subMgmtPppoe
Num Entries              : 1
Lcl Deleted Entries      : 0
Alarm Entries            : 0
OMCR Standby Entries     : 0
OMCR Alarm Entries       : 0
-----
Rem Num Entries          : 1
Rem Lcl Deleted Entries  : 0
Rem Alarm Entries        : 0
Rem OMCR Standby Entries : 0
Rem OMCR Alarm Entries   : 0
-----

---snip---

=====
Ports synced on peer 192.0.2.3
=====
Port/Encap                Tag
-----
lag-1
  4-4                      mclagdata
  5-5                      mclagcontrol
  4094-4094                mclagmulticast
=====

---snip---

*A:BNG-1#
```

To check the details of the sync data across the BNGs, a tools command giving a detailed description of the IGMP information synced across MCS can be used.

```
*A:BNG-1# tools dump redundancy multi-chassis sync-database application igmp detail

If no entries are present for an application, no detail will be displayed.

FLAGS LEGEND: ld - local delete; da - delete alarm; pd - pending global delete;
              oal - omcr alarmed; ost - omcr standby

Peer Ip 192.0.2.3

Application IGMP
Sap-id      Client Key      DLen  Flags      timeStamp
SyncTag     deleteReason code and description      #ShRec
-----
lag-1:4     Host=10.0.0.15, HostGroup=239.255.1.1
```

```
mclagdata          20  -- -- -- -- 06/28/2017 09:53:28
0x0                0
lag-1:4            Host=10.0.0.16, HostGroup=239.255.1.1
mclagdata          20  -- -- -- -- 06/28/2017 09:53:37
0x0                0
lag-1:4094         Group=239.255.1.1
mclagmulticast     20  -- -- -- -- 06/28/2017 09:53:28
0x0                0

The following totals are for:
peer ip ALL, port/lag/sdp ALL, sync-tag ALL, application IGMP
Valid Entries:      3
Locally Deleted Entries: 0
Locally Deleted Alarmed Entries: 0
Pending Global Delete Entries: 0
Omcrr Alarmed Entries: 0
Omcrr Standby Entries: 0
Associated Shared Records (ALL): 0
Associated Shared Records (LD): 0
*A:BNG-1#
```

ESM IGMP Debug

Debug facilities allow for real-time monitoring of events happening on the system. This includes tools for debugging ESM multicast streams.

First enable the required debug on the system, then send an IGMP message to join a multicast group (S,G). The message used in this example is an IGMPv3 message with SSM.

The following is the debug information for an ESM IGMP report message at packet level.

```
debug
router
  igmp
    packet mode egr-ingr-and-dropped
  exit
exit
exit
```

```
13 2017/06/28 10:09:32.86 CEST MINOR: DEBUG #2001 Base IGMP[1]
"IGMP[1]: RX-PKT
[000 22:17:54.700] IGMP host 10.0.0.15 V3 PDU: 10.0.0.15 -> 224.0.0.22 pduLen 20
  Type: V3 REPORT maxrespCode 0x0 checksum 0x2352
  Num Group Records: 1
    Group Record 0
      Type: ALW_NEW_SRCS, AuxDataLen 0, Num Sources 1
      Mcast Addr: 239.255.1.1
      Source Address List
        192.168.4.2
"
```

The following is the debug information for an ESM IGMP host. The multicast stream is redirected to the LAG interface and that an MCS entry is installed for the new IGMP group.

```
debug
router
  igmp
    host "10.0.0.15"
  exit
```

```
exit
exit
```

```
22 2017/06/28 10:11:53.84 CEST MINOR: DEBUG #2001 Base IGMP[Base inst 1]
"IGMP[Base inst 1]: igmpIfGroupAdd
Adding 239.255.1.1 to IGMP host 10.0.0.15 database"

23 2017/06/28 10:11:53.84 CEST MINOR: DEBUG #2001 Base IGMP[Base inst 1]
"IGMP[Base inst 1]: igmpProcessGroupRec
Process group rec ALW_NEW_SRCS received on host 10.0.0.15 for group 239.255.1.1
in mode INCLUDE. Num srcs 1"

24 2017/06/28 10:11:53.84 CEST MINOR: DEBUG #2001 Base IGMP[Base inst 1]
"IGMP[Base inst 1]: igmpIfSrcAdd
Adding i/f source entry for host 10.0.0.15 (192.168.4.2,239.255.1.1) to IGMP fwd
List Database, redir if interface int-LAG-REDIRECTED [ifIndex 7]"

25 2017/06/28 10:11:53.84 CEST MINOR: DEBUG #2001 Base IGMP[Base inst 1]
"IGMP[Base inst 1]: igmpMcsAddIfGroup
Building MCS entry for host 10.0.0.15, group 239.255.1.1"

26 2017/06/28 10:11:56.31 CEST MINOR: DEBUG #2001 Base IGMP[Base inst 1]
"IGMP[Base inst 1]: igmpProcessGroupRec
Process group rec ALW_NEW_SRCS received on host 10.0.0.15 for group 239.255.1.1
in mode INCLUDE. Num srcs 1"

27 2017/06/28 10:11:56.30 CEST MINOR: DEBUG #2001 Base IGMP[Base inst 1]
"IGMP[Base inst 1]: igmpMcsAddIfGroup
Building MCS entry for host 10.0.0.15, group 239.255.1.1"
```

The following debug information is received for ESM IGMP when MCS sync is enabled. The MCS sends a sync message for the redirect interface.

```
debug
router
  igmp
    mcs "int-LAG-REDIRECTED"
  exit
exit
exit
```

```
46 2017/06/28 10:19:06.23 CEST MINOR: DEBUG #2001 Base IGMP MCS[1]
"IGMP MCS[1]: TX-MCS Data
interface int-LAG-REDIRECTED [ifIndex 7]
Key Type: Group, Len: 9, Grp Addr: 239.255.1.1
Data Type: Group, Len: 20, Ver: 0, RecType: 1, Compat Mode: 3,
Num Fwd Srcs: 1, Num Blk Srcs: 0
Fwd Sources:
  192.168.4.2
"

47 2017/06/28 10:19:11.05 CEST MINOR: DEBUG #2001 Base IGMP MCS[1]
"IGMP MCS[1]: TX-MCS Data
interface int-LAG-REDIRECTED [ifIndex 7]
Key Type: Group, Len: 9, Grp Addr: 239.255.1.1
Data Type: Group, Len: 20, Ver: 0, RecType: 1, Compat Mode: 3,
Num Fwd Srcs: 1, Num Blk Srcs: 0
Fwd Sources:
  192.168.4.2
"
```

The corresponding debug information for ESM IGMP MCS sync on BNG-2 looks as follows:

```
1 2017/06/28 10:19:06.24 CEST MINOR: DEBUG #2001 Base IGMP MCS[1]
"IGMP MCS[1]: RX-MCS Data
interface int-LAG-REDIRECTED [ifIndex 8]
Key Type: Group, Len: 9, Grp Addr: 239.255.1.1
Data Type: Group, Len: 20, Ver: 0, RecType: 1, Compat Mode: 3,
Num Fwd Srcs: 1, Num Blk Srcs: 0
Fwd Sources:
    192.168.4.2
"
```

The same debug commands can be used for viewing IGMP leave messages. The following debug information is received for an ESM IGMP leave at the packet level. The leave report message received over the subscriber SAP results in the multicast stream being stopped on the redirected interface, after ensuring no other CPE devices still require the multicast streams (by means of a query).

```
debug
router
    igmp
        packet mode egr-ingr-and-dropped
    exit
exit
exit
```

```
64 2017/06/28 10:26:25.31 CEST MINOR: DEBUG #2001 Base IGMP[1]
"IGMP[1]: RX-PKT
[000 22:34:47.160] IGMP host 10.0.0.15 V3 PDU: 10.0.0.15 -> 224.0.0.22 pduLen 20
Type: V3 REPORT maxrespCode 0x0 checksum 0x2252
Num Group Records: 1
Group Record 0
Type: BLK_OLD_SRCS, AuxDataLen 0, Num Sources 1
Mcast Addr: 239.255.1.1
Source Address List
    192.168.4.2
"
```

```
65 2017/06/28 10:26:25.31 CEST MINOR: DEBUG #2001 Base IGMP[1]
"IGMP[1]: TX-PKT
[000 22:34:47.160] IGMP interface int-LAG-REDIRECTED [ifIndex 7] V3 PDU: 192.168
.10.253 -> 239.255.1.1 pduLen 16
Type: QUERY maxrespCode 0xa checksum 0x36cc
GroupAddr: 239.255.1.1
S bit 0, QRV 2, Encoded-QQIC 125, NumSources 1
Source Address List:
    192.168.4.2
"
```

```
66 2017/06/28 10:26:26.86 CEST MINOR: DEBUG #2001 Base IGMP[1]
"IGMP[1]: TX-PKT
[000 22:34:48.710] IGMP interface int-LAG-REDIRECTED [ifIndex 7] V3 PDU: 192.168
.10.253 -> 239.255.1.1 pduLen 16
Type: QUERY maxrespCode 0xa checksum 0x36cc
GroupAddr: 239.255.1.1
S bit 0, QRV 2, Encoded-QQIC 125, NumSources 1
Source Address List:
    192.168.4.2
"
```

The following debug information is received for an ESM IGMP host showing various IGMP events. The MCS also signals the removal of the IGMP entry in the database.

```
debug
  router
    igmp
      host "10.0.0.15"
    exit
  exit
exit
```

```
73 2017/06/28 10:30:29.95 CEST MINOR: DEBUG #2001 Base IGMP[Base inst 1]
"IGMP[Base inst 1]: igmpProcessGroupRec
Process group rec BLK_OLD_SRCS received on host 10.0.0.15 for group 239.255.1.1
in mode INCLUDE. Num srcs 1"

74 2017/06/28 10:30:29.95 CEST MINOR: DEBUG #2001 Base IGMP[Base inst 1]
"IGMP[Base inst 1]: igmpProcessIfSrcTimerExp
Source Timer expired for IGMP host 10.0.0.15 (192.168.4.2,239.255.1.1)"

75 2017/06/28 10:30:29.95 CEST MINOR: DEBUG #2001 Base IGMP[Base inst 1]
"IGMP[Base inst 1]: igmpIfSrcDel
Deleting i/f source entry for host 10.0.0.15 (192.168.4.2,239.255.1.1) from IGMP
Database. DeleteFromAvl: 1 Redir 0"

76 2017/06/28 10:30:29.95 CEST MINOR: DEBUG #2001 Base IGMP[Base inst 1]
"IGMP[Base inst 1]: igmpIfGroupDel
Deleting 239.255.1.1 from IGMP host 10.0.0.15 database"

77 2017/06/28 10:30:29.95 CEST MINOR: DEBUG #2001 Base IGMP[Base inst 1]
"IGMP[Base inst 1]: igmpMcsDelIfGroup
Building MCS entry for host 10.0.0.15, group 239.255.1.1"

78 2017/06/28 10:30:29.95 CEST MINOR: DEBUG #2001 Base IGMP[Base inst 1]
"IGMP[Base inst 1]: igmpMcsDelIfGroup
Deleting MCS entry for host 10.0.0.15, group 239.255.1.1, Glb"
```

The following debug information is received when MCS removes the entry on BNG-1. MCS also triggers the backup BNG to remove the multicast stream.

```
debug
  router
    igmp
      mcs "int-LAG-REDIRECTED"
    exit
  exit
exit
```

```
84 2017/06/28 10:34:35.06 CEST MINOR: DEBUG #2001 Base IGMP MCS[1]
"IGMP MCS[1]: TX-MCS Data
interface int-LAG-REDIRECTED [ifIndex 7]
Key Type: Group, Len: 9, Grp Addr: 239.255.1.1
Data Type: Group, Len: 20, Ver: 0, RecType: 1, Compat Mode: 3,
Num Fwd Srcs: 1, Num Blk Srcs: 0
Fwd Sources:
  192.168.4.2
"

85 2017/06/28 10:34:36.86 CEST MINOR: DEBUG #2001 Base IGMP MCS[1]
"IGMP MCS[1]: TX-MCS Data (GlbDel)
interface int-LAG-REDIRECTED [ifIndex 7]
```

```
Key Type: Group, Len: 9, Grp Addr: 239.255.1.1
Data Type: Group, Len: 16, Ver: 0, RecType: 1, Compat Mode: 3,
Num Fwd Srcs: 0, Num Blk Srcs: 0
"
```

The debug information on BNG-2 shows the sync message received over MCS for the removal of the multicast (S,G).

```
9 2017/06/28 10:34:22.88 CEST MINOR: DEBUG #2001 Base IGMP MCS[1]
"IGMP MCS[1]: RX-MCS Data
interface int-LAG-REDIRECTED [ifIndex 8]
Key Type: Group, Len: 9, Grp Addr: 239.255.1.1
Data Type: Group, Len: 20, Ver: 0, RecType: 1, Compat Mode: 3,
Num Fwd Srcs: 1, Num Blk Srcs: 0
Fwd Sources:
    192.168.4.2
"
```

Conclusion

Multicast is an essential part of Triple Play Services. The SR OS TPSDA solution is much more than a baseline multicast delivery, it includes individual subscriber awareness and offers a full state redundancy option. Subscriber awareness allows for fine tuning of subscriber multicast settings and for troubleshooting on a per subscriber basis. Full state redundancy reduces failover time and ensures high availability of multicast services. This example provided a complete configuration walkthrough of both the IPoE and PPPoE SRRP model with redirection. All multicast streams can be redirected to a dedicated interface for all subscribers to receive.

ESM IPv4: Multicast with SRRP

This chapter describes ESM IPv4 multicast with SRRP configurations.

Topics in this chapter include:

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

Applicability

This chapter applies to SR OS routers, and was originally written for Release 11.0.R1. The CLI is updated to Release 15.0.R3. It covers multicast with Subscriber Routed Redundancy Protocol (SRRP) for IPoE and PPPoE subscribers.


Overview

Triple Play Service Delivery Architecture (TPSDA) has allowed operators to integrate High Speed Internet (HSI), voice, and video services within a single network. The goal of this chapter is to walk through the configuration of a redundant TPSDA multicast architecture and the configuration of multicast filters. The topics are divided into the following sections:

- Enhanced Subscriber Management (ESM) multicast baseline configuration
 - IGMP configuration on ESM group interface
 - ESM IGMP-policy configuration
- IPoE ESM multicast walkthrough
- PPPoE ESM multicast walkthrough
- IGMP Subscriber Router Redundancy Protocol (SRRP)
 - Multi-Chassis Synchronization (MCS) walkthrough
- IGMP Debugging
- IGMP Control Plane Filters
- IGMP Data Plane Filters

The network topology displayed in [Figure 48: Network Topology Overview](#) shows a typical TPSDA setup. It consists of three 7750s and a single 7450. Two 7750s are configured as Broadband Network Gateways (BNGs) and the third 7750 is configured as a P router. The 7450 is used as an aggregation switch to aggregate all subscribers.

Figure 48: Network Topology Overview



al_0261

Both BNGs are configured with SRRP to provide redundancy. SRRP is only used for redundancy purposes. SRRP is not required for supporting multicast. The P router is connected to the multicast source and to the network side of both BNGs. The connections between the BNGs and the P router are also running PIM to provide multicast delivery. On the access side, the two BNGs are connected to an aggregation switch which aggregates the traffic originating from both PPPoE and IPoE subscribers. The BNGs are IGMP capable and will respond to subscribers' IGMP requests.

There are two requirements to enable multicast delivery using ESM. First, the ESM group interface must have IGMP enabled. Second, the ESM subscribers must be configured with an IGMP-policy to receive multicast. When both requirements are met, the BNG will process the subscribers' IGMP messages, otherwise, IGMP messages are simply ignored and dropped. All customer premise equipment (CPE) IGMP messages are aggregated via the 7450 and passed to the BNGs. Because the BNGs are running SRRP, the SRRP master is the only BNG processing and answering the IGMP messages. Protocol Independent Multicasting (PIM) is then used between the BNG and the P router to request the multicast stream. If PIM is successful in retrieving the multicast group, the multicast stream is forwarded toward the subscribers. This is the typical multicast delivery model for TPSDA.


Configuration

This section expects a basic knowledge of ESM.

ESM SRRP Baseline Configuration

[Figure 49: Example Topology](#) shows the addressing scheme used in the setup. The example uses numbered SRRP subscriber interfaces with static SAPs serving both IPoE and PPPoE subscribers. The configuration of the RADIUS server is out of the scope of this example.

Figure 49: Example Topology



The baseline configuration for BNG-1 follows, excluding the IGMP configuration. In this example, the subscriber-interface is configured in an IES, although it is possible to configure the subscriber-interface in a VPRN. OSPF and PIM are provisioned to provide routing and multicast capabilities. The SRRP configuration with priority 200 ensures BNG-1 is the master when both BNGs are active because the SRRP priority for BNG-2 is lower.

```
# BNG-1
configure
service
  ies 1 customer 1 create
  redundant-interface "mclink-BNG-1-BNG-2" create
  address 192.168.11.0/31
  ip-mtu 1500
  spoke-sdp 23:1 create
  no shutdown
  exit
exit
subscriber-interface "sub-int-1" create
address 10.255.255.252/8 gw-ip-address 10.255.255.254 track-srrp 1
group-interface "grp-int-1" create
srrp-enabled-routing
dhcp
  option
    action replace
    circuit-id
    no remote-id
  exit
  server 192.168.0.1
  lease-populate 10
  client-applications dhcp ppp
  gi-address 10.255.255.252
  no shutdown
exit
authentication-policy "auth-1"
```

```
        redundant-interface "mclink-BNG-1-BNG-2"
        sap 1/1/4:4 create
            sub-sla-mgmt
                def-sub-id use-sap-id
                def-sub-profile "sub-profile-1"
                def-sla-profile "sla-profile-1"
                sub-ident-policy "sub-ident-1"
                multi-sub-sap 10
                no shutdown
            exit
        exit
        sap 1/1/4:5 create
        exit
        srrp 1 create
            message-path 1/1/4:5
            priority 200
            no preempt
            no shutdown
        exit
        pppoe
            no shutdown
        exit
    exit
exit
no shutdown
exit
exit
exit
```

```
configure
router
    ospf
        area 0.0.0.0
            interface "system"
                no shutdown
            exit
            interface "int-BNG-1-BNG-2"
                interface-type point-to-point
                metric 10000
                no shutdown
            exit
            interface "int-BNG-1-P-1"
                interface-type point-to-point
                no shutdown
            exit
            interface "sub-int-1"
                no shutdown
            exit
        exit
    exit
    no shutdown
    exit
    pim
        interface "int-BNG-1-P-1"
            exit
        no shutdown
    exit
exit
exit
```

The baseline configuration for BNG-2 follows, excluding the IGMP configuration. The SRRP priority for BNG-2 is lower than the SRRP priority for BNG-1, thus BNG-2 will be in standby mode.

```
# BNG-2
configure
service
    ies 1 customer 1 create
        no shutdown
        redundant-interface "mclink-BNG-2-BNG-1" create
            address 192.168.11.1/31
            ip-mtu 1500
            spoke-sdp 32:1 create
                no shutdown
            exit
        exit
    subscriber-interface "sub-int-1" create
        address 10.255.255.253/8 gw-ip-address 10.255.255.254 track-srrp 1
        group-interface "grp-int-1" create
            srrp-enabled-routing
            dhcp
                option
                    action replace
                    circuit-id
                    no remote-id
                exit
                server 192.168.0.1
                lease-populate 10
                client-applications dhcp ppp
                gi-address 10.255.255.253
                no shutdown
            exit
        authentication-policy "auth-1"
        redundant-interface "mclink-BNG-2-BNG-1"
        sap 1/1/1:4 create
            sub-sla-mgmt
                def-sub-id use-sap-id
                def-sub-profile "sub-profile-1"
                def-sla-profile "sla-profile-1"
                sub-ident-policy "sub-ident-1"
                multi-sub-sap 10
                no shutdown
            exit
        exit
    sap 1/1/1:5 create
        exit
    srrp 1 create
        message-path 1/1/1:5
        priority 150
        no preempt
        no shutdown
    exit
    pppoe
        no shutdown
    exit
    exit
    exit
    no shutdown
exit
exit
exit
exit
```

```
configure
router
```

```
ospf
  area 0.0.0.0
    interface "system"
      no shutdown
    exit
    interface "int-BNG-2-BNG-1"
      interface-type point-to-point
      metric 10000
      no shutdown
    exit
    interface "int-BNG-2-P-1"
      interface-type point-to-point
      no shutdown
    exit
    interface "sub-int-1"
      no shutdown
    exit
  exit
no shutdown
exit
pim
  interface "int-BNG-2-P-1"
  exit
  no shutdown
exit
exit
exit
```

The baseline configuration for the aggregation switch AGG-1 follows. Two VPLS services are configured. VPLS 5 is responsible for passing SRRP control traffic over VLAN 5. VPLS 4 is responsible for passing all subscriber data traffic over VLAN 4.

```
# AGG-1
configure
  service
    vpls 4 customer 1 create
      description "for user traffic"
      stp
        shutdown
      exit
      sap 1/1/6:4 create
        no shutdown
      exit
      sap 2/1/1:4 create
        no shutdown
      exit
      sap 2/1/2:4 create
        no shutdown
      exit
      no shutdown
    exit
    vpls 5 customer 1 create
      description "for SRRP/redundancy"
      stp
        shutdown
      exit
      sap 2/1/1:5 create
        no shutdown
      exit
      sap 2/1/2:5 create
        no shutdown
      exit
      no shutdown
```

```
        exit
    exit
exit
```

The baseline configuration on the P router is as follows. The P router has a local DHCP server configured and performs the DHCP address assignment. The P router also is attached to the multicast source and uses PIM to deliver multicast streams.

```
# P-1
configure
router Base
  dhcp
    local-dhcp-server "dhcp-local-server" create
    use-gi-address scope pool
    pool "pool-1" create
    subnet 10.0.0.0/8 create
    options
      subnet-mask 255.0.0.0
      default-router 10.255.255.254
    exit
    address-range 10.0.0.10 10.0.0.254
  exit
  exit
  no shutdown
exit
interface "int-DHCP-lb1"
  address 192.168.0.1/32
  loopback
  local-dhcp-server "dhcp-local-server"
  no shutdown
exit
interface "int-P-1-BNG-1"
  address 192.168.24.2/30
  port 1/1/1:4094
  no shutdown
exit
interface "int-P-1-BNG-2"
  address 192.168.34.2/30
  port 1/1/2:4094
  no shutdown
exit
interface "int-P-1-mcast-source"
  address 192.168.4.1/30
  port 1/1/5:4092
  no shutdown
exit
interface "system"
  address 192.0.2.4/32
  no shutdown
exit
ospf 0
  traffic-engineering
  area 0.0.0.0
    interface "system"
      no shutdown
    exit
    interface "int-DHCP-lb1"
      no shutdown
    exit
    interface "int-P-1-BNG-1"
      interface-type point-to-point
      no shutdown
```

```

        exit
        interface "int-P-1-BNG-2"
            interface-type point-to-point
            no shutdown
        exit
        interface "int-P-1-mcast-source"
            passive
            no shutdown
        exit
    exit
    no shutdown
exit
pim
    interface "int-P-1-BNG-1"
    exit
    interface "int-P-1-BNG-2"
    exit
    interface "int-P-1-mcast-source"
    exit
    ---snip---
    no shutdown
    exit
exit
exit
exit

```

Enable IGMP on Group Interfaces

The configuration below adds the group interface to IGMP. If the subscriber-interface is configured in a VPRN, each VPRN will have its individual IGMP instance. Add the group-interface to the IGMP instance.

```

# on BNG-1 and BNG-2
configure
    router
        igmp
            group-interface "grp-int-1"
            no shutdown
        exit
        no shutdown
    exit
exit
exit
exit

```

Placing the group-interface into IGMP is the first step required to deliver multicast streams. The options available in this IGMP context can be classified into two categories:

1. Bandwidth and multicast group management
2. Interoperability

[no] disable-router*	- Enable/disable the IGMP router alert check option
[no] import	- Import a policy to filter IGMP packets
[no] max-groups	- Configure the maximum number of groups for this group-interface
[no] max-grp-sources	- Configure the maximum number of group sources for this group-interface
[no] max-sources	- Configure the maximum number of sources for this group-interface
[no] mcac	+ Configure multicast CAC policy and constraints for this interface
[no] query-interval	- Configure the frequency at which Host-Query packets are transmitted
[no] query-last-mem*	- Configure the frequency at which Group-Specific-Query packets are transmitted


```
[no] query-response* - Configure the time to wait to receive a response to the
                        Host-Query message from the host
[no] query-src-ip    - Configure the IP source address used in IGMP queries for this
                        group interface
[no] shutdown        - Administratively enable/disable the interface
[no] sub-hosts-only  - Enable/disable the IGMP traffic from known hosts only
[no] subnet-check    - Enable/disable local subnet checking for IGMP
[no] version         - Configure the version of IGMP
```

The bandwidth and multicast group management options are:

- Import — Used for white-listing or black-listing multicast groups in the IGMP control plane. More configuration detail is offered in a later section.
- Max-groups — Controls the maximum number of groups (channels) allowed on the group interface.
- Max-sources — Controls the maximum number of sources of the multicast streams on a group interface.
- Max-grp-sources — Specifies the maximum number of multicast group and source pairs for a group-interface.
- MCAC — Multicast Connection Admission Control (MCAC) is a bandwidth management feature to control the amount of multicast streams a group interface is allowed to receive. It can also be applied at subscriber level to offer a hierarchical control. Multicast bandwidth can be controlled on a per subscriber basis.
- query-interval — defines the frequency at which general host-query messages are sent out by the querier.
- query-last-member-interval — defines the frequency at which the querier sends group-specific queries including messages sent in response to leave-group messages.
- query-response-interval — defines how long the querier waits to receive a response to a host-query message from a host.
- query-src-ip — defines the query source IP address for the group interface.

The interoperability options available are:

- Router alert — enable or disable router alert processing.
- Sub-hosts-only — Only subscriber originated IGMP messages are accepted and anything else is ignored. Sometimes, IGMP message might not arrive directly from the subscriber. For example, an aggregation switch or DSLAM residing between the CPE and the BNG might perform IGMP proxy. The switch/DSLAM will insert its own source IP-address in place of the subscriber.

It should be noted that, when an IGMP proxy is used, the identity of the subscriber is lost (because the original source IP of the IGMP message is replaced).

- Subnet-check — IGMP messages will be checked against the group interface subnet. All IGMP messages with a source address that is not in the local subnet are dropped.
- Version — The RFCs define three versions of IGMP, all of them are supported by SR OS.

It must be noted that when subscribers are sending IGMPv1 or v2 in a bridged LAN, suppression of IGMP messages can occur. If an IGMP host detects the presence of another host reporting for the same multicast group, it will suppress its own IGMP report message and silently receive the multicast stream. When IGMP messages are suppressed, the BNG might not be able to account for the real multicast bandwidth consumption of each subscriber. IGMPv3, on the other hand, forces all hosts to send IGMP reports. This guarantees that the BNG identifies each subscriber's IGMP request.

ESM IGMP Policy

In addition to enabling IGMP on the group interface, the subscriber must be allowed to receive multicast streams through an IGMP policy. For this purpose, the IGMP-policy is associated with the subscriber's subscriber-profile. Therefore during authentication, either RADIUS, the local user database (LUDB), or the default-sub-profile should return a sub-profile with an IGMP policy. The provisioning requires two steps:

First create the IGMP policy, as follows:

```
configure
  subscriber-mgmt
    igmp-policy "igmp-policy-1" create
  exit
exit
```

Then add the IGMP policy to a subscriber-profile, as follows:

```
configure
  subscriber-mgmt
    sub-profile "sub-profile-1" create
    igmp-policy "igmp-policy-1"
  exit
exit
```

The above configuration is the minimum requirement for a subscriber to receive multicast streams. The different options inside an IGMP policy are:

[no] description	- Description for the policy
[no] disable-router*	- Enable/disable the router alert check option
[no] egress-rate-mo*	- Configure the egress rate modification
[no] fast-leave	- Enable/disable fast-leave processing
[no] import	- Specify the import policy to filter packets
[no] max-num-groups	- Configure the max number of multicast groups
[no] max-num-grp-so*	- Configure the max number of multicast group sources
[no] max-num-sources	- Configure the max number of multicast sources
[no] mcast-reporting +	- Configure the mcast reporting
[no] per-host-repli*	- Enable/disable per-host-replication processing
[no] query-interval	- Configure the frequency at which Host-Query packets are transmitted
[no] query-last-mem*	- Configure the frequency at which Group-Specific-Query packets are transmitted
[no] query-response*	- Configure the time to wait to receive a response to the Host-Query message from the host
[no] redirection-po*	- Specify the redirection policy
static	+ Add/remove static group membership
[no] version	- Configure the version

Again, two groups of options are available: the bandwidth and multicast group management options, and the interoperability options.

Bandwidth and multicast group management options:

- Import — Used for white-listing and black-listing multicast groups. This allows the import policy to be defined per subscriber.
- Max-num-group — Limits the maximum multicast groups for the group interface. This limits the groups per subscriber.

- **Max-num-sources** — Limits the maximum multicast sources for the group interface. This limits the sources per subscriber.
- **Max-num-grp-sources** — Limits the maximum multicast group and source pairs for the group interface.
- **Egress-rate-modify** — This feature adjusts the subscriber queue bandwidth according to multicast consumption. It is used in conjunction with MCAC. An MCAC policy defines the bandwidth consumption per multicast group. As a subscriber joins a multicast group, the bandwidth of the multicast channel is subtracted from the subscriber queue bandwidth. The remaining bandwidth is what the subscriber can use for all other services.
- **query-interval** — defines the frequency at which general host-query messages are sent out by the querier.
- **query-last-member-interval** — defines the frequency at which the querier sends group-specific queries including messages sent in response to leave-group messages.
- **query-response-interval** — defines how long the querier waits to receive a response to a host-query message from a host.

Interoperability options:

- **Fast-leave** — Enables the router to withdraw the multicast group quickly when receiving an IGMP leave message without any last query. This should be used in a subscriber per SAP (dot1q or qinq) model.
- **Static** — This allows the provisioning of static multicast groups that the subscriber will receive regardless of any IGMP report. The static multicast group can be Source Specific Multicast (SSM) based.
- **Per-host-replication** — SR OS has the capability to replicate a multicast source per subscriber. For example, if 10 subscribers are requesting the same multicast group, then 10 multicast streams are replicated and delivered individually to each subscriber. PPPoE requires service delivery to be point to point. To achieve this, the Ethernet header destination address for the multicast stream is the subscriber's source MAC. The IP layer destination is the same as the multicast group that the subscriber requested. For IPoE, without per-host replication, the standard multicast MAC representing the multicast group is used as the destination MAC address when delivering the multicast streams to the subscribers. When per-host replication is used for IPoE subscribers, the destination MAC address will be the host's own MAC address and the IP destination will be the same as the multicast group that the subscriber is interested in. The multicast stream is then delivered to the subscriber via MAC learning as a unicast stream.

When per-host-replication is enabled, all multicast streams will be using the subscriber queues. It is no longer necessary to use egress-rate-modify as mentioned before.

- **Redirection-policy** — Another popular model for multicast deployment is to redirect all multicast streams to another interface instead of sending the content directly to the subscriber. All subscribers use a common VLAN to receive the multicast streams.

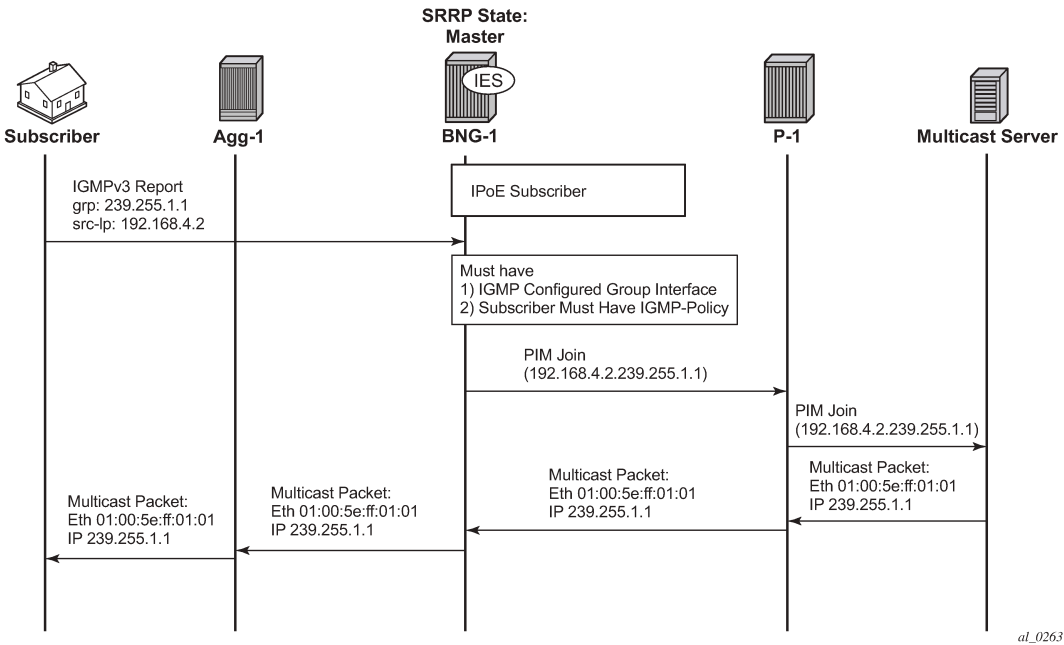
ESM IGMP IPoE Walkthrough

With the baseline configuration applied, the BNG is ready to process IGMP messages and deliver multicast. [Figure 50: IPoE Subscriber Multicast Flow](#) shows a flow for IPoE subscribers requesting and receiving multicast traffic. The key items are highlighted in dotted box:

1. The ESM group-interface must have IGMP enabled
2. The subscriber must be associated with an IGMP-policy via sub-profile.

The subscriber sends an IGMPv3 report using (192.168.4.2, 239.255.1.1).

Figure 50: IPoE Subscriber Multicast Flow



To verify that the group interface is ready for multicast, use the following **show** command. Remember that the IES service id is 1 and the group-interface name is *grp-int-1*.

```
*A:BNG-1# show router igmp group-interface

=====
IGMP Group-Interfaces
=====
FwdSvc Group-Interface      Adm/Opr-State  Import-Policy
      SAP                  Adm/Opr-Version Num-Groups
-----
1      grp-int-1            Up/Up          none
      1/1/4:4               3/3            0
      1/1/4:5               3/3            0
-----
Group-Interfaces = 1, SAPs = 2
=====
*A:BNG-1#
```

Ensure the subscriber is associated with an IGMP-policy. Since the IGMP-policy is associated with a subscriber-profile, verifying the IGMP-policy is done through the subscriber profile, as follows:

```
*A:BNG-1# show subscriber-mgmt sub-profile "sub-profile-1"

=====
Subscriber Profile sub-profile-1
=====
Description      : (Not Specified)
I. Sched. Policy : N/A
E. Sched. Policy : N/A
I. Policer Ctrl. : N/A
E. Agg Rate Limit: Max
```

```

E. Policer Ctrl. : N/A
I. vport-hashing : Disabled
I. sec-sh-hashing: Disabled
Q Frame-Based Ac*: Disabled
Acct. Policy      : N/A                      Collect Stats    : Disabled
ANCP Pol.        : N/A
Accu-stats-pol   : (Not Specified)
HostTrk Pol.     : N/A
IGMP Policy      : igmp-policy-1
MLD Policy       : N/A
PIM Policy       : N/A
Sub. MCAC Policy : N/A
NAT Policy       : N/A
Firewall Policy  : N/A
UPnP Policy      : N/A
NAT Prefix List  : N/A
Def. Encap Offset: none                     Encap Offset Mode: none
Avg Frame Size   : N/A
Vol stats type   : full
Preference       : 5
LAG hash class   : 1
LAG hash weight  : 1
-----
Radius Accounting
-----
Policy           : N/A
Session Opti.Stop: False
-----
HSM DA-2
-----
I. Qos Policy    : 1                      E. Qos Policy    : 1
                                           E. Agg Rate Limit: Max
E. WRR Policy    : N/A                     Pkt Byte Offset  : add 0*
-----
Last Mgmt Change : 06/21/2017 11:58:50
=====
* indicates that the corresponding row element may have been truncated.

---snip---

*A:BN G-1#

```

After the verification, the BNGs are ready to deliver multicast streams.

First, initiate an IGMP report from a subscriber requesting a multicast channel. In this example, IGMPv3 SSM is used. IPoE by default replicates per-SAP. If the IGMP message was successfully received and processed, an (S,G) binding will be associated with the subscriber SAP.

In this case, the IGMPv3 SSM message requests (192.168.4.2, 239.255.1.1). The subscriber host is assigned an IP address of 10.0.0.11. To verify the IGMP message was successfully processed, check that the (S,G) is learned on the IGMP instance. The following example shows a successful IGMP message processed by the BNG, the (S,G) is registered against the subscriber SAP and the host.

```

*A:BN G-1# show router igmp group
=====
IGMP Interface Groups
=====
No Matching Entries
=====
IGMP Host Groups
=====
(192.168.4.2,239.255.1.1)
  Fwd List   : 10.0.0.11                      UpTime: 0d 00:05:27

```

```

-----
Entries : 1
=====
IGMP SAP Groups
=====
(192.168.4.2,239.255.1.1)
  Fwd List : 1/1/4:4                                UpTime: 0d 00:05:27
-----
Entries : 1
=====
*A:BNG-1#

```

For more IGMP details on the group interface, including maximum multicast groups and bandwidth management, use the following command:

```

*A:BNG-1# show router igmp group-interface detail
=====
IGMP Group-Interfaces
=====
FwdSvc/Grp-Intf      : 1/grp-int-1
Admin-Status         : Up                               Oper-Status         : Up
Import-Policy        : none                             Subnet-Check        : Enabled
Router-Alert-Check   : Enabled                          Sub-Hosts-Only      : Enabled
MCAC Policy Name     :                                 MCAC Const Adm St   : Enable
MCAC Max Unconst BW  : no limit                         MCAC Max Mand BW    : no limit
MCAC In use Mand BW  : 0                               MCAC Avail Mand BW  : unlimited
MCAC In use Opnl BW  : 0                               MCAC Avail Opnl BW  : unlimited
Qry Interval         : None                             Qry Last Mbr Inter*: None
Qry Resp Interval    : None
MCAC If-Policy Name:
-----
SAP                  : 1/1/4:4
Admin/Oper version: 3/3                                Num Groups          : 1
Max Groups Allowed: No Limit                           Max Groups Till Now: 1
Max Sources Allow*: No Limit
Max Grp Srcs Allo*: No Limit
Qry Interval        : 125                               Qry Last Memb Inte*: 1
Qry Resp Interval   : 10
-----
Group-Address        : 239.255.1.1                      Up Time             : 0d 00:05:54
Expires              : N/A                              Mode                : include
V1 Host Timer        : Not running                      Type                : dynamic
V2 Host Timer        : Not running                      Compat Mode         : IGMP Version 3
-----
GrpSrc-Address       Expires                               Type                Fwd/Blk
-----
192.168.4.2          0d 00:04:18          dynamic              Fwd
-----
SAP                  : 1/1/4:5
Admin/Oper version: 3/3                                Num Groups          : 0
Max Groups Allowed: No Limit                           Max Groups Till Now: 0
Max Sources Allow*: No Limit
Max Grp Srcs Allo*: No Limit
Qry Interval        : 125                               Qry Last Memb Inte*: 1
Qry Resp Interval   : 10
-----
Group-Interfaces = 1, SAPs = 2
=====
* indicates that the corresponding row element may have been truncated.
*A:BNG-1#

```

If the subscriber fails to receive multicast traffic, check if the subscriber has an associated IGMP policy with the following command. If the subscriber entry is missing, make sure the subscriber has a sub-profile that is tied to an IGMP-policy.

```
*A:BNG-1# show service active-subscribers igmp detail

=====
Active Subscribers Detail
=====
Subscriber                               IGMP-Policy
HostAddr                                GrpItf
GrpAddr                                Type          Up-Time
SrcAddr                                Type
-----
subscr-1                                igmp-policy-1
10.0.0.11                               grp-int-1
239.255.1.1                             Dynamic       0d 00:08:50
192.168.4.2                             Dynamic
-----
Number of Subscribers : 1
=====
*A:BNG-1#
```

Another possibility for failing to receive multicast traffic could be due to the control mechanisms inside the IGMP-policy such as: bandwidth control, multicast groups restrictions, and interoperability options. Use the following command to view the IGMP policy configured control parameters.

```
*A:BNG-1# show subscriber-mgmt igmp-policy "igmp-policy-1"

=====
IGMP Policy igmp-policy-1
=====
Import Policy                           :
Admin Version                           : 3
Num Subscribers                         : 1
Host Max Group                          : No Limit
Host Max Sources                        : No Limit
Host Max Group Sources                  : No Limit
Query Interval                          : None
Query Last Member Interval              : None
Query Response Interval                 : None
Router Alert Check                      : Enabled
Fast Leave                             : yes
Redirection Policy                      :
Per Host Replication                    : no
Egress Rate Modify                      : no
Mcast Reporting Destination Name        :
Mcast Reporting Admin State             : Disabled
=====
*A:BNG-1#
```

The command to view the (S,G)s that all subscribers are requesting is as follows. Since the system has only one subscriber, this example only shows one host.

```
*A:BNG-1# show router igmp hosts detail

=====
IGMP Host 10.0.0.11
=====
Oper Status      : Up      MacAddress      : 00:10:94:00:00:13
Oper version     : 3       Subscriber      : subscr-1
```

```
Num Groups      : 1          GrpItf       : grp-int-1
Max Grps Till Now: 1        IGMP-Policy  : igmp-policy-1
PPPoE SessionId : N/A      Next query time: 0d 00:01:29
FwdSvcId        : 1        Max Srcs Allow*: No Limit
Max Grps Allowed : No Limit Max Grp Srcs A*: No Limit
Qry Interval     : 125      Qry Last Mbr I*: 1
Qry Resp Interval: 10      Router Alert C*: Enabled

-----
IGMP Group
-----
Group Address    : 239.255.1.1  Up Time       : 0d 00:10:30
Expires          : Not running  Mode           : Include
V1 Host Timer    : Not running  Type           : Dynamic
V2 Host Timer    : Not running  Compat Mode: IGMP Version 3
Redir.SvcId      : N/A         Redir.Intf    : N/A
-----
Source Address   Expires      Type        Fwd/Blk
-----
192.168.4.2      0d 00:03:46  Dynamic     Fwd
-----
Hosts : 1
=====
* indicates that the corresponding row element may have been truncated.
*A:BNG-1#
```

To check for an individual subscriber and its requested (S,G)s, the following command can be used.


```
*A:BNG-1# show service active-subscribers igmp subscriber "subscr-1" detail

=====
Active Subscribers Detail
=====
Subscriber      IGMP-Policy
HostAddr        GrpItf
GrpAddr         Type        Up-Time      NumGroups
SrcAddr         Type        Mode
                Blk/Fwd
-----
subscr-1        igmp-policy-1
10.0.0.11       grp-int-1
239.255.1.1     Dynamic     0d 00:13:31  1
192.168.4.2     Dynamic     Include
                Fwd
-----
Number of Subscribers : 1
=====
*A:BNG-1#
```

ESM IGMP PPPoE Walkthrough

IGMP message processing and delivery of multicast streams to PPPoE subscribers is considered next. [Figure 51: PPPoE Multicast Flow](#) shows the message flow for multicast stream delivery to PPPoE subscribers.

Figure 51: PPPoE Multicast Flow



As stated earlier, the important configuration aspects are highlighted in the dotted box. The main difference between IPoE subscribers and PPPoE subscribers is the encapsulation on the last stretch of the multicast data path. PPPoE subscribers receive multicast streams via Ethernet unicast while IPoE subscribers receive multicast streams via Ethernet multicast. PPPoE natively does not have a multicast mechanism and requires all data traffic to be unicasted. Even if the subscribers are on the same SAP, multicast streams are replicated per subscriber. To achieve this, the IP header indicates a multicast address while the Ethernet header destination MAC address is changed to the subscriber's MAC address.

Verify the group interface. It is very similar to the output for the IPoE group interface.

```
*A:BNG-1# show router igmp group-interface detail
```

```
=====
IGMP Group-Interfaces
=====
FwdSvc/Grp-Intf   : 1/grp-int-1
Admin-Status      : Up
Import-Policy     : none
Router-Alert-Check : Enabled
MCAC Policy Name  :
MCAC Max Unconst BW: no limit
MCAC In use Mand BW: 0
MCAC In use Opnl BW: 0
Qry Interval      : None
Qry Resp Interval : None
MCAC If-Policy Name:
Bonding connection : N/A

=====
SAP                : 1/1/4:4
Admin/Oper version: 3/3
Max Groups Allowed: No Limit
Max Sources Allow*: No Limit
Max Grp SrCs Allo*: No Limit

Oper-Status        : Up
Subnet-Check       : Enabled
Sub-Hosts-Only     : Enabled
MCAC Const Adm St  : Enable
MCAC Max Mand BW   : no limit
MCAC Avail Mand BW : unlimited
MCAC Avail Opnl BW : unlimited
Qry Last Mbr Inter*: None

Num Groups         : 0
Max Groups Till Now: 1
```

```

Qry Interval      : 125          Qry Last Memb Inte*: 1
Qry Resp Interval : 10
-----
SAP                : 1/1/4:5
Admin/Oper version: 3/3          Num Groups          : 0
Max Groups Allowed: No Limit     Max Groups Till Now: 0
Max Sources Allow*: No Limit
Max Grp Srcs Allo*: No Limit
Qry Interval      : 125          Qry Last Memb Inte*: 1
Qry Resp Interval : 10
-----
Group-Interfaces = 1, SAPs = 2
=====
* indicates that the corresponding row element may have been truncated.
*A:BNG-1#

```

Next an IGMPv3 message is sent toward the BNG. The (S,G) is (192.168.4.2, 239.255.1.1). The PPPoE subscriber is assigned an IP address of 10.0.0.15.

The following output shows the key difference between a PPPoE subscriber and an IPoE subscriber. A PPPoE multicast stream is replicated per host and not per SAP. This output shows this clearly because the multicast group is associated with the host and not with the SAP.

```

*A:BNG-1# show router igmp group
=====
IGMP Interface Groups
=====
No Matching Entries
=====
IGMP Host Groups
=====
(192.168.4.2,239.255.1.1)
  Fwd List   : 10.0.0.15          UpTime: 0d 00:03:39
-----
Entries : 1
=====
IGMP SAP Groups
=====
No Matching Entries
=====
*A:BNG-1#

```

The next command shows all of the subscribers and all the (S,G)s joined. In this case there is only one PPPoE subscriber.

```

*A:BNG-1# show router igmp hosts detail
=====
IGMP Host 10.0.0.15
=====
Oper Status      : Up           MacAddress       : 00:10:94:00:00:14
Oper version     : 3             Subscriber       : pppoe-sub-1
Num Groups       : 1             GrpItf          : grp-int-1
Max Grps Till Now: 1             IGMP-Policy     : igmp-policy-1
PPPoE SessionId  : 1             Next query time: 0d 00:01:02
FwdSvcId         : 1             Max Srcs Allow*: No Limit
Max Grps Allowed : No Limit      Max Grp Srcs A*: No Limit
Qry Interval     : 125           Qry Last Mbr I*: 1
Qry Resp Interval: 10           Router Alert C*: Enabled
-----
IGMP Group

```

```

-----
Group Address      : 239.255.1.1    Up Time       : 0d 00:05:03
Expires           : Not running    Mode          : Include
V1 Host Timer     : Not running    Type          : Dynamic
V2 Host Timer     : Not running    Compat Mode: IGMP Version 3
Redir.SvcId       : N/A           Redir.Intf    : N/A
-----
Source Address     Expires         Type          Fwd/Blk
-----
192.168.4.2       0d 00:03:22    Dynamic       Fwd
-----
Hosts : 1
=====
* indicates that the corresponding row element may have been truncated.
*A:BNG-1#

```

To view each individual subscriber and their respective (S,G)s, use the following command.

```

*A:BNG-1# show service active-subscribers igmp subscriber "pppoe-sub-1" detail

=====
Active Subscribers Detail
=====
Subscriber                               IGMP-Policy
HostAddr      GrpItf                      NumGroups
GrpAddr       Type                      Mode
SrcAddr       Type                      Blk/Fwd
-----
pppoe-sub-1                               igmp-policy-1
10.0.0.15    grp-int-1                          1
239.255.1.1  Dynamic                      0d 00:07:31 Include
192.168.4.2  Dynamic                      Fwd
-----
Number of Subscribers : 1
=====
*A:BNG-1#

```

ESM IGMP MCS

The BNGs are configured with SRRP for both IPoE and PPPoE subscribers. This provides stateful redundancy when the master BNG fails. The master BNG will be the only one processing and answering IGMP messages. The standby BNG does not perform any IGMP processing and receives updates through MCS for all subscribers in real time. In the event of a failure, the standby will become active and starts processing all IGMP messages. The standby will also immediately trigger PIM joins for all of the subscribers's (S,G)s. This is all possible because the standby is always synchronized with the master BNG prior to the failover. Restoration of all multicast channels should happen quickly after the failover and depends on both the PIM configuration and the underlying routing infrastructure.

The key parameters for MCS for ESM multicast are: syncing of subscribers (ipoe, pppoe), SRRP and IGMP. The redundancy configuration for BNG-1 is as follows. The configuration for BNG-2 is similar.

```

configure
  redundancy
    multi-chassis
      peer 192.0.2.3 create
      sync
        igmp
        srrp

```

```

sub-mgmt ipoe pppoe
port 1/1/4 create
    range 4-4 sync-tag "sub"
    range 5-5 sync-tag "srrp"
exit
no shutdown
exit
no shutdown
exit
exit
exit
exit
exit

```

The following command displays the number of entries being synced across the BNGs.

```

*A:BNG-1# show redundancy multi-chassis sync peer 192.0.2.3 detail
=====
Multi-chassis Peer Table
=====
Peer
-----
Peer IP Address      : 192.0.2.3
Description          : (Not Specified)
Authentication       : Disabled
Source IP Address    : 192.0.2.2
Admin State          : Enabled
Warm standby         : No
Remote warm standby  : No
-----
Sync-status
-----
Client Applications  : IGMP SUBMGMT-IPOE SUBMGMT-PPPOE SRRP
Sync Admin State     : Up
Sync Oper State      : Up
Sync Oper Flags      :
DB Sync State        : inSync
Num Entries          : 29
Lcl Deleted Entries  : 0
Alarm Entries        : 0
OMCR Standby Entries : 0
OMCR Alarm Entries   : 0
Rem Num Entries      : 29
Rem Lcl Deleted Entries : 0
Rem Alarm Entries    : 0
Rem OMCR Standby Entries : 0
Rem OMCR Alarm Entries : 0
-----
MCS Application Stats
=====
Application          : igmp
Num Entries          : 2
Lcl Deleted Entries  : 0
Alarm Entries        : 0
OMCR Standby Entries : 0
OMCR Alarm Entries   : 0
-----
Rem Num Entries      : 2
Rem Lcl Deleted Entries : 0
Rem Alarm Entries    : 0
Rem OMCR Standby Entries : 0
Rem OMCR Alarm Entries : 0

```

```
-----
Application      : igmpSnooping
Num Entries      : 0
Lcl Deleted Entries : 0
Alarm Entries    : 0
OMCR Standby Entries : 0
OMCR Alarm Entries : 0
-----
Rem Num Entries  : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries : 0
-----
Application      : subMgmtIpoe
Num Entries      : 1
Lcl Deleted Entries : 0
Alarm Entries    : 0
OMCR Standby Entries : 0
OMCR Alarm Entries : 0
-----
Rem Num Entries  : 1
Rem Lcl Deleted Entries : 0
Rem Alarm Entries : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries : 0
-----
Application      : srrp
Num Entries      : 26
Lcl Deleted Entries : 0
Alarm Entries    : 0
OMCR Standby Entries : 0
OMCR Alarm Entries : 0
-----
Rem Num Entries  : 26
Rem Lcl Deleted Entries : 0
Rem Alarm Entries : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries : 0
-----

---snip---

*A:BNG-1#
```

To check the details of the synchronized data across the BNGs, use the following command. It provides a detailed description of the IGMP information synced across MCS.

```
*A:BNG-1# tools dump redundancy multi-chassis sync-database application igmp detail

If no entries are present for an application, no detail will be displayed.

FLAGS LEGEND: ld - local delete; da - delete alarm; pd - pending global delete;
              oal - omcr alarmed; ost - omcr standby

Peer Ip 192.0.2.3

Application IGMP
Sap-id          Client Key
SyncTag         DLen  Flags          timeStamp
deleteReason code and description      #ShRec
-----
```

```
1/1/4:4      SapGroup=239.255.1.1, 0.0.0.0
sub          20      -- -- -- -- 06/22/2017 12:42:51
0x0          0
1/1/4:4      Host=10.0.0.16, HostGroup=239.255.1.1
sub          20      -- -- -- -- 06/22/2017 12:42:51
0x0          0

The following totals are for:
peer ip ALL, port/lag/sdp ALL, sync-tag ALL, application IGMP
Valid Entries:      2
Locally Deleted Entries: 0
Locally Deleted Alarmed Entries: 0
Pending Global Delete Entries: 0
Omcrr Alarmed Entries: 0
Omcrr Standby Entries: 0
Associated Shared Records (ALL): 0
Associated Shared Records (LD): 0
*A:BNG-1#
```

ESM IGMP Debug

There are many debug features for ESM multicast. Debug allows real-time monitoring of all events happening on the system and can assist operators with troubleshooting. First enable debug on the system, then send an IGMP message to join a multicast group (S,G). Again the IGMP message used in this case is IGMPv3 with SSM. The following is the debug information for ESM IGMP at packet level.

```
debug
router
igmp
packet mode egr-ingr-and-dropped
exit
exit
exit
```

```
1050 2017/06/22 15:58:25.90 CEST MINOR: DEBUG #2001 Base IGMP[1]
"IGMP[1]: RX-PKT
[001 02:05:49.110] IGMP host 10.0.0.17 V3 PDU: 10.0.0.17 -> 224.0.0.22 pduLen 20
Type: V3 REPORT maxrespCode 0x0 checksum 0x2352
Num Group Records: 1
Group Record 0
Type: ALW_NEW_SRCS, AuxDataLen 0, Num Sources 1
Mcast Addr: 239.255.1.1
Source Address List
192.168.4.2
"
```

Below is the debug information for ESM IGMP at host level and the associated IGMP events.

```
debug
router
igmp
host "10.0.0.17"
exit
exit
exit
```

```
1055 2017/06/22 16:00:37.72 CEST MINOR: DEBUG #2001 Base IGMP[Base inst 1]
"IGMP[Base inst 1]: igmpIfGroupAdd
```

```
Adding 239.255.1.1 to IGMP host 10.0.0.17 database"
```

```
1056 2017/06/22 16:00:37.72 CEST MINOR: DEBUG #2001 Base IGMP[Base inst 1]
"IGMP[Base inst 1]: igmpProcessGroupRec
Process group rec ALW_NEW_SRCS received on host 10.0.0.17 for group 239.255.1.1
in mode INCLUDE. Num srcs 1"

1057 2017/06/22 16:00:37.72 CEST MINOR: DEBUG #2001 Base IGMP[Base inst 1]
"IGMP[Base inst 1]: igmpIfSrcAdd
Adding i/f source entry for host 10.0.0.17 (192.168.4.2,239.255.1.1) to IGMP fwd
List Database, redir if sap 1/1/4:4"
```

Below is the debug information for ESM IGMP if MCS synchronization is enabled.

```
debug
  router
    igmp
      mcs "grp-int-1"
    exit
  exit
exit
```

```
1109 2017/06/22 16:17:24.49 CEST MINOR: DEBUG #2001 Base IGMP[Base inst 1]
"IGMP[Base inst 1]: igmpMcsAddIfGroup
Building MCS entry for host 10.0.0.17, group 239.255.1.1"
lear screen
```

The same debug commands can be used for viewing subscribers IGMP leave messages. The following is the debug information for ESM IGMP at packet level.

```
debug
  router
    igmp
      packet mode egr-ingr-and-dropped
    exit
  exit
exit
```

```
1091 2017/06/22 16:11:49.53 CEST MINOR: DEBUG #2001 Base IGMP[1]
"IGMP[1]: RX-PKT
[001 02:19:12.740] IGMP host 10.0.0.18 V3 PDU: 10.0.0.18 -> 224.0.0.22 pduLen 20
  Type: V3 REPORT maxrespCode 0x0 checksum 0x2251
  Num Group Records: 1
    Group Record 0
      Type: BLK_OLD_SRCS, AuxDataLen 0, Num Sources 1
      Mcast Addr: 239.255.1.2
      Source Address List
        192.168.4.2
"
```

The following is the debug information for ESM IGMP at host level and the associated IGMP events.

```
debug
  router
    igmp
      host "10.0.0.17"
    exit
  exit
```

```
exit
```

```
1113 2017/06/22 16:21:04.08 CEST MINOR: DEBUG #2001 Base IGMP[Base inst 1]
"IGMP[Base inst 1]: igmpProcessGroupRec
Process group rec BLK_OLD_SRCS received on host 10.0.0.17 for group 239.255.1.1
in mode INCLUDE. Num srcs 1"

1114 2017/06/22 16:21:04.08 CEST MINOR: DEBUG #2001 Base IGMP[Base inst 1]
"IGMP[Base inst 1]: igmpProcessIfSrcTimerExp
Source Timer expired for IGMP host 10.0.0.17 (192.168.4.2,239.255.1.1)"

1115 2017/06/22 16:21:04.08 CEST MINOR: DEBUG #2001 Base IGMP[Base inst 1]
"IGMP[Base inst 1]: igmpIfSrcDel
Deleting i/f source entry for host 10.0.0.17 (192.168.4.2,239.255.1.1) from IGMP
Database. DeleteFromAvl: 1 Redir 0"

1116 2017/06/22 16:21:04.08 CEST MINOR: DEBUG #2001 Base IGMP[Base inst 1]
"IGMP[Base inst 1]: igmpIfGroupDel
Deleting 239.255.1.1 from IGMP host 10.0.0.17 database"
```

The following debug information is seen when MCS removes the entry on the standby BNG.

```
debug
  router
    igmp
      mcs "grp-int-1"
    exit
  exit
exit
```

```
1117 2017/06/22 16:21:04.08 CEST MINOR: DEBUG #2001 Base IGMP[Base inst 1]
"IGMP[Base inst 1]: igmpMcsDelIfGroup
Building MCS entry for host 10.0.0.17, group 239.255.1.1"

1118 2017/06/22 16:21:04.08 CEST MINOR: DEBUG #2001 Base IGMP[Base inst 1]
"IGMP[Base inst 1]: igmpMcsDelIfGroup
Deleting MCS entry for host 10.0.0.17, group 239.255.1.1, Glb"
```

IGMP Control Plane Filters

IGMP control plane filtering can be applied at the router level and/or subscriber level (IGMP-policy). The filter list contains multicast groups (S,G) and is provisioned at the router level in the policy-options context. The filter can be applied either as a black-list or a white-list.

Provision a prefix list for the multicast group (G). The following configuration is an example showing the various options possible for the prefix list. The only one used in this configuration is the prefix 239.255.1.1/32.

```
configure
  router
    policy-options
      begin
        prefix-list "igmp-prefix-list-1"
        prefix 239.255.1.1/32 exact
        prefix 239.255.2.0/24 longer
        prefix 239.255.3.0/24 prefix-length-range 24-25
        prefix 239.255.4.0/24 through 25
      exit
```



```
        commit
    exit
exit
exit
```

A white-list router policy is configured as follows, allowing only the prefix list specified and rejecting everything else. Source-address configuration is also possible for IGMP v3 (S,G). The white-list is used for the demonstration later in this chapter.

```
configure
router
  policy-options
  begin
    policy-statement "igmp-white-list-1"
      entry 10
        from
          group-address "igmp-prefix-list-1"
          source-address 192.168.4.2
        exit
        action accept
      exit
    exit
    default-action reject
  exit
exit
commit
exit
exit
exit
```

An example black-list router policy is configured as follows, denying the prefix list and accepting everything else. Again, source-address configuration is possible for IGMP v3 (S,G). The use of this black-list is not demonstrated.

```
configure
router
  policy-options
  begin
    policy-statement "igmp-black-list-1"
      entry 10
        from
          group-address "igmp-prefix-list-1"
          source-address 192.168.4.2
        exit
        action reject
      exit
    exit
    default-action accept
  exit
exit
commit
exit
exit
exit
```

Hierarchical filter

The filter can be applied in two places. First, at router/group-interface level, this will apply to all subscribers connected to the group interface.

```
configure
router
  igmp
```

```
        group-interface "grp-int-1"
            import "igmp-white-list-1"
            no shutdown
        exit
    no shutdown
exit
exit
exit
```

The group-interface filter takes precedence over the subscriber level filter. After the group-interface applies its filter against the incoming IGMP messages, the individual subscriber defined IGMP filters will be applied to the remaining IGMP messages.

```
configure
    subscriber-mgmt
        igmp-policy "igmp-policy-1" create
        import "igmp-white-list-1"
    exit
exit
exit
```

Use the **debug** command to verify that the policy is working correctly for the host. Group 239.255.1.2 is not in the white-list and so is dropped.

```
debug
    router "Base"
        igmp
            group-interface "grp-int-1"
            host "10.0.0.16"
            packet mode egr-ingr-and-dropped
        exit
    exit
exit
```

```
390 2017/06/22 13:15:27.83 CEST MINOR: DEBUG #2001 Base IGMP[1]
"IGMP[1]: RX-PKT
[000 23:22:51.040] IGMP host 10.0.0.16 V3 PDU: 10.0.0.16 -> 224.0.0.22 pduLen 20
  Type: V3 REPORT maxrespCode 0x0 checksum 0x2751
  Num Group Records: 1
    Group Record 0
      Type: MODE_IS_INCL, AuxDataLen 0, Num Sources 1
      Mcast Addr: 239.255.1.2
      Source Address List
        192.168.4.2
"
```

```
391 2017/06/22 13:15:27.83 CEST MINOR: DEBUG #2001 Base IGMP[Base inst 1]
"IGMP[Base inst 1]: igmpParseV3Report
IGMP V3 policy DROP on host 10.0.0.16, from host 10.0.0.16, grpAddr 239.255.1.2,
srcAddr 192.168.4.2"
"
```

IGMP Data Plane Filters

IGMP data plane filter utilize the ip-filter defined in the sla-profile. Again the filter can be used as a black-list or a white-list.

Configure an ip-filter. The following is an example of a black-list filter.

```
configure
  filter
    ip-filter 1 create
      default-action forward
    entry 1 create
      match
        dst-ip 239.255.1.1/32
      exit
    action
      drop
    exit
  exit
exit
exit
exit
```

Apply the configured ip filter into an sla-profile. Because multicast content is sent toward the subscriber, it is applied to the sla-profile egress.

```
configure
  subscriber-mgmt
    sla-profile "sla-profile-1" create
      egress
        ip-filter 1
      exit
    exit
  exit
exit
exit
```

To view the statistics of the filter applied to the subscribers, use the following command.

```
*A:BNG-1# show filter ip 1 counters

=====
IP Filter
=====
Filter Id       : 1                               Applied       : Yes
Scope          : Template                         Def. Action   : Forward
System filter   : Unchained
Radius Ins Pt   : n/a
CrCtl. Ins Pt   : n/a
RadSh. Ins Pt   : n/a
PccRl. Ins Pt   : n/a
Entries        : 1
Description     : (Not Specified)
-----
Filter Match Criteria : IP
-----
Entry          : 1
Ing. Matches    : 0 pkts
Egr. Matches    : 18103 pkts (1918918 bytes)
=====
*A:BNG-1#
```

Conclusion

Multicast is an essential part of Triple Play Services. The TPSDA solution offers much more than a baseline multicast delivery, it includes individual subscriber awareness and a full state redundancy option. Subscriber awareness allows for the fine tuning of each subscriber's multicast experience and also for troubleshooting on a per subscriber basis. Full state redundancy reduces failover time and ensures high availability of the services offered. This example provides a complete configuration walkthrough of both the IPoE and PPPoE SRRP models.

For operators wanting to further control and restrict individual subscriber's multicast content, ESM has a comprehensive set of both control path filtering and data path filtering.

ESM SLAAC Prefix Assignment via Local Address Server

This chapter provides information about ESM SLAAC prefix assignment via local address server.

Topics in this chapter include:

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

Applicability


The information and configuration in this chapter was based on SR OS Release 13.0.R1. The CLI in the current edition is SR OS Release 16.0.R7 based. Both Internet Protocol over Ethernet (IPoE) and Point-to-Point Protocol over Ethernet (PPPoE) are supported.

Overview

Triple Play Service Delivery Architecture (TPSDA) supports IPv6 address/prefix assignment through Dynamic Host Configuration Protocol (DHCP), Point-to-Point Protocol (PPP), and Stateless Address Auto-Configuration (SLAAC). This chapter provides configuration examples of SLAAC prefix assignment via the local address server.

The network topology shown in [Figure 52: TPSDA Network Topology](#) shows a TPSDA setup. The setup consists of a 7750 SR serving as a Broadband Network Gateway (BNG). The 7450 is used as a Layer 2 switch aggregating all subscriber traffic.

Figure 52: TPSDA Network Topology



There are two methods available for subscriber SLAAC prefix assignment. The first method, not covered in this chapter, is to pre-define a static SLAAC prefix for each subscriber on the BNG, in a Local User Database (LUDB) or via a RADIUS AAA server. With such a configuration, the database would contain hundreds of thousands of /64 SLAAC prefixes, each with their associated host. Every time a subscriber moved to a new location (a new subnet), the allocation of a new prefix within the new subnet would be required, along with a manual update of the database.

The second method, covered in this chapter, is to simplify SLAAC prefix assignment. A local address server is configured to dynamically assign SLAAC prefixes to hosts. Only a SLAAC pool name is obtained from RADIUS or LUDB after a successful subscriber authentication. This pool name is then used to assign a SLAAC prefix to the subscriber, out of the named address pool.

Using a local pool for SLAAC prefix assignment provides the following advantages:

- Reduces the configuration required on the RADIUS server, the local user database (LUDB), and the BNG to a few lines;
- Removes the complexity of managing actual prefixes in a database;
- Reduces configuration errors; for example, accidentally assigning the same prefix to two different subscribers.

The local address server already has tools, logs, and monitoring features for prefix management, such as prefix depletion and subnet migration. Service providers can rely on the local address server to assist in SLAAC prefix assignment.

Configuration

This guide assumes a basic knowledge of ESM.

Different Types of SLAAC Hosts

SLAAC is supported for both PPP and IP over Ethernet hosts. The local address server can be enabled for either or both host types on a group interface level.

PPP SLAAC Hosts

PPP IPv4 hosts rely on IPCP to retrieve an IPv4 address. However, IPv6CP does not assign IPv6 addresses or prefixes to the host. PPP hosts rely on router solicitations (RSs) or DHCPv6 to obtain IPv6 addresses/prefixes.

PPP SLAAC hosts creation requires three configuration steps.

Step 1. Following is a baseline PPP subscriber management configuration on the BNG.

```
# on BNG-1
configure
  service
    ies 1 customer 1 create
      subscriber-interface "sub-int-1" create
        ipv6
          subscriber-prefixes
            prefix 2001:db8::/32 wan-host
          exit
        exit
      exit
    exit
  exit
exit
```

```
exit
group-interface "group-int-1" create
  ipv6
    router-advertisements
      prefix-options
        autonomous
      exit
    no shutdown
  exit
exit
sap 1/1/1:1 create
  sub-sla-mgmt
    def-sub-id use-sap-id
    def-sub-profile "sub-profile-1"
    def-sla-profile "sla-profile-1"
    sub-ident-policy "sub-ident-policy-1"
    multi-sub-sap 10
    no shutdown
  exit
exit
pppoe
  no shutdown
exit
exit
exit
exit
```

Step 2. A DHCPv6 server is used as the local address server to assign SLAAC prefixes. It is possible to reuse the same pool for both DHCPv6 and SLAAC address/prefix assignment. For SLAAC hosts, the keyword **wan-host** is required.

```
# on BNG-1
configure
  router
    dhcp6
      local-dhcp-server "dhcp6-server-1" create
      use-pool-from-client
      pool "pool-v6-1" create
      prefix 2001:db8::/32 wan-host create
      exit
    exit
  no shutdown
  exit
exit
```

Step 3. On the PPP group interface, configure the local address server. Specify that the local address server is to be used for client application ppp-slaac. The server name must match the name configured for the DHCPv6 server (step 2). The DHCPv6 server is reused as the local address server.

```
# on a/BNG1
configure
  service
    ies 1
      subscriber-interface "sub-int-1" create
      group-interface "group-int-1" create
      local-address-assignment
        ipv6
          client-application ppp-slaac ipoe-slaac
          server "dhcp6-server-1"
        exit
      no shutdown
    exit
```

```
        exit
    exit
exit
```

There are two options for supplying a SLAAC pool name for a PPP host: RADIUS and LUDB.

Option 1: During PPP authentication, RADIUS can return the SLAAC pool name attribute along with other subscriber attributes. Note: Remove the user-db configuration from pppoe when using RADIUS authentication. Add an authentication policy to the group interface to enable RADIUS authentication.

```
# on a/BNG1
configure
service
  ies 1
    subscriber-interface "sub-int-1" create
    group-interface "group-int-1" create
    authentication-policy "auth-policy-1"
  exit
exit
exit
```

Then, add the attribute Alc-SLAAC-IPv6-Pool to the RADIUS user database. The following is an example from a freeradius clients file.

```
user_ppp_01 Auth-Type := CHAP,    Cleartext-Password := password
        Alc-SLA-Prof-Str = "sla-profile-1",
        Alc-Subsc-ID-Str = "home-ppp-1",
        Alc-Subsc-Prof-Str = "sub-profile-1",
        Alc-SLAAC-IPv6-Pool = pool-v6-1,
        Alc-PPP-Force-IPv6CP = 1
```

Option 2: During PPP authentication, an LUDB can return the SLAAC pool name attribute along with other subscriber attributes. Note: Remove the authentication policy from the group interface when using LUDB.

First, create an LUDB and add a user to the LUDB. This LUDB is configured with a default host for all PPPoE hosts and returns a default SLAAC pool name.

```
# on a/BNG1
configure
  subscriber-mgmt
    local-user-db "pppoe-ludb-lookup" create
    ppp
      match-list username
      host "default" create
      ipv6-slaac-prefix-pool "pool-v6-1"
      no shutdown
    exit
  exit
  no shutdown
exit
exit
exit
```

Then, refer to this LUDB from the group interface.

```
# on a/BNG1
configure
service
```



```

        ies 1
            subscriber-interface "sub-int-1" create
            group-interface "group-int-1" create
            pppoe
                user-db "pppoe-ludb-lookup"
                no shutdown
            exit
        exit
    exit
exit

```

With the preceding configuration, this group interface supports SLAAC prefix assignment through the local address pool. The following is the result of a PPP host being assigned a SLAAC prefix by the local address server.

```

*A:BNG-1# show service active-subscribers hierarchy

=====
Active Subscribers Hierarchy
=====
-- 1/1/1:1 (sub-profile-1)
|
+-- sap:1/1/1:1 - sla:sla-profile-1
|
+-- PPP-session - mac:00:00:00:11:11:11 - sid:1 - svc:1
|   circuit-id:circuit1
|
+-- 2001:db8::/64 - SLAAC

-----
Number of active subscribers : 1
Flags: (N) = the host or the managed route is in non-forwarding state
=====
*A:BNG-1#

```

In the **show pppoe session**, the IPv6 prefix is from the local address pool and the pool name is from authentication (RADIUS or LUDB).

```

*A:BNG-1# show service id 1 pppoe session detail

=====
PPPoE sessions for svc-id 1
=====
Sap Id          Mac Address      Sid  Up Time      Type
IP/L2TP-Id/Interface-Id      MC-Stdby
-----
1/1/1:1         00:00:00:11:11:11 1      0d 00:02:29  local
02:00:00:FF:FE:11:11:11

LCP State       : Opened
IPCP State      : Closed
IPv6CP State    : Opened
PPP MTU         : 1492
PPP Auth-Protocol : CHAP
PPP User-Name   : sub@domain

Subscriber-interface : sub-int-1
Group-interface     : group-int-1

IP Origin       : none
DNS Origin      : none

```

```

NBNS Origin      : none

Subscriber       : "1/1/1:1"
Sub-Profile-String : ""
SLA-Profile-String : ""
SPI group ID     : (Not Specified)
ANCP-String      : ""
Int-Dest-Id      : ""
App-Profile-String : ""
Category-Map-Name : ""
Acct-Session-Id  : "0217FF0000000C5CB82FA7"
Sap-Session-Index : 1

IP Address       : N/A
Primary DNS      : N/A
Secondary DNS     : N/A
Primary NBNS     : N/A
Secondary NBNS   : N/A
Address-Pool     : N/A

IPv6 Prefix      : 2001:db8::/64
IPv6 Prefix Origin : local-pool
IPv6 Prefix Pool  : "pool-v6-1"
IPv6 Del.Pfx.    : N/A
IPv6 Del.Pfx. Origin : none
IPv6 Del.Pfx. Pool : ""
IPv6 Address     : N/A
IPv6 Address Origin : none
IPv6 Address Pool : ""
Primary IPv6 DNS  : N/A
Secondary IPv6 DNS : N/A
Router adv. policy : N/A

Ignoring DF bit   : false
Radius sub-if prefix : N/A

Circuit-Id       : circuit1
Remote-Id        :

Radius Session-T0 : N/A
Radius Class      :
Radius User-Name  : sub@domain
Logical-Line-Id   :
Service-Name      :

-----
Number of sessions : 1
=====
*A:BNG-1#

```

ICMP6 debugging can be used to show the SLAAC address assignment process.

```

debug
  router
    ip
      icmp6
    exit
  exit

```

```

43 2019/04/18 10:04:55.642 CEST MINOR: DEBUG #2001 Base TIP
"TIP: ICMP6_PKT
ICMP6 egressing on group-int-1 (Base):
fe80::17:ffff:fe00:0 -> ff02::1
Type: Router Advertisement (134)

```

```
Code: No Code (0)
Hop Limit      : 64
Flags          :
Retrans Time   : 0
Def Life Time  : 4500
Reachable Time : 0
Option : Prefix      : 2001:db8::/64
        Flags        : On Link Autoconfig
        Valid Life Time: 86400
        Pref Life Time: 3600
"
```

IPoE SLAAC Hosts

IPoE offers two methods to create an SLAAC host:

1. Triggered by a successful IPv4 host creation
2. Triggered by an RS request

SLAAC Host Creation via IPv4 Host

A successful IPv4 host creation can subsequently trigger the creation of a SLAAC host; this is known as IPoE-linking. The SLAAC prefix for the host must be provided through either RADIUS or LUDB during the IPv4 host authentication.

IPoE SLAAC host creation through IPoE linking requires four steps.

Step 1. Following is a baseline IPoE subscriber management configuration on the BNG.

```
# on BNG-1
configure
service
  ies 1 customer 1 create
    description "BNG-1"
    subscriber-interface "sub-int-1" create
      address 10.255.255.253/8
      ipv6
        subscriber-prefixes
          prefix 2001:db8::/32 wan-host
        exit
      exit
    group-interface "group-int-1" create
      dhcp
        server 192.168.0.1
        lease-populate 10
        client-applications dhcp
        gi-address 10.255.255.253
        no shutdown
      exit
      ipv6
        router-advertisements
          prefix-options
            autonomous
          exit
          no shutdown
        exit
      exit
    sap 1/1/1:1 create
```

```

sub-sla-mgmt
  def-sub-id use-sap-id
  ---snip---
  multi-sub-sap 10
  no shutdown
exit
exit
exit
exit

```

Step 2. Enable IPoE-linking to allow SLAAC host creation after a successful IPv4 host creation. Several options should be enabled for the SLAAC host to function. Gratuitous router advertisement will send unsolicited router advertisements with a SLAAC prefix for the host to use. The BNG uses the gratuitous router advertisement to let the subscriber know the assigned prefix to auto-configure. In this case, where prefixes are dynamically assigned, the subscriber will not know the prefix ahead of time, so the gratuitous router advertisement must be enabled. Shared-circuit-id will allow the SLAAC host to use the same circuit ID as the IPv4 host.

```

# on BNG-1
configure
  service
    ies 1
      subscriber-interface "sub-int-1"
      group-interface "group-int-1"
      ipoe-linking
      shared-circuit-id
      gratuitous-rtr-adv
      no shutdown
    exit

```

Step 3. As with PPP hosts, the DHCPv6 server is reused as the local address server for SLAAC prefix assignment. It is possible to reuse the same pool for both DHCPv6 and SLAAC subscribers. For SLAAC hosts, the keyword wan-host is required. In this case, an IPv4 host must be created first to trigger the creation of the IPv6 SLAAC host. The following example uses the local DHCPv4 server for IPv4 address assignment, but it is possible to use other methods for IPv4 address assignment, such as through LUDB and RADIUS proxy.

```

configure
  router
    dhcp
      local-dhcp-server "dhcp-server-1" create
      use-gi-address scope pool
      pool "pool-v4-1" create
      subnet 10.0.0.0/8 create
      options
        subnet-mask 255.0.0.0
        default-router 10.255.255.253
      exit
      address-range 10.0.0.10 10.0.0.254
    exit
  exit
  no shutdown
exit
dhcp6
  local-dhcp-server "dhcp6-server-1" create
  use-pool-from-client
  pool "pool-v6-1" create
  prefix 2001:db8::/32 wan-host create
exit

```

```

        exit
      no shutdown
    exit
  exit
exit

```

Step 4. On the group interface, configure the local address server. Specify that the local address-server is to be used for client application ipoe-slaac. The server name must match the name configured for the DHCPv6 server (Step 3). The local address server reuses the local DHCPv6 server.

```

configure
service
  ies 1
    subscriber-interface "sub-int-1" create
    group-interface "group-int-1" create
    local-address-assignment
      ipv6
        client-application ipoe-slaac
        server "dhcp6-server-1"
    exit
  no shutdown
exit

```

There are two options for supplying a SLAAC pool name for the DHCPv4 host: RADIUS and LUDB.

Option 1: During authentication, RADIUS can return the SLAAC pool name attribute along with other subscriber attributes. Note: Remove the user-db configuration from the DHCP and IPOE-session context when using RADIUS authentication. First, add an authentication policy to the group interface to allow RADIUS authentication.

```

# on a/BNG1
configure
service
  ies 1
    subscriber-interface "sub-int-1" create
    group-interface "group-int-1" create
    authentication-policy "auth-policy-1"
  exit
exit
exit

```

Then, add the attribute Alc-SLAAC-IPv6-Pool to the subscriber host RADIUS user database. The following is an example using the client file on freeradius.

```

00:00:10:10:12:13 Cleartext-Password := password
Alc-SLA-Prof-Str = "sla-profile-1",
Alc-Subsc-ID-Str = "home-ipoe-1",
Alc-Subsc-Prof-Str = "sub-profile-1",
Alc-SLAAC-IPv6-Pool = pool-v6-1

```

Option 2: During authentication, LUDB can return the SLAAC pool name attribute along with other subscriber attributes. Note: Remove the authentication policy from the group interface when using LUDB.

First, create an LUDB and add a user in the LUDB. This LUDB is configured with a default host for all DHCPv4 hosts and returns a default SLAAC pool name.

```

# on a/BNG1
configure

```

```

subscriber-mgmt
  local-user-db "ipoe-ludb-lookup" create
  ipoe
    match-list sap-id
    host "default" create
    ipv6-slaac-prefix-pool "pool-v6-1"
    no shutdown
  exit
exit
no shutdown
exit
exit
exit

```

Then, refer to this LUDB from the group interface. The LUDB can be referred to in two places.

Nokia recommends that IPoE subscribers use an IPoE session. In this case, the LUDB is referenced from the group interface ipoe-session context.

```

# on a/BNG1
configure
  service
    ies 1
      subscriber-interface "sub-int-1" create
      group-interface "group-int-1" create
      ipoe-session
        user-db "ipoe-ludb-lookup"
        no shutdown
      exit
    exit
  exit
exit

```

For operators that do not use IPoE sessions (not recommended), the LUDB is referenced from the group interface dhcp context.

```

# on a/BNG1
configure
  service
    ies 1
      subscriber-interface "sub-int-1" create
      group-interface "group-int-1" create
      dhcp
        user-db "ipoe-ludb-lookup"
        no shutdown
      exit
    exit
  exit
exit

```

With the preceding configuration, the local address server on the group interface is ready to assign SLAAC prefixes. Start a DHCPv4 session to the group interface SAP.

```

*A:BNG-1# show service active-subscribers hierarchy
=====
Active Subscribers Hierarchy
=====
-- 1/1/1:1 (sub-profile-1)
  |
  +-- sap:1/1/1:1 - sla:sla-profile-1
    |

```

```

+-- IPoE-session - mac:00:00:00:22:22:22 - svc:1
|
|-- 10.0.0.10 - DHCP
|
+-- 2001:db8::/64 - SLAAC

-----
Number of active subscribers : 1
Flags: (N) = the host or the managed route is in non-forwarding state
=====
*A:BNG-1#

```

In the show ipoe session, the IPv6 prefix origin is from the local address pool and the pool name is from authentication (RADIUS or LUDB).

```

*A:BNG-1# show service id 1 ipoe session detail

=====
IPoE sessions for service 1
=====

SAP                               : 1/1/1:1
Mac Address                       : 00:00:00:22:22:22
Circuit-Id                       : 11
Remote-Id                        : AA
Session Key                       : sap-mac

MC-Standby                       : No

Subscriber-interface              : sub-int-1
Group-interface                   : group-int-1

Termination Type                 : local
Up Time                          : 0d 00:07:55
Session Time Left                 : N/A
Last Auth Time                   : 04/18/2019 10:45:43
Min Auth Intvl (left)            : infinite (N/A)
Persistence Key                   : N/A

Subscriber                       : "1/1/1:1"
Sub-Profile-String                : ""
SLA-Profile-String                : ""
SPI group ID                     : (Not Specified)
ANCP-String                      : ""
Int-Dest-Id                      : ""
App-Profile-String               : ""
Category-Map-Name                : ""
Acct-Session-Id                  : "0217FF000000215CB83937"
Sap-Session-Index                : 1

IP Address                       : 10.0.0.10/8
IP Origin                        : DHCP
Primary DNS                      : N/A
Secondary DNS                    : N/A
Primary NBNS                     : N/A
Secondary NBNS                   : N/A
Address-Pool                     : N/A

IPv6 Prefix                      : 2001:db8::/64
IPv6 Prefix Origin               : LclPool
IPv6 Prefix Pool                 : "pool-v6-1"
IPv6 Del.Pfx.                   : N/A
IPv6 Del.Pfx. Origin             : None

```

```
IPv6 Del.Pfx. Pool      : ""
IPv6 Address           : N/A
IPv6 Address Origin    : None
IPv6 Address Pool      : ""
Primary IPv6 DNS        : N/A
Secondary IPv6 DNS      : N/A
Router adv. policy     : N/A
Radius sub-if prefix    : N/A

Radius Session-T0       : N/A
Radius Class            :
Radius User-Name        :

GTP IMSI                :
GTP APN                 : (Not Specified)
-----
Number of sessions : 1
=====
*A:BNG-1#
```

ICMP6 debugging can be used to show the SLAAC address assignment process.

```
debug
  router
    ip
      icmp6
    exit
  exit
exit
```

```
63 2019/04/18 10:45:44.212 CEST MINOR: DEBUG #2001 Base TIP
"TIP: ICMP6_PKT
ICMP6 egressing on group-int-1 (Base):
 fe80::17:ffff:fe00:0 -> ff02::1
Type: Router Advertisement (134)
Code: No Code (0)
  Hop Limit      : 64
  Flags          :
  Retrans Time   : 0
  Def Life Time  : 4500
  Reachable Time: 0
  Option  : Src Link Layer Addr 02:17:01:01:00:01
  Option  : Prefix              : 2001:db8::/64
           : Flags              : On Link Autoconfig
           : Valid Life Time: 86400
           : Pref Life Time: 3600
"
```

SLAAC Host Creation via RS Trigger

An IPv6 SLAAC host can be created through a host router originated solicit message, which removes the dependency of a SLAAC host on successful DHCPv4 host creation.

SLAAC hosts creation via RS trigger requires four configuration steps.

Step 1. The following is a baseline IPoE subscriber management configuration on the BNG.

```
# on BNG-1
configure
```



```

service
  ies 1
    description "BNG-1"
    subscriber-interface "sub-int-1" create
    ipv6
      subscriber-prefixes
        prefix 2001:db8::/32 pd wan-host
      exit
    exit
    group-interface "group-int-1" create
    ipv6
      router-advertisements
        prefix-options
          autonomous
        exit
      no shutdown
    exit
    exit
    sap 1/1/1:1 create
    sub-sla-mgmt
      def-sub-id use-sap-id
      def-sub-profile "sub-profile-1"
      def-sla-profile "sla-profile-1"
      sub-ident-policy "sub-ident-policy-1"
      multi-sub-sap 10
      no shutdown
    exit
  exit
exit
exit
exit

```

Step 2. Enable the group interface to process router solicit messages. There are a few options available for router solicit triggered hosts. The inactivity timer will remove the host if the global unique address of the host is not learned through Neighbor Solicitation (NS), Router Solicitation (RS), or Duplicate Address Detection (DAD) messages within the time specified. The min-auth-interval is the interval that a subscriber must wait before the next router-solicit messages is used for re-authentication. Re-authentication can occur if the first RS was lost, or the BNG/RADIUS system was queued up with requests.

```

# on a/BNG1
configure
  service
    ies 1
      subscriber-interface "sub-int-1" create
      group-interface "group-int-1" create
      ipv6
        router-solicit
          inactivity-timer min 5
          min-auth-interval min 5
          user-db "ipoe-ludb-lookup"
          no shutdown
        exit
      exit
    exit

```

Step 3. A DHCPv6 server is used as the local address server for SLAAC prefix assignment. It is possible to reuse the same pool for both DHCPv6 and SLAAC subscribers. For SLAAC hosts, the keyword **wan-host** is required.

```

# on BNG-1
configure
  router
    dhcp6

```

```

        local-dhcp-server "dhcp6-server-1" create
        use-pool-from-client
        pool "pool-v6-1" create
        prefix 2001:db8::/32 wan-host create
        exit
    exit
    no shutdown
exit
exit

```

Step 4. On the group interface, configure the local address server. Specify that the local address server is to be used for client application ipoe-slaac. The server name must match the name configured for the DHCPv6 server. The local address server reuses the local DHCPv6 server.

```

# on BNG-1
configure
    service
        ies 1 customer 1 create
        subscriber-interface "sub-int-1" create
        group-interface "group-int-1" create
        local-address-assignment
            ipv6
                client-application ppp-slaac ipoe-slaac
                server "dhcp6-server-1"
            exit
        no shutdown
    exit
exit

```

There are two options for supplying the SLAAC pool name for the DHCPv4 host: RADIUS and LUDB.

Option 1: During authentication, RADIUS can return the SLAAC pool name attribute along with other subscriber attributes. Note: Remove the user-db configuration from the router-solicit and ipoe-session when using RADIUS authentication. First, add an authentication policy to the group interface to allow RADIUS authentication.

```

# on a/BNG1
configure
    service
        ies 1
        subscriber-interface "sub-int-1" create
        group-interface "group-int-1" create
        authentication-policy "auth-policy-1"
    exit

```

Then, add the attribute Alc-SLAAC-IPv6-Pool to the subscriber host RADIUS user database.

```

00:00:10:10:12:13 Cleartext-Password := password
    Alc-SLA-Prof-Str = "sla-profile-1",
    Alc-Subsc-ID-Str = "home-ipoe-1",
    Alc-Subsc-Prof-Str = "sub-profile-1",
    Alc-SLAAC-IPv6-Pool = pool-v6-1,

```

Option 2: During authentication, LUDB can return the SLAAC pool name attribute along with other subscriber attributes. Note: Remove the authentication policy from the group interface when using LUDB.

First, create an LUDB and add a user in the LUDB. The LUDB configures a default host for all SLAAC hosts and returns a default SLAAC pool name.

```

# on BNG-1

```

```
configure
  subscriber-mgmt
    local-user-db "ipoe-ludb-lookup" create
    ipoe
      match-list sap-id
      host "default" create
      ipv6-slaac-prefix-pool "pool-v6-1"
      no shutdown
    exit
  exit
  no shutdown
exit
exit
exit
```

Then, refer to this LUDB from the group interface. The LUDB can be referred to in two places.

Nokia recommends that IPE subscribers use an IPE session. In this case, the LUDB is referenced from the group interface ipoe-session context.

```
# on BNG-1
configure
  service
    ies 1
      subscriber-interface "sub-int-1" create
      group-interface "group-int-1" create
      ipoe-session
        user-db "ipoe-ludb-lookup"
        no shutdown
      exit
    exit
  exit
exit
```

For operators that do not enable IPE sessions on the BNG (not recommended), the LUDB can be referred to from the group interface in the router-solicit context.

```
# on a/BNG1
configure
  service
    ies 1
      subscriber-interface "sub-int-1" create
      group-interface "group-int-1" create
      ipv6
        router-solicit
          user-db "ipoe-ludb-lookup"
          no shutdown
        exit
      exit
    exit
  exit
exit
```

With the preceding configuration, the group interface is ready to assign SLAAC prefixes from the local address pool. Let the host trigger a router solicit packet.

```
*A:BNG-1# show service active-subscribers hierarchy
```

```
=====
Active Subscribers Hierarchy
=====
```

```
-- 1/1/1:1 (sub-profile-1)
|
+-- sap:1/1/1:1 - sla:sla-profile-1
|
+-- IPoE-session - mac:00:00:00:33:33:33 - svc:1
|
+-- 2001:db8::/64 - SLAAC

-----
Number of active subscribers : 1
Flags: (N) = the host or the managed route is in non-forwarding state
=====
*A:BNG-1#
```

In the show ipoe session, the IPv6 prefix is from the local address pool and the pool name is from authentication (RADIUS or LUDB).

```
*A:BNG-1# show service id 1 ipoe session detail

=====
IPoE sessions for service 1
=====

SAP                               : 1/1/1:1
Mac Address                       : 00:00:00:33:33:33
Circuit-Id                       :
Remote-Id                        :
Session Key                       : sap-mac

MC-Standby                       : No

Subscriber-interface              : sub-int-1
Group-interface                   : group-int-1

Termination Type                  : local
Up Time                           : 0d 00:02:21
Session Time Left                 : N/A
Last Auth Time                   : 04/18/2019 11:07:52
Min Auth Intvl (left)            : infinite (N/A)
Persistence Key                   : N/A

Subscriber                        : "1/1/1:1"
Sub-Profile-String                : ""
SLA-Profile-String                : ""
SPI group ID                     : (Not Specified)
ANCP-String                      : ""
Int-Dest-Id                      : ""
App-Profile-String                : ""
Category-Map-Name                : ""
Acct-Session-Id                  : "0217FF000000255CB83E67"
Sap-Session-Index                : 1

IP Address                       : N/A
IP Origin                        : None
Primary DNS                      : N/A
Secondary DNS                    : N/A
Primary NBNS                     : N/A
Secondary NBNS                   : N/A
Address-Pool                     : N/A

IPv6 Prefix                      : 2001:db8::/64
IPv6 Prefix Origin                : LclPool
IPv6 Prefix Pool                  : "pool-v6-1"
```

```
IPv6 Del.Pfx.      : N/A
IPv6 Del.Pfx. Origin : None
IPv6 Del.Pfx. Pool  : ""
IPv6 Address       : N/A
IPv6 Address Origin : None
IPv6 Address Pool   : ""
Primary IPv6 DNS    : N/A
Secondary IPv6 DNS  : N/A
Router adv. policy  : N/A
Radius sub-if prefix : N/A

Radius Session-T0   : N/A
Radius Class        :
Radius User-Name     :

GTP IMSI            :
GTP APN              : (Not Specified)
-----
Number of sessions : 1
=====
*A:BNG-1#
```

ICMP6 debugging can be used to show the SLAAC address assignment process.

```
debug
router
ip
icmp6
exit
exit
exit
```

```
71 2019/04/18 11:07:52.372 CEST MINOR: DEBUG #2001 Base TIP
"TIP: ICMP6_PKT
ICMP6 egressing on group-int-1 (Base):
 fe80::17:ffff:fe00:0 -> ff02::1
Type: Router Advertisement (134)
Code: No Code (0)
Hop Limit      : 64
Flags          :
Retrans Time   : 0
Def Life Time  : 4500
Reachable Time : 0
Option  : Src Link Layer Addr 02:17:01:01:00:01
Option  : Prefix                : 2001:db8::/64
          Flags                  : On Link Autoconfig
          Valid Life Time: 86400
          Pref Life Time: 3600
"
```

Conclusion

7750 SR TPSDA offers a variety of address assignment options such as PPP, DHCPv4, DHCPv6, and SLAAC. These options allow service providers to pick the address assignment scheme that best fits their networks. SLAAC address assignment is an essential IPv6 address assignment protocol. Having to assign a static prefix per subscriber host in advance could be a challenge for operators. This chapter provides a complete configuration example of using the local address server to assign prefixes dynamically to IPoE

and PPPoE subscriber hosts. This efficient way to assign SLAAC prefixes to subscribers enables operators to achieve a faster time to market for new IPv6 services.

ESMv4: PPPoE Hosts

This chapter describes advanced IPv4 Enhanced Subscriber Management (ESM) PPPoE host configurations.

Topics in this chapter include:

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

Applicability

This chapter applies to SR OS routers and was written for Release 8.0.R4. The CLI is updated to Release 15.0.R2.

This chapter describes support of PPP termination and aggregation (PTA) hosts. L2TP-hosts are out of the scope of this chapter. Only IPv4 PPPoE hosts are handled in this chapter.

PPPoE hosts are only supported in a Routed CO model (IES or VPRN) using Ethernet SAPs with null, dot1q, or QinQ encapsulation.

Overview

The delivery of services to residential customers encompassing voice, video, and data is covered by Triple Play Service Delivery Architecture (TPSDA).

In the TPSDA, a subscriber is defined as a collection of hosts pertaining to a single access connection (for example, DSL line) and identified by a subscriber identifier. A subscriber host is an end user terminal within the subscriber home (PC, set-top box, home gateway) that is identified in the network with a unique (IP address/MAC address) tuple for IPoE or (PPPoE session ID; MAC address) tuple for PPPoE.

The following host types are distinguished:

Static hosts

- IP-MAC
- IP only

Dynamic hosts

- ARP host
- DHCP host
- PPPoE host

This chapter provides configuration and troubleshooting commands for PPPoE-hosts and will use a local user database (LUDB) for host authentication and ESM (Enhanced Subscriber Management) string assignments.

The IP information in this chapter is retrieved from a Local DHCP server.


The authentication, IP information, and ESM strings can come all from an LUDB, a RADIUS server, a (local) DHCP server, or any combination of them. These combinations are out of the scope of this chapter.

Knowledge of the TPSDA concept is assumed throughout this document.

PPPoE hosts in a routed CO environment

The network topology for a Routed CO environment is displayed in [Figure 53: Routed CO network topology](#).

Figure 53: Routed CO network topology



The following configuration tasks should already be configured and are not detailed or explained in this chapter. See the appropriate user guide.

- Basic service router configuration (system interface, IGP, MPLS, BGP)
- Routed CO service topology: VPRN or IES service with subscriber- and group-interface on BSR-1
- ESM
- Local User Data Base (LUDB)
- Local Dynamic Host Configuration Protocol (DHCP) server

In the Routed CO model, PPPoE hosts can be instantiated in routed services, such as the base router, IES, and VPRN. The configuration section of this chapter focuses on PPPoE hosts instantiated in a VPRN servicelocated in BSR-1 (Routed CO).

Review of the PPPoE protocol

PPPoE, Point-to-Point Protocol over Ethernet, is a network protocol for encapsulating PPP frames inside Ethernet frames. The protocol is described in RFC 2516, *A Method for Transmitting PPP Over Ethernet (PPPoE)*, and is based on RFC 1661, *The Point-to-Point Protocol (PPP)*, which provides a standard method for transporting multi-protocol datagrams over point-to-point links.

PPP has three main components:

- A method for encapsulating multi-protocol datagrams.

- A link control protocol (LCP) for establishing, configuring, and testing the data-link connection.
- A family of network control protocols (NCP) for establishing and configuring different network-layer protocols.

Ethernet networks are packet-based and have no concept of a connection or circuit. By using PPPoE, users can virtually dial from one machine to another over an Ethernet network; establish a point to point connection between them and then transport data packets over the connection.

In a typical wire-line solution with broadband access, PPPoE is used between a client (PC or modem) and a Network Access Server (NAS) (also called Broadband Network Gateway (BNG) or Broadband Service Router (BSR)) through an access node, like a Broadband Service Access Node (BSAN).

PPPoE consists of two phases, the Discovery Stage and the Session Stage.

Discovery stage

The discovery phase offers a stateless client-server model. When the Discovery Stage completes, both peers know the PPPoE SESSION_ID and the peer's Ethernet address, which together uniquely define the PPPoE session. There are four steps in the Discovery Stage:

1. PPPoE Active Discovery Initiation (PADI)

Initiation (Host broadcast) — This broadcast packet is used by the client to search for an active server (BNG/BSR/NAS) providing access to a service.

Additional attributes on the PADI message could be added if a BSAN is situated between the client and the BRAS.

2. PPPoE Active Discovery Offer (PADO)

Access concentrator unicast — If the access server can provide the service it will respond with a unicast PADO to signal the client it may request connectivity.

Multiple servers may respond and the client may choose a server to connect.

3. PPPoE Active Discovery Request (PADR):

Host unicast — After the client receives a PADO it will send a PADR unicast packet to connect to a server.

4. PPPoE Active Discovery Session-Confirmation (PADS)

Access concentrator unicast — A server will respond to the client with this unicast packet to establish the session and provide the session-id. Once the PADS was provided, the Session Stage begins.

Discovery PPPoE Ethernet frames have the ETHER_TYPE field set to the value 0x8863.

PPPoE tags

IANA has set up a registry of PPPoE tag values (16-bit values). PPPoE tag values already in use are specified as reserved values as shown in [Table 7: Reserved PPPoE tags](#). All other tag values between 0 and 65535 are to be assigned by IANA.

Table 7: Reserved PPPoE tags

Tag Value	Tag Name
0 0x0000	End-Of-List
257 0x0101	Service-Name
258 0x0102	AC-Name
259 0x0103	Host-Uniq
260 0x0104	AC-Cookie
261 0x0105	Vendor-Specific
262 0x0106	Credits
263 0x0107	Metrics
264 0x0108	Sequence Number
272 0x0110	Relay-Session-Id
273 0x0111	HURL
274 0x0112	MOTM
288 0x0120	PPP-Max-Payload
289 0x0121	IP_Route_Add
513 0x0201	Service-Name-Error
514 0x0202	AC-System-Error
515 0x0203	Generic-Error

Explanations for some PPPoE tags (RFC 2516) are shown in the PPPoE discovery debug messages:

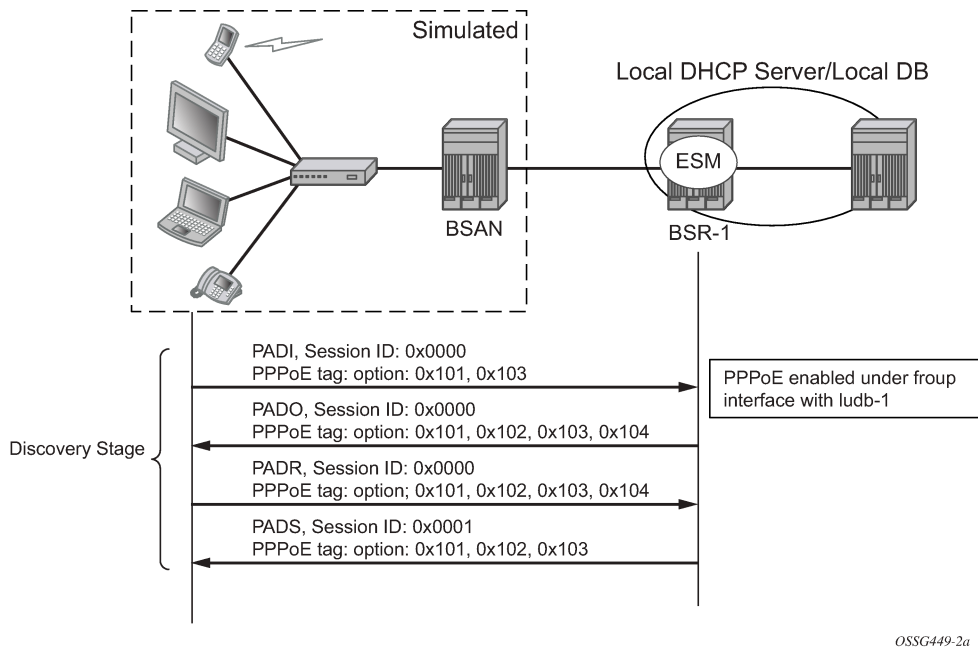
(0x0101) Service-Name — This tag indicates that a service name follows. The tag_value is an UTF-8 string that is not null terminated. When the tag_length is zero, this tag is used to indicate that any service is acceptable. Examples of the use of the **service-name** tag are to indicate an ISP name or a class or quality of service.

(0x0102) AC-Name — This tag indicates that a string follows which uniquely identifies this particular Access Concentrator unit from all others. It may be a combination of trademark, model, and serial id information, or simply an UTF-8 rendition of the MAC address of the box. It is not null terminated.

(0x0103) Host-Uniq — This tag is used by a host to uniquely associate an access concentrator response (PADO or PADS) to a particular host request (PADI or PADR). The tag_value is binary data of any value and length that the host chooses. It is not interpreted by the Access Concentrator. The host may include a host-uniq tag in a PADI or PADR. If the access concentrator receives this tag, it must include the tag unmodified in the associated PADO or PADS response.

(0x0104) AC-Cookie — This tag is used by the access concentrator to aid in protecting against denial of service attacks. The access concentrator may include this tag in a PADO packet. If a host receives this tag, it must return the tag unmodified in the following PADR. The tag_value is binary data of any value and length and is not interpreted by the host.

Figure 54: Discovery stage messages



Session stage

This next stage after Discovery is called the Session Stage. Once the MAC address of the peer is known and a session-id is exchanged, the two end points have all the information needed to start building a point-to-point connection over Ethernet and exchange packets over the connection.

This stage can be divided into the following sections:

- [Setup](#)
- [Maintenance](#)
- [Termination](#)

Setup

PPP Link Control Protocol (LCP)

Both the NAS and the user open the PPP session based on LCP packets. All post-discovery PPPoE Ethernet frames have the ETHER_TYPE field set to the value 0x8864.

The authentication method and the MRU are negotiated during this phase.

RFC 2516 mandates a maximum negotiated Maximum Receive Unit (MRU) of 1492.

RFC 4638, *Accommodating a Maximum Transit Unit/Maximum Receive Unit (MTU/MRU) Greater Than 1492 in the Point-to-Point Protocol over Ethernet (PPPoE)*, relaxes this restriction and allows a maximum negotiated MRU greater than 1492 to minimize fragmentation in next-generation broadband networks.

The SR OS implementation follows RFC 4638 when the client implements these extensions.

LCP uses config_request and config_ack/nack to negotiate parameters:

- LCP goes to final state opened when configure-ack is sent & received.
- The own options are proposed in configure request.

There are three cases for the LCP negotiations parameters:

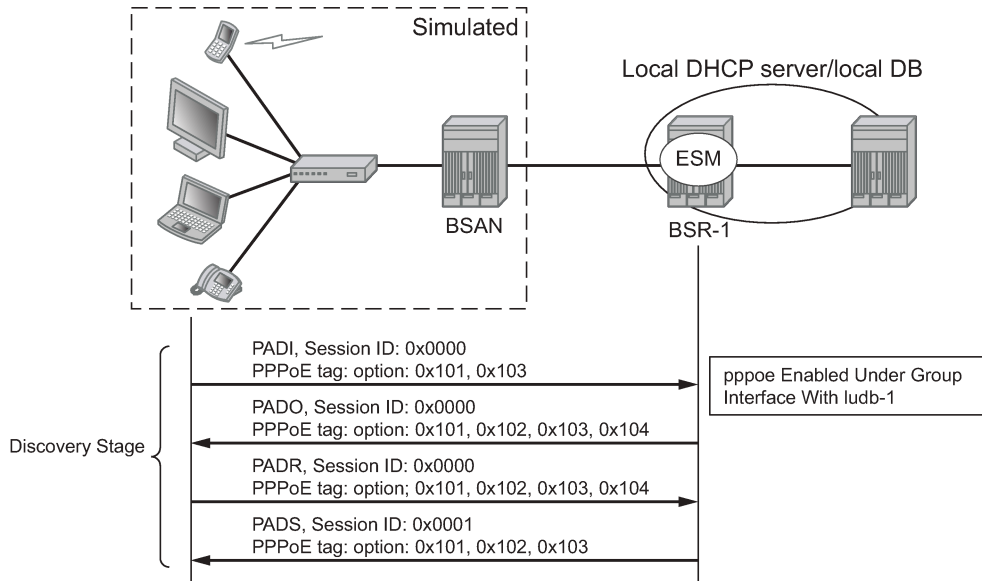
- Peer supports the option and its content.
 - Peer will agree and send config-ack.
- Peer does not support an option
 - Peer will send configure-reject with the option that is not supported.
 - Resend of configure-request without that option.
 - Peer agrees and sends config-ack.
- Peer does support the option, but not the content.
 - Peer will send config-nack with the option and its new content.
 - Resend of configure-request with same options but new content.
 - Peer agrees and sends config-ack.

Table 8: LCP and IPCP code

Code	Packet type
1	Configure-Request
2	Configure-Ack
3	Configure-Nak
4	Configure-Reject
5	Terminate-Request
6	Terminate-Ack
7	Code-Reject
8	Protocol-Reject
9	Echo-Request
10	Echo-Reply
11	Discard-Request
12	Identification

Code	Packet type
13	Time-Remaining
14	Reset-request CCP
15	Reset-Ack CCP

Figure 55: LCP phase messages



OSSG449

Authentication phase

The client authenticates itself through PAP (PPP Password Authentication Protocol) or CHAP (Challenge Handshake Authentication Protocol) to check for access permission. For the CHAP authentication, the BSR initiates the authentication as shown in [Figure 56: CHAP handshaking overview process](#). The password is hashed on the link and plain text in the RADIUS Access-Request message.

The diagram illustrates the Local DHCP server/local DB architecture and the challenge-response protocol flow.

Architecture:

- Simulated:** A group of client devices (phone, laptop, desktop, printer) connected to a central router.
- BSAN:** A Base Station Access Network unit.
- Local DHCP server/local DB:** A system containing an **ESM** (Entity Set Manager) and a **BSR-1** (Base Station Router).

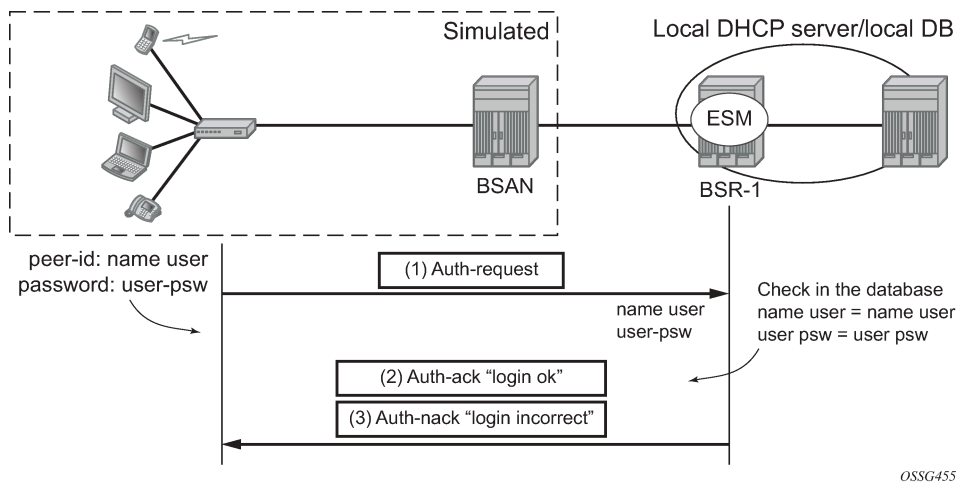
Protocol Flow:

- (1) Challenge:** The BSR-1 sends a challenge to the BSAN. The challenge contains:
 - Name BSR
 - Session id x
 - Random-y
- (2) Response:** The BSAN responds to the BSR-1. The response contains:
 - Hash
 - User-name
 - Sid x
- (3) Success:** The BSR-1 sends a success message to the BSAN.
- (4) Failure:** The BSR-1 sends a failure message to the BSAN.

Processing Details:

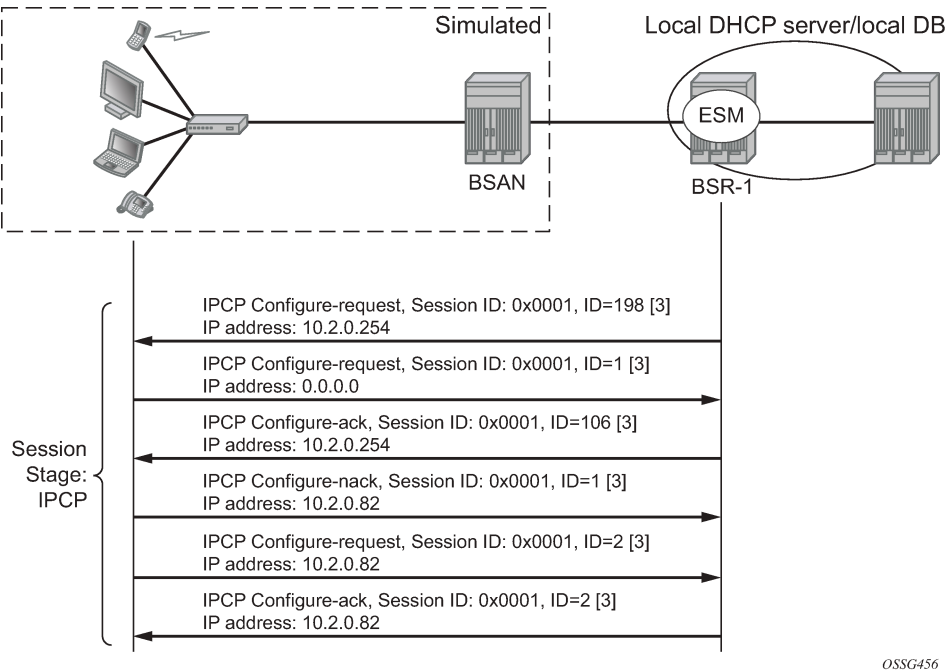
- BSAN Processing:** The BSAN receives the challenge and sends a response. The response is generated by an **MDS** (Message Digest Service) block, which takes **user-psw**, **sid x**, and **Random-y** as input. The response is then hashed using the **Hash** function.
- BSR-1 Processing:** The BSR-1 receives the response and sends a success or failure message. The BSR-1 also performs a **Lookup User-name** operation, which involves hashing the **user-psw**, **Random-y**, and **sid x** using the **MDS** block to produce a **Hash = ? Hash**.

Figure 57: PAP overview process



At this stage, the user requests an IP address to be used for data transmission. During this negotiation, the client will also receive a Domain Name Server (DNS), NBNS (Netbios Name Server) address, etc. if they are requested.

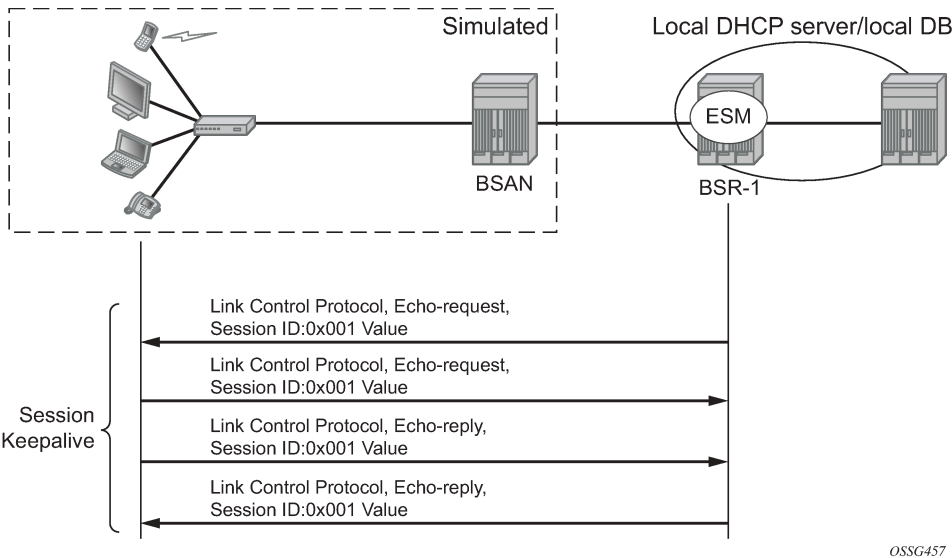
Figure 58: IPCP phase messages



Maintenance

PPP uses keepalives in order to maintain the integrity of the connection. This keepalive mechanism uses an echo-request that is sent to remote PPP peer, following which the remote PPP peer should respond with an echo-reply. The connection is considered down if a number of echo replies are missed. Both sides can initiate keepalives which run independently.

Figure 59: Keepalive messages




Termination

Link termination phase

A PPPoE session can be terminated by either the client or the BRAS and consists of a Terminate-request followed by a PADT.

Figure 60: Link termination phase



Configuration

Session setup, operation, and release

Enable PPPoE termination under the group-interface context.

Enable the local user database under the PPPoE node of the group-interface.

```
configure
service
  vprn 1 customer 1 create
    subscriber-interface "sub-int-1" create
    address 10.2.0.254/16
    group-interface "grp-int-1" create
      sap 1/1/1:1 create
      sub-sla-mgmt
      sub-ident-policy "sub-id-default"
      no shutdown
    exit
  exit
  pppoe
    policy "ppp-policy-1"
    user-db "ludb-1"
    no shutdown
```



```

        exit
    exit
exit
exit

```

The local user database is configured with the following parameters.

```

configure
  subscriber-mgmt
    local-user-db "ludb-1" create
    ppp
      match-list username
      host "user1" create
      host-identification
        username "user1@domain1"
      exit
      address pool "pool-1"
      password chap letmein
      identification-strings 254 create
        subscriber-id "PPPoE-host-user1@domain1"
        sla-profile-string "sla-profile-1"
        sub-profile-string "sub-profile-1"
      exit
      no shutdown
    exit
  ---snip---
exit
exit

```

The PPPoE policy *ppp-policy-1* is defined as follows:

```

*A:BSR-1# configure subscriber-mgmt ppp-policy "ppp-policy-1"
*A:BSR-1>config>subscr-mgmt>ppp-policy# info detail
-----
    no description
    no default-pap-password
    no default-user-name
    disable-cookies
    no force-ppp-mtu-gt-1492
    no ipcp-subnet-negotiation
    keepalive 30 hold-up-multiplier 3
    no pado-ac-name
    pado-delay 30
    no ppp-initial-delay
    no ppp-mtu
    no lcp-ignore-magic-numbers
    max-sessions-per-mac 63
    mlppp
      no accept-mrru
      no endpoint
      no short-sequence-numbers
    exit
    reply-on-padt
    ppp-authentication pref-chap
    ppp-chap-challenge-length min 32 max 64
    no unique-sid-per-sap
    no re-establish-session
    no reject-disabled-ncp
    no session-timeout
    ppp-options
    exit
-----

```

```
*A:BSR-1>config>subscr-mgmt>ppp-policy#
```

The PPPoE policy defines the parameters which are used when establishing the PPPoE session, and includes:

- Disable-cookies — This parameter disables the use of cookies.
- Keepalive — This command defines the keepalive interval and the number of keepalives that can be missed before the session is declared down for this PPPoE policy.
 - [10 — 300] seconds: Specifies the keepalive interval in seconds.
 - hold-up-multiplier [1 — 5]: Specifies the number of keepalives that can be missed.
- PADO-delay — This parameter configures the delay timeout before sending a PPPoE Active Discovery Offer (PADO) packet.
 - [1 — 30] deciseconds
- PPP-mtu — This parameter configures the maximum PPP MTU size.
 - [512 — 9212]: possible values for MTU size.
- Max-sessions-per-mac — This parameter sets the maximum PPPoE sessions that can be opened for the MAC address.
 - [1 — 63]: possible PPPoE sessions per MAC address.
- Reply-on-PADT — Some of the PPPoE clients expect reply on PPPoE Active Discovery Terminate (PADT) message before the context of the session is cleared up. To support such client, a command enabling reply to PADT is provided.
 - [Default] **no reply-on-padt**
- PPP-options — This parameter enables the context to configure PPP options which is not supported by default

These parameters will be explained later in details according to their existence in the respective PPPoE phase.

Multiple PPPoE policies may be configured, but the default policy cannot be modified or deleted.

The PPPoE policy is applied within the PPPoE context under the group interface.

Troubleshooting the PPPoE discovery messages (PADI, PADO, PADR, PADS, and PADT) is done through PPPoE debugging:

```
*A:BSR-1# debug service id 1 ppp packet discovery
- discovery [padi] [pado] [padr] [pads] [padt]
- no discovery

<padi>           : keyword - debug PADI packets
<pado>           : keyword - debug PADO packets
<padr>           : keyword - debug PADR packets
<pads>           : keyword - debug PADS packets
<padt>           : keyword - debug PADT packets
```

```
*A:BSR-1#
```

```
*A:BSR-1# show debug
debug
  service
    id 1
      ppp
```

```
        packet
        mode egr-ingr-and-dropped
        discovery
        ppp
        dhcp-client
    exit
exit
exit
exit
exit
*A:BSR-1#
```

To display the debugging information, a dedicated log should be created:

```
configure
log
    log-id 1
    from debug-trace
    to session
    no shutdown
exit
exit
exit
```

Discovery stage

The following is an example of PPPoE (PADI discovery packet) debug log output:

```
1 2017/05/16 12:07:16.94 CEST MINOR: DEBUG #2001 vprn1 PPPoE
"PPPoE: RX Packet
  VPRN 1, SAP 1/1/1:1

  DMAC: ff:ff:ff:ff:ff:ff
  SMAC: 00:00:67:14:01:02
  Ether Type: 0x8863 (Discovery)

  PPPoE Header:
  Version: 1          Type      : 1
  Code   : 0x09 (PADI) Session-Id: 0x0000 (0)
  Length : 65

  PPPoE Tags:
  [0x0101] Service-Name: "AGILENT"
  [0x0103] Host-Uniq: len = 1, value = 31
  [0x0105] Vendor-Specific: vendor-id = 0x0de9 (ADSL Forum)
    [0x01] Agent-Circuit-Id: "circuit10"
    [0x02] Agent-Remote-Id: "remotel0"
    [0x81] Actual-Upstream: 1024
    [0x82] Actual-Downstream: 16384
    [0x90] Access-Loop-Encap: 01 01 00
"
```

PPPoE policy parameters

Service-Name — The client can ask a particular service. Empty means that any service is acceptable. The service name can indicate an ISP name, class, QoS.

```
1 2017/05/16 12:07:16.94 CEST MINOR: DEBUG #2001 vprn1 PPPoE
"PPPoE: RX Packet
  VPRN 1, SAP 1/1/1:1

  DMAC: ff:ff:ff:ff:ff:ff
  SMAC: 00:00:67:14:01:02
  Ether Type: 0x8863 (Discovery)

  PPPoE Header:
  Version: 1                      Type      : 1
  Code   : 0x09 (PADI)           Session-Id: 0x0000 (0)
  Length : 65

  PPPoE Tags:
  [0x0101] Service-Name: "AGILENT"
  ---snip---
```

The BSR echoes the service name present in the PADI message. Empty means that any service is acceptable.

```
2 2017/05/16 12:07:19.97 CEST MINOR: DEBUG #2001 vprn1 PPPoE
"PPPoE: TX Packet
  VPRN 1, SAP 1/1/1:1

  DMAC: 00:00:67:14:01:02
  SMAC: 14:f2:01:01:00:01
  Ether Type: 0x8863 (Discovery)

  PPPoE Header:
  Version: 1                      Type      : 1
  Code   : 0x07 (PADO)           Session-Id: 0x0000 (0)
  Length : 48

  PPPoE Tags:
  [0x0101] Service-Name: "AGILENT"
  ---snip---
```

Host-Uniq — The host can include a unique tag of any length inserted in PADI or PADR. The AC should echo back this tag in the PADO or PADS.

```
1 2017/05/16 12:07:16.94 CEST MINOR: DEBUG #2001 vprn1 PPPoE
"PPPoE: RX Packet
  VPRN 1, SAP 1/1/1:1

  DMAC: ff:ff:ff:ff:ff:ff
  SMAC: 00:00:67:14:01:02
  Ether Type: 0x8863 (Discovery)

  PPPoE Header:
  Version: 1                      Type      : 1
  Code   : 0x09 (PADI)           Session-Id: 0x0000 (0)
  Length : 65
```

```
PPPoE Tags:
---snip---
[0x0103] Host-Uniq: len = 1, value = 31
---snip---
"
```

Vendor-specific information —The following parameters can optionally be added to the PADI by the PPPoE intermediate agent (BSAN):

- Agent-Circuit-Id
- Agent-Remote-Id
- Access-loop-Encapsulation
- Access loop characteristics (actual-upstream, actual-downstream)

The debug output:

```
1 2017/05/16 12:07:16.94 CEST MINOR: DEBUG #2001 vprn1 PPPoE
"PPPoE: RX Packet
  VPRN 1, SAP 1/1/1:1

  DMAC: ff:ff:ff:ff:ff:ff
  SMAC: 00:00:67:14:01:02
  Ether Type: 0x8863 (Discovery)

  PPPoE Header:
  Version: 1                      Type      : 1
  Code   : 0x09 (PADI)           Session-Id: 0x0000 (0)
  Length : 65

  PPPoE Tags:
  ---snip---
  [0x0105] Vendor-Specific: vendor-id = 0x0de9 (ADSL Forum)
    [0x01] Agent-Circuit-Id: "circuit10"
    [0x02] Agent-Remote-Id: "remote10"
    [0x81] Actual-Upstream: 1024
    [0x82] Actual-Downstream: 16384
    [0x90] Access-Loop-Encap: 01 01 00
"
```

Cookies — The cookies can be seen in the PADO message. This tag of any value and length may be included by the AC and is echoed back by the client to the AC in the next PADR.

```
2 2017/05/16 12:07:19.97 CEST MINOR: DEBUG #2001 vprn1 PPPoE
"PPPoE: TX Packet
  VPRN 1, SAP 1/1/1:1

  DMAC: 00:00:67:14:01:02
  SMAC: 14:f2:01:01:00:01
  Ether Type: 0x8863 (Discovery)

  PPPoE Header:
  Version: 1                      Type      : 1
  Code   : 0x07 (PADO)           Session-Id: 0x0000 (0)
  Length : 18

  PPPoE Tags:
  ---snip---
  [0x0104] ACCookie: len = 16, value = d7 91 ---snip--- ba 74
"
```

When **disable-cookies** is configured, the use of cookies will be disabled, when omitted the **no-disable-cookies** will be used.

```
*A:BSR-1# configure subscriber-mgmt ppp-policy "ppp-policy-1" disable-cookies
```

The cookies are encoded back by the client in the next PADR message.

AC-Name — The string that uniquely identifies the access concentrator (AC).

```
2 2017/05/16 12:07:19.97 CEST MINOR: DEBUG #2001 vprn1 PPPoE
"PPPoE: TX Packet
  VPRN 1, SAP 1/1/1:1

  DMAC: 00:00:67:14:01:02
  SMAC: 14:f2:01:01:00:01
  Ether Type: 0x8863 (Discovery)

  PPPoE Header:
  Version: 1                      Type      : 1
  Code   : 0x07 (PADO)           Session-Id: 0x0000 (0)
  Length : 48

  PPPoE Tags:
  ---snip---
  [0x0102] AC-Name: "BSR-1"
  ---snip---
"
```

PADO-delay — SR OS has the possibility to delay the sending of the PADO message to the client. This feature could be used if the client is dual homed to two BSRs and is explained later in the chapter.

When PADO-delay is configured, the configured value equals the delay timeout before sending PADO, when omitted the PADO-delay value of 0 msec will be used.

```
*A:BSR-1# configure subscriber-mgmt ppp-policy "ppp-policy-1" pado-delay ?
- no pado-delay
- pado-delay <deciseconds>

<deciseconds>      : [1..30]

*A:BSR-1#
```

Authentication — PPPoE hosts are authenticated based on username-password information (PAP/CHAP authentication) or on information embedded in the PADI message in case of PADI authentication.

For CHAP authentication, the **min** and **max** values for the **ppp-chap-challenge-length** are defined when enabling **ppp-chap-challenge-length**. When omitted, a **min** of 32 and **max** of 64 are used.

```
*A:BSR-1# configure subscriber-mgmt ppp-policy "ppp-policy-1" ppp-chap-challenge-length
- ppp-chap-challenge-length min <minimum-length> max <maximum-length>
- no ppp-chap-challenge-length

<minimum-length>   : [8..64]
<maximum-length>   : [8..64]

*A:BSR-1#
```

To complete the discovery phase, the server must provide a session-id to the client and SR OS allocates session-id 1 when the MACs are different.

```
*A:BSR-1# show service id 1 pppoe session
=====
PPPoE sessions for svc-id 1
=====
Sap Id          Mac Address      Sid   Up Time          Type
IP/L2TP-Id/Interface-Id          MC-Stdby
-----
1/1/1:1         00:00:67:14:01:02 1      0d 00:01:32     local
10.2.0.1
-----
Number of sessions : 1
=====
*A:BSR-1#
```



Note:

In the VLAN per service model (N: 1 VLAN), the MACs are the same and the PPPoE interworking is done at the BSAN.

Session stage

LCP

During the link establishment phase, client and server negotiate options and need to come to an agreement on these options. Options that are unknown by the peer are rejected whereas known options with unknown content are nack'd. In the latter case, the peer needs to resend the same option, but with another content. In case of a reject, the peer should remove that option. An Ack will be sent if there is a full agreement.

One of the more important options that is exchanged is the maximum receive unit (MRU) and the authentication protocol that will be used later in the authentication phase. The first option, the MRU value (minus overhead) is sent from the BSR toward the client and is the lowest value between the port MTU and the optional configured ppp-mtu in the ppp-policy.

RFC 2516 mandates a maximum negotiated Maximum Receive Unit (MRU) of 1492, but RFC 4638 relaxes this restriction and allows a maximum negotiated MRU greater than 1492 to minimize fragmentation in next-generation broadband networks. The SR OS implementation follows RFC 4638 when the client implements these extensions.

If a PPPoE client wants to use MRU>1492 in the LCP-config request, it should include the **ppp-max-payload** tag with the higher MTU value in the initial PADI message.

```
*A:BSR-1# configure subscriber-mgmt ppp-policy "ppp-policy-1" ppp-mtu ?
- no ppp-mtu
- ppp-mtu <mtu-bytes>

<mtu-bytes>          : [512..9212]

*A:BSR-1#
```

For LCP, the PPPoE debug output is as follows:

```
5 2017/05/16 12:07:20.12 CEST MINOR: DEBUG #2001 vprn1 PPPoE
```

```
"PPPoE: TX Packet
  VPRN 1, SAP 1/1/1:1

  DMAC: 00:00:67:14:01:02
  SMAC: 14:f2:01:01:00:01
  Ether Type: 0x8864 (Session)

  PPPoE Header:
  Version: 1                      Type      : 1
  Code   : 0x00                  Session-Id: 0x0001 (1)
  Length : 21

  PPP:
  Protocol : 0xc021 (LCP)
  Code     : 1 (Configure-Request)
  Identifier: 239          Length   : 19

  Options:
  [1] MRU: 1492
  ---snip---
```

The second important option, the authentication method used in the authentication phase is exchanged between client and server and can be PAP or CHAP authentication. The authentication method is not exchanged when PADI authentication is done. PADI authentication means that the BSR will authenticate the user based on parameters in the PADI message. Authentication based on PADI and PAP/CHAP is possible.

The debug output for CHAP authentication protocol is as follows:

```
8 2017/05/16 12:07:20.89 CEST MINOR: DEBUG #2001 vprn1 PPPoE
"PPPoE: RX Packet
  VPRN 1, SAP 1/1/1:1

  DMAC: 14:f2:01:01:00:01
  SMAC: 00:00:67:14:01:02
  Ether Type: 0x8864 (Session)

  PPPoE Header:
  Version: 1                      Type      : 1
  Code   : 0x00                  Session-Id: 0x0001 (1)
  Length : 21

  PPP:
  Protocol : 0xc021 (LCP)
  Code     : 2 (Configure-Ack)
  Identifier: 239          Length   : 19

  Options:
  [1] MRU: 1492
  [3] Authentication-Protocol: 0xc223 (CHAP), Algorithm = 5 (MD5)
  [5] Magic-Number: 0x772a29d2
"
```

The debug output for PAP authentication is as follows:

```
72 2017/05/16 14:46:50.68 CEST MINOR: DEBUG #2001 vprn1 PPPoE
"PPPoE: RX Packet
  VPRN 1, SAP 1/1/1:1

  DMAC: 14:f2:01:01:00:01
```



```
SMAC: 00:00:67:13:01:02
Ether Type: 0x8864 (Session)
PPPoE Header:
Version: 1                      Type      : 1
Code   : 0x00                  Session-Id: 0x0001 (1)
Length : 20

PPP:
Protocol : 0xc021 (LCP)
Code     : 2 (Configure-Ack)
Identifier: 134              Length    : 18

Options:
[1] MRU: 1492
[3] Authentication-Protocol: 0xc023 (PAP)
[5] Magic-Number: 0x0538a9d5
"
```

CHAP to PAP fallback case

For user authentication, with pap-chap-access, always try CHAP first; if that doesn't succeed, try PAP.

The option to be used first (CHAP/PAP) is defined when enabling the ppp-authentication. When omitted, CHAP is always preferred.

```
*A:BSR-1# configure subscriber-mgmt ppp-policy "ppp-policy-1" ppp-authentication ?
- no ppp-authentication
- ppp-authentication {pap|chap|pref-chap|pref-pap}

*A:BSR-1#
```

PPPoE clients that implement undocumented options also require an agreement on those unknown options. By default, the 7750 SR will reject unknown options, but the **ppp-options** feature in the **pppoe-policy** allows for support of undocumented client LCP or IPCP/IPv6CP options. If the received LCP or IPCP/IPv6CP option matches the configured options in the pppoe-policy, an ack will be sent instead of a reject.

```
*A:BSR-1# configure subscriber-mgmt ppp-policy "ppp-policy-1"
                                     ppp-options custom-option ?
- custom-option <protocol> <option-number> address <ip-address>
- custom-option <protocol> <option-number> hex <hex-string>
- custom-option <protocol> <option-number> string <ascii-string>
- no custom-option <protocol> <option-number>

<protocol>          : lcp|ipcp|ipv6cp
<option-number>     : [0..255]
<ip-address>        : a.b.c.d
<ascii-string>      : [127 chars max]
<hex-string>        : [0x0..0xFFFFF...(max 254 hex nibbles)]

*A:BSR-1#
```

Troubleshooting the PPPoE LCP session messages is done with PPPoE debugging:

```
A:BSR-1# debug service id 1 ppp packet ppp ?
- no ppp
- ppp [lcp] [pap] [chap] [ipcp] [ipv6cp]

<lcp>                : keyword - debug LCP packets
<pap>                : keyword - debug PAP packets
<chap>              : keyword - debug CHAP packets
```

```
<ipcp>          : keyword - debug IPCP packets
<ipv6cp>        : keyword - debug IPv6CP packets
```

A:BSR-1#

At the end of the LCP phase, authentication is started.

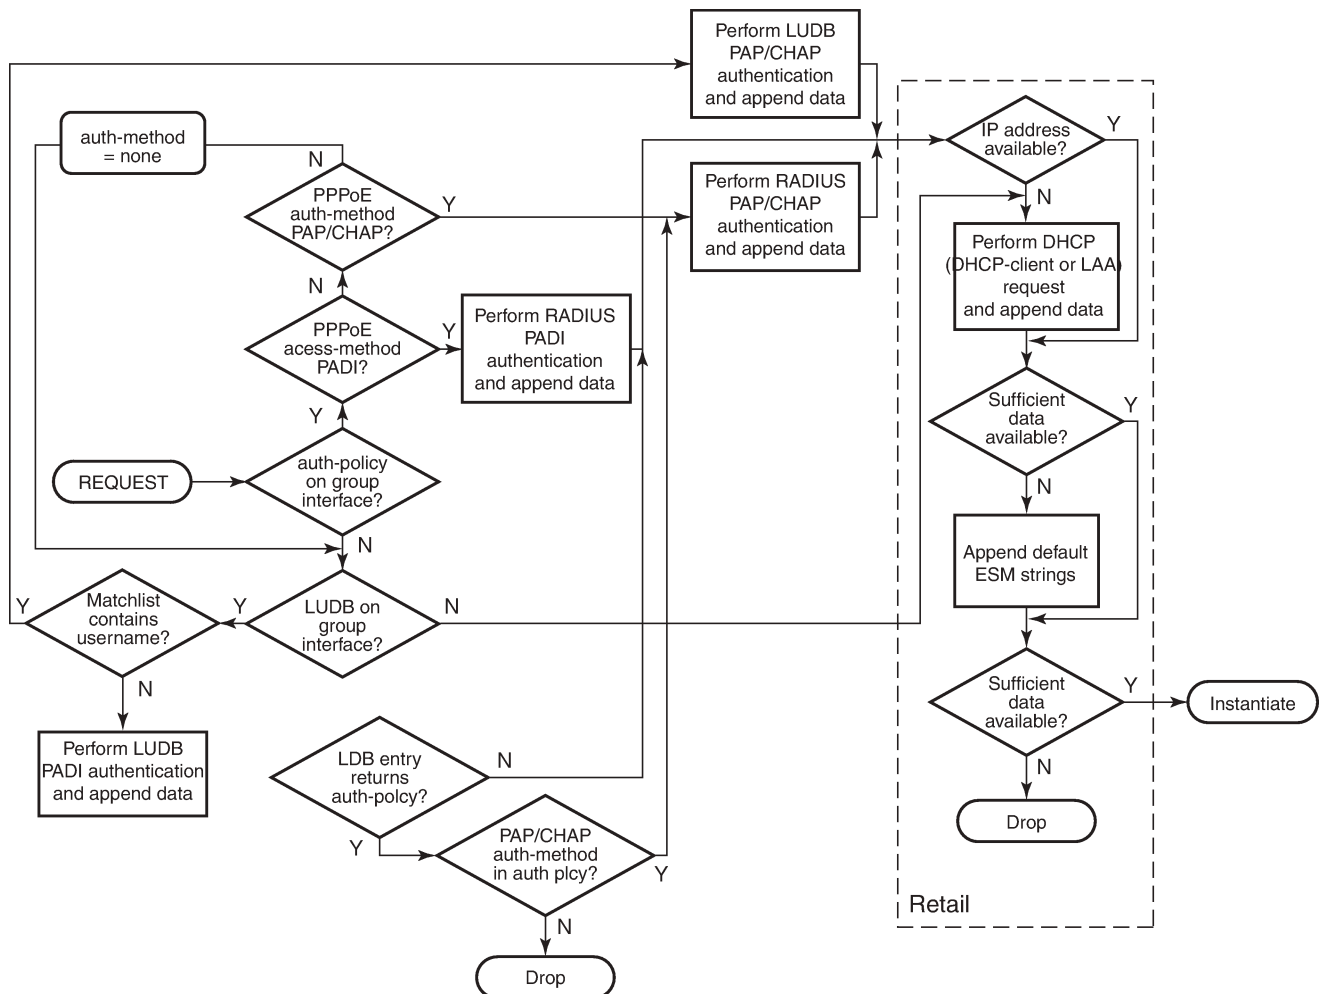
SR OS supports three main methods for PPPoE authentication.

- PADI or PAP/CHAP authentication through RADIUS
- PADI or PAP/CHAP authentication through a LUDB
- PADI authentication via LUDB then PAP/CHAP pre-authentication through RADIUS

DHCP server authentication will not be explained in this section because this is more authorization than authentication.

The flow chart of the PPPoE host authentication process is shown in [Figure 61: Authentication flow chart](#) and starts with the request to the middle left.

Figure 61: Authentication flow chart



OSSG460

PPPoE users get authenticated via the LUDB if this LUDB is configured under the group-interface, as follows:

```
configure
  service
    vprn 1
      subscriber-interface "sub-int-1" create
      group-interface "grp-int-1" create
      pppoe
        user-db "ludb-1"
```

Users get authenticated via RADIUS if an authentication-policy is configured under the same group-interface. RADIUS has precedence if both are configured, as follows:

```
configure
  service
    vprn 1
      subscriber-interface "sub-int-1" create
      group-interface "grp-int-1" create
      authentication-policy "auth-1"
      pppoe
        user-db "ludb-1"
        no shutdown
      exit
    exit
```

Users that get authenticated via the LUDB can still go to RADIUS if the authentication-policy is moved from the group-interface to the LUDB.

```
configure
  subscriber-mgmt
    local-user-db "ludb-1" create
    ppp
      match-list username
      ---snip---
      host "user2" create
      host-identification
        username "user2@domain1"
      exit
      auth-policy "auth-1"
      address pool "pool-1"
      password pap letmein
      identification-strings 254 create
        subscriber-id "PPPoE-host-user2@domain1"
        sla-profile-string "sla-profile-1"
        sub-profile-string "sub-profile-1"
      exit
      no shutdown
    exit
  exit
  exit
  exit
exit
```

This last mechanism could be used to pick up parameters like **pado-delay** or to check some variables such as **circuit-id**, **remote-id** from the LUDB during discovery phase and subsequently to use RADIUS for PAP/CHAP authentication.

```
*A:BSR-1# configure subscriber-mgmt local-user-db "ludb-1" ppp host "user2"
```

```

host-identification ?
- host-identification

[no] circuit-id      - Configure the circuit id of this host
[no] derived-id     - Configure the host ID to be derived by a python script from
                    DHCP packets during a DHCP transaction
[no] encap-tag-range - Configure the SAP encap range tags for this host
[no] mac            - Configure the MAC address of this host
[no] remote-id      - Configure the remote id of this host
[no] sap-id         - Configure the SAP identifier of this host
[no] service-name   - Configure the service name of this host
[no] username       - Configure the user name of this host

*A:BSR-1#

```

Both RADIUS and LUDB support PADI or PAP/CHAP authentication.

PPPoE users that are authenticated through the LUDB and have in the LUDB a match-list other than **username** will get authenticated based on PADI parameters like mac, circuit-id, remote-id.

```

*A:BSR-1# configure subscriber-mgmt local-user-db "ludb-1" ppp match-list ?
- no match-list
- match-list <ppp-match-type-1> [<ppp-match-type-2>...(up to 3 max)]

<ppp-match-type>      : circuit-id|derived-id|mac|remote-id|sap-id|
                    encap-tag-range|service-name|username

*A:BSR-1#

```

PPPoE users that have in the LUDB a match-list equal to **username** will use the PAP/CHAP authentication method.

```

configure
  subscriber-mgmt
    local-user-db "ludb-1" create
      ppp
        match-list username
        host "user1" create
          host-identification
            username "user1@domain1"
          exit
          address pool "pool-1"
          password chap letmein
          identification-strings 254 create
            subscriber-id "PPPoE-host-user1@domain1"
            sla-profile-string "sla-profile-1"
            sub-profile-string "sub-profile-1"
          exit
          no shutdown
        exit
      ---snip---
    exit
  exit
exit

```

PPPoE users that are authenticated through RADIUS and have in the authentication policy, a pppoe-access-method equal to PADI will use the **mac** or **circuit-id** information from the PADI in their request to RADIUS.

```

*A:BSR-1# configure subscriber-mgmt authentication-policy "auth-1"

```

```

                                     pppoe-access-method ?
- no pppoe-access-method
- pppoe-access-method {none|padi|pap-chap}

*A:BSR-1#

```

The selection for mac or circuit-id can be altered via the parameter user-name-format.

```

*A:BSR-1# configure subscriber-mgmt authentication-policy "auth-1"
                                     user-name-format ?
- no user-name-format
- user-name-format <format> [mac-format <mac-format>]
- user-name-format <format> append [<domain-name>]
                                     [mac-format <mac-format>]
- user-name-format <format> append domain-name
- user-name-format <format> default-domain <domain-name>
                                     [mac-format <mac-format>]
- user-name-format <format> replace <domain-name>
                                     [mac-format <mac-format>]
- user-name-format <format> strip [mac-format <mac-format>]

<format>                : ascii-converted-circuit-id|ascii-converted-tuple|
                           circuit-id|dhcp-client-vendor-opts|mac|
                           mac-giaddr|ppp-user-name|tuple
<domain-name>           : max 128 chars, no @ needed
<mac-format>            : (only when format is dhcp-client-vendor-opts)
                           like ab:   for 00:0c:f1:99:85:b8
                           or XY-    for 00-0C-F1-99-85-B8
                           or mmmm.  for 0002.03aa.abff
                           or xx     for 000cf19985b8

*A:BSR-1#

```

PPPoE users that are authenticated through RADIUS and having a pppoe-access-method equal to pap-chap in the authentication policy use the username from the authentication phase in their request to RADIUS. The parameter user-name-format is irrelevant in this last case.

RADIUS authentication

When authentication is provided through RADIUS, two methods can be used to authenticate the PPPoE session.

- PADI authentication
- PAP/CHAP authentication

The RADIUS policy specifies which parameters are provided in the RADIUS access-request message.

The following parameters can be configured:

- Circuit-id: Provided through the PADI/PADR PPPoE relay vendor specific tag as specified in TR-101 of the DSL Forum.
- Remote-id: Provided through the PADI/PADR PPPoE relay vendor specific tag as specified in TR-101 of the DSL Forum.
- NAS-port-id: SAP ID on which the PPPoE session terminates (e.g. 1/1/3:1).
- NAS-identifier: System name of the NAS or BNG.
- PPPoE-service-name: Provided through the PPPoE PADI packet.

- access-loop-options: Provided through the PAD/PADR PPPoE extensions as specified in TR-101 of the DSL Forum.

The option to use PADI authentication or PAP/CHAP authentication is selected with the following configuration in the RADIUS policy:

```
*A:BSR-1# configure subscriber-mgmt authentication-policy "auth-1" pppoe-access-method ?
- no pppoe-access-method
- pppoe-access-method {none|padi|pap-chap}

*A:BSR-1#
```

When PADI authentication is used the MAC address or the PPPoE relay tag (Circuit-ID) or a combination of MAC address and circuit ID can be used to identify the subscriber in the RADIUS server.

The attributes included in the RADIUS Access-Request message is configured in the RADIUS policy using the following command:

```
*A:BSR-1# configure subscriber-mgmt authentication-policy "auth-" include-radius-attribute ?
- include-radius-attribute
- no include-radius-attribute

[no] access-loop-op* - Enable/disable generation of the DSL Forum access loop
characteristics RADIUS attributes
[no] acct-session-id - Enable/disable generation of the Acct-Session-Id RADIUS
attribute
[no] called-station* - Enable/disable generation of the called-station-id RADIUS
attribute
[no] calling-statio* - Enable/disable generation of the calling-station-id RADIUS
attribute
[no] circuit-id      - Enable/disable generation of the agent-circuit-id RADIUS
attribute
[no] dhcp-options   - Enable/disable generation of the dhcp-options RADIUS attribute
[no] dhcp-vendor-cl* - Enable/disable generation of the dhcp-vendor-class-id RADIUS
attribute
[no] dhcp6-options  - Enable/disable generation of the dhcp6-options RADIUS attribute
[no] mac-address     - Enable/disable generation of the client MAC address RADIUS
attribute
[no] nas-identifier  - Enable/disable generation of the NAS-Identifier RADIUS
attribute
[no] nas-port        - Enable/disable include of the NAS-Port attribute
[no] nas-port-id     - Enable/disable generation of the NAS-Port-Id RADIUS attribute
[no] nas-port-type   - Enable/disable generation of the NAS-Port-Type RADIUS attribute
[no] pppoe-service-* - Enable/disable generation of the pppoe-service-name RADIUS
attribute
[no] remote-id       - Enable/disable generation of the agent-remote-id RADIUS
attribute
[no] sap-session-in* - Enable/disable generation of the per-SAP unique session index
[no] tunnel-server-* - Enable/disable generation of the tunnel-server-attrs RADIUS
attribute
[no] wifi-num-attac* - Enable/disable including the Alc-Num-Attached-UEs RADIUS
attribute
[no] wifi-ssid-vlan  - Enable/disable including the per-SSID VLAN ID in Alc-Wlan
-SSID-VLAN

*A:BSR-1#
```

Local User Database authentication

A second authentication option for PPPoE termination is to use the local user database of the BSR 7750 for PAP/CHAP authentication (this option will be used in this example). With this authentication method, the client's PPPoE session is authenticated locally on the BSR without any constraint of an external radius server.

The local user database is configured with the following parameters:

```
configure
  subscriber-mgmt
    local-user-db "ludb-1" create
    ppp
      match-list username
      host "user1" create
      host-identification
        username "user1@domain1"
      exit
      address pool "pool-1"
      password chap letmein
      identification-strings 254 create
        subscriber-id "PPPoE-host-user1@domain1"
        sla-profile-string "sla-profile-1"
        sub-profile-string "sub-profile-1"
      exit
      no shutdown
    exit
  exit
  ---snip---
```

With the LUDB authentication method, authentication can be provided by the username/password combination.

To enable the local authentication method, the local user database is configured under the PPPoE node of the group-interface as follows.

```
configure
  service
    vprn 1
      subscriber-interface "sub-int-1" create
      group-interface "grp-int-1" create
      pppoe
        policy "ppp-policy-1"
        user-db "ludb-1"
        no shutdown
      exit
    exit
  exit
```

The properties of LUDB user user-1 are as follows:

```
*A:BSR-1# show subscriber-mgmt local-user-db "ludb-1" ppp-host "user1"

=====
PPP Host "user1"
=====
Admin State      : Up
Last Mgmt Change : 05/16/2017 11:57:57

Host Identification
Mac Address      : N/A
```

```
Circuit Id      : N/A
Remote Id      : N/A
Sap Id         : N/A
Service Name    : N/A
User Name      : user1@domain1
Encap Tag Range : N/A
Derived Id     : N/A

Matched Objects : userName

Address         : pool "pool-1"
Password Type   : CHAP
PAD0 Delay      : 0msec
Pre Auth Policy : N/A
Auth Policy     : N/A
Padi Auth Policy : N/A
---snip---

DHCPv6 lease times
Renew timer     : > 9999 days
Rebind timer    : > 9999 days
Preferred lifetime : 0d 00:00:00
Valid lifetime  : 0d 00:00:00

Identification Strings (option 254)
Subscriber Id   : PPPoE-host-user1@domain1
SLA Profile String : sla-profile-1
Sub Profile String : sub-profile-1
App Profile String : N/A
ANCP String     : N/A
Inter Destination Id: N/A
Category Map Name : N/A

---snip---

Access loop info
Circuit ID format : none
Circuit ID       : N/A
Remote ID format  : none
Remote ID        : N/A
=====
*A:BSR-1#
```

To debug the LUDB, use the command as follows:

```
*A:BSR-1# debug subscriber-mgmt local-user-db "ludb-1" detail ?
- detail {all|failed}
- no detail

*A:BSR-1#
```

See the LUBD Basics chapter for further information.

DHCP client authentication

A third authentication method for PPPoE termination is to perform PPPoE to DHCP transformation (where the router acts as a DHCP client on behalf of the PPP session) and to use a DHCP server for session authentication. This method is useful when a similar authentication is used for DHCP based clients.

The PPPoE to DHCP authentication method can provide authentication on the basis of MAC address, circuit ID or remote ID.

IPCP

IP and DNS information can be obtained from different sources like LUDB and RADIUS for fixed IP addressing or (local) DHCP for dynamic IP pool management.

If IP information is returned from a DHCP server, then the PPPoE options such as the DNS name are retrieved from the DHCP ACK and provided to the PPPoE client.

Local DHCP server

This chapter uses a local DHCP server as a source for the IP address information of the PPPoE host.

```
configure
  service
    vprn 1 customer 1 create
    dhcp
      local-dhcp-server "server-1" create
      use-pool-from-client
      pool "pool-1" create
      subnet 10.2.0.0/16 create
      exclude-addresses 10.2.0.254 10.2.0.255
      address-range 10.2.0.1 10.2.0.253
      exit
    exit
  no shutdown
  exit
exit
exit
exit
exit
exit
```

To check the DHCP server summary:

```
*A:BSR-1# show router 1 dhcp local-dhcp-server "server-1" summary
=====
DHCP server server-1 router 1
=====
Admin State           : inService
Operational State     : inService
Persistency State     : shutdown
User Data Base        : N/A
Use gateway IP address : disabled
Use pool from client  : enabled

---snip---

-----
Pool name : pool-1
-----
Failover Admin State   : outOfService
Failover Oper State    : shutdown
Failover Persist Key   : N/A
Administrative MCLT    : 0h10m0s
Operational MCLT       : 0h10m0s
Startup wait time      : 0h2m0s
```

```

Partner down delay      : 23h59m59s
Ignore MCLT             : disabled
-----
Subnet                  Free      %      Stable  Declined Offered  Rem-pend Drain
-----
10.2.0.0/16             252      99%    1         0         0         0         N
Totals for pool         252      99%    1         0         0         0
-----
Totals for server       252      99%    1         0         0         0
-----

Interface associations
Interface                Admin
-----
local-dhcp-server-1      Up
-----

Local Address Assignment associations
Group interface          Admin
-----
=====
*A:BSR-1#

```

To debug the DHCP server:

```

debug
  router "1"
    local-dhcp-server "server-1"
    detail-level high
    mode egr-ingr-and-dropped
  exit
exit
exit

```

Keepalive

The keepalive timer (defined in seconds) and the hold-up-multiplier are defined when enabling keepalives. When omitted, a 30 second keepalive timer and three hold-up-multipliers are used.

```

*A:BSR-1# configure subscriber-mgmt ppp-policy "ppp-policy-1" keepalive ?
- keepalive <seconds> [hold-up-multiplier <multiplier>]
- no keepalive

<seconds>          : [4..300]
<multiplier>       : [1..5]

*A:BSR-1#

```

The keepalive defines the interval between LCP Echo Requests in seconds. The hold-up-multiplier defines the number of missed replies before the PPPoE session is considered dead.

To check the keepalive statistics:

```

*A:BSR-1# show service id 1 pppoe session session-id 1 mac 00:00:67:14:01:02
                                           statistics
=====
PPPoE sessions for svc-id 1
=====

```

Sap Id	Mac Address	Sid	Up Time	Type
IP/L2TP-Id/Interface-Id				MC-Stdby
1/1/1:1	00:00:67:14:01:02	1	0d 00:05:45	local
10.2.0.4				
Packet Type	Received	Transmitted		
LCP Configure-Request	1	1		
LCP Configure-Ack	1	1		
LCP Configure-Nak	0	0		
LCP Configure-Reject	0	0		
LCP Terminate-Request	0	0		
LCP Terminate-Ack	0	0		
LCP Code-Reject	0	0		
LCP Echo-Request	35	0		
LCP Echo-Reply	0	35		
LCP Protocol-Reject	0	0		
LCP Discard-Request	0	0		
PAP Authenticate-Request	0	-		
---snip---				
CHAP Challenge	-	1		
CHAP Response	1	-		
---snip---				
IPCP Configure-Request	2	1		
IPCP Configure-Ack	1	1		
---snip---				
IPv6CP Configure-Request	0	0		
IPv6CP Configure-Ack	0	0		
---snip---				
Unknown Protocol	0	-		
Number of sessions	: 1			
=====				
*A:BSR-1#				

The BSR supports an optimized implementation of keepalive mechanism; this is a mechanism where client and/or server can check the aliveness of the peer. This LCP echo-request is sent on expiration of a timer, derived from the configured **pppoe-policy keepalive** value.

An LCP echo reply is returned to the client after a LCP echo request is received and the preceding timer on the BSR is reset to the initial keepalive value.

The preceding mechanism results in an optimized mechanism if the keepalive timers from the client are smaller than the configured values on the BSR.

The client or BSR will terminate the session with a PADT if no LCP echo-reply is received within the time specified by the hold-up-multiplier.

An example for Echo Request from BSR and Echo Reply from the PPPoE host is as follows:

```
15:36:27.907363 00:00:67:14:01:02 > 14:f2:01:01:00:01, ethertype 802.1Q (0x8100),
length 64: vlan 1, p 7, ethertype PPPoE S, PPPoE [ses 0x1] LCP (0xc021),
length 10: LCP, Echo-Request (0x09), id 249, length 10
        encoded length 8 (=Option(s) length 4)
        0x0000: c021 09f9 0008
        Magic-Num 0x00000000
15:36:27.907844 14:f2:01:01:00:01 > 00:00:67:14:01:02, ethertype 802.1Q (0x8100),
```

```
length 60: vlan 1, p 5, ethertype PPPoE S, PPPoE [ses 0x1] LCP (0xc021),
length 10: LCP, Echo-Reply (0x0a), id 249, length 10
  encoded length 8 (=Option(s) length 4)
    0x0000: c021 0af9 0008
    Magic-Num 0x3298fadc
```

To check the PPPoE session for a particular service, use the **show service id <service-id> pppoe session** command. Detailed output as well as additional output filtering is available:

```
*A:BSR-1# show service id 1 pppoe session ?
- session [interface <ip-int-name|ip-address> | sap <sap-id>]
  [type <pppoe-session-type>] [session-id <session-id>] [mac <ieee-address>]
  [ip-address <ip-prefix[/prefix-length]>] [port <port-id>]
  [no-inter-dest-id | inter-dest-id <intermediate-destination-id>]
  [steering-profile <steering-profile>] [router-advertisement-policy
  <router-adv-policy>] [detail|statistics]
- session l2tp-connection-id <connection-id> [detail|statistics]
```

The details of the PPPoE session for MAC address 00:00:67:14:01:02 are as follows:

```
*A:BSR-1# show service id 1 pppoe session mac 00:00:67:14:01:02 detail

=====
PPPoE sessions for svc-id 1
=====
```

Sap Id	Mac Address	Sid	Up Time	Type
IP/L2TP-Id/Interface-Id				MC-Stdby
1/1/1:1 10.2.0.4	00:00:67:14:01:02	1	0d 00:14:05	local

```

LCP State           : Opened
IPCP State          : Opened
IPv6CP State        : Closed
PPP MTU             : 1492
PPP Auth-Protocol   : CHAP
PPP User-Name       : user1@domain1

Subscriber-interface : sub-int-1
Group-interface     : grp-int-1

IP Origin           : dhcp
DNS Origin          : none
NBNS Origin         : none

Subscriber          : "PPPoE-host-user1@domain1"
Sub-Profile-String  : "sub-profile-1"
SLA-Profile-String  : "sla-profile-1"
ANCP-String         : ""
Int-Dest-Id        : ""
App-Profile-String  : ""
Category-Map-Name   : ""
Acct-Session-Id     : "14F2FF00000009591AFDFA"
Sap-Session-Index   : 1

IP Address          : 10.2.0.4/32
Primary DNS         : N/A
Secondary DNS       : N/A
Primary NBNS        : N/A
Secondary NBNS      : N/A
Address-Pool        : pool-1
```

```
IPv6 Prefix      : N/A
IPv6 Prefix Origin : none
IPv6 Prefix Pool  : ""
IPv6 Del.Pfx.    : N/A
IPv6 Del.Pfx. Origin : none
IPv6 Del.Pfx. Pool : ""
IPv6 Address     : N/A
IPv6 Address Origin : none
IPv6 Address Pool : ""
Primary IPv6 DNS  : N/A
Secondary IPv6 DNS : N/A
Router adv. policy : N/A

Ignoring DF bit    : false
Radius sub-if prefix : N/A

Circuit-Id       : circuit10
Remote-Id        : remotel0

Radius Session-T0 : N/A
Radius Class      :
Radius User-Name  : user1@domain1
Logical-Line-Id   :
Service-Name      :
Data link        : ethernet
Encaps 1         : untagged-ethernet
Encaps 2         : not-available
Origin           : tags
Link Rate Down   : 16384
Rate Origin      : tags
-----
Number of sessions : 1
=====
*A:BSR-1#
```

An event will be generated when a PPPoE host has been created in the system.

```
*A:BSR-1# show log log-id 99

=====
Event Log 99
=====
Description : Default System Log
Memory Log contents [size=500  next event=75  (not wrapped)]

73 2017/05/16 15:26:18.19 CEST WARNING: SVCNMR #2500 Base Subscriber created
"Subscriber PPPoE-host-user1@domain1 has been created in the system"

---snip---

*A:BSR-1#
```

The PPPoE host will appear in the subscriber-host table for the service with origin set to IPCP.

```
*A:BSR-1# show service id 1 subscriber-hosts

=====
Subscriber Host table
=====
Sap      Subscriber
IP Address
MAC Address  PPPoE-SID Origin  Fwding State
-----
```

```
1/1/1:1          PPPoE-host-user1@domain1
10.2.0.4
00:00:67:14:01:02  1          IPCP          Fwding
-----
Number of subscriber hosts : 1
=====
*A:BSR-1#
```

A host route (/32) for its IP address is inserted in the routing table toward the appropriate group-interface.

```
*A:BSR-1# show router 1 route-table

=====
Route Table (Service: 1)
=====
Dest Prefix[Flags]          Type  Proto  Age      Pref
Next Hop[Interface Name]    Metric
-----
10.2.0.0/16                 Local  Local   03h47m06s  0
sub-int-1                   0
10.2.0.4/32                 Remote Sub Mgmt 00h18m50s 0
[grp-int-1]                  0
172.16.0.1/32               Local  Local   03h47m12s  0
local-dhcp-server-1         0
-----
No. of Routes: 3
Flags: n = Number of times nexthop is repeated
      B = BGP backup route available
      L = LFA nexthop available
      S = Sticky ECMP requested
=====
*A:BSR-1#
```

To advertise the PPPoE host subnets to other protocol/network, a policy statement should be defined with using **from protocol direct**.

```
configure
router
  policy-options
  begin
    policy-statement "policy-1"
    entry 10
    from
      protocol direct
    exit
    ---snip---
  exit
exit
commit
exit
exit
exit
```

Terminate

Some PPPoE clients expect a reply on PADT message before the context of the session is cleared up. To support such clients, a command enabling reply to PADT is configured.

When reply-on-padt is configured, the BSR will reply with PADT message, when omitted, no PADT will be sent from the BSR as a reply on the client's PADT.

```
*A:BSR-1# configure subscriber-mgmt ppp-policy "ppp-policy-1" reply-on-padt
```

The PPPoE debug output:

```
116 2017/05/16 15:47:16.67 CEST MINOR: DEBUG #2001 vprn1 PPPoE
"PPPoE: RX Packet
  VPRN 1, SAP 1/1/1:1

  DMAC: 14:f2:01:01:00:01
  SMAC: 00:00:67:14:01:02
  Ether Type: 0x8863 (Discovery)

  PPPoE Header:
  Version: 1                      Type      : 1
  Code   : 0xa7 (PADT)           Session-Id: 0x0001 (1)
  Length : 0
"
```

```
117 2017/05/16 15:47:16.68 CEST MINOR: DEBUG #2001 vprn1 PPPoE
"PPPoE: TX Packet
  VPRN 1, SAP 1/1/1:1

  DMAC: 00:00:67:14:01:02
  SMAC: 14:f2:01:01:00:01
  Ether Type: 0x8863 (Discovery)

  PPPoE Header:
  Version: 1                      Type      : 1
  Code   : 0xa7 (PADT)           Session-Id: 0x0001 (1)
  Length : 0
"
```

A PPPoE host can be manually deleted from the system using following clear command:

```
*A:BSR-1# clear service id 1 ppp session ?
- session [sap-id <sap-id>] [interface <ip-int-name|ip-address>]
  [mac <ieee-address>] [session-id <session-id>]
  [type <pppoe-session-type>] [ip-address <ip-prefix[/prefix-length]>]
  [port <port-id>] [inter-dest-id <intermediate-destination-id>]
  [no-inter-dest-id] [user-name <user-name>]
  [sub-ppp-type {oa|oe|oeoa|ol2tp}] [no-padt]
- session all [no-padt]

---snip---

*A:BSR-1#
```

The logs indicate a cause for terminating the PPP session:

- **Admin Reset** — Use the **clear** command or a RADIUS Disconnect Request.
TERMINATE CAUSE [49] 4 Admin Reset(6)
- **User Request** — User disconnects the session.
TERMINATE CAUSE [49] 4 User Request(1)
- **Accounting OFF**
 - When accounting policy has been removed from sap/interface/sub-profile.

- The VPRN service which is transporting accounting information has been shutdown.
- The last RADIUS accounting server has been removed from an already applied accounting policy.

TERMINATE CAUSE [49] 4 NAS Request(10)

- **PPPoE keepalive timeout**

TERMINATE CAUSE [49] 4 Lost Carrier(2)

- **RADIUS session timeout**

PPPoE hosts advanced topics

QoS Aspects

VLAN encapsulated downstream PPPoE control traffic is generated by default with dot1p value 7. This value can be overruled with the following commands:

In case of the PPPoE hosts instantiated in the Base routing instance using an IES service.

```
*A:BSR-1# configure router sgt-qos application pppoe dot1p 5
```

In case of the PPPoE hosts instantiated in a VPRN service subscriber-interface.

```
*A:BSR-1# configure service vprn 1 sgt-qos application pppoe dot1p 5
```

The **show router sgt-qos** command displays the configured and default DSCP and default dot1p values per application. Because PPPoE is a Layer 2 protocol we will see only the dot1p settings. The default dot1p value **none** corresponds with value 7.

```
*A:BSR-1# show router 1 sgt-qos application pppoe
```

```
=====
Dot1p Application Values
=====
Application          Configured Dot1p Value      Default Dot1p Value
-----
pppoe                 5                            7
=====
*A:BSR-1#
```

Limiting the number of PPPoE hosts

The maximum number of PPPoE sessions can be controlled by the parameters session-limit, sap-session-limit, host-limit, multi-sub-sap limit and max-sessions-per mac.

session-limit — The maximum number of PPPoE sessions for an IES/VPRN group-interface is defined when enabling session-limit. When omitted, a single PPPoE session is allowed.

```
configure
  service
    vprn 1 customer 1 create
      subscriber-interface "sub-int-1" create
```



```
group-interface "grp-int-1" create
  pppoe
    policy "ppp-policy-1"
    session-limit 1
    user-db "ludb-1"
    no shutdown
  exit
exit
```



Note:

The discovery phase is completed before the session-limit check is executed.

The debug log shows a message when the session-limit is reached, as follows:

```
199 2017/05/16 16:03:22.14 CEST MINOR: DEBUG #2001 vprn1 PPPoE
"PPPoE: Dropped Packet
  VPRN 1, SAP 1/1/1:1

  Problem: Reached the interface session limit (1) for "grp-int-1"

  DMAC: ff:ff:ff:ff:ff:ff
  SMAC: 00:00:67:13:01:02
  Ether Type: 0x8863 (Discovery)

  PPPoE Header:
  Version: 1                      Type      : 1
  Code   : 0x09 (PADI)           Session-Id: 0x0000 (0)
  Length : 37

  PPPoE Tags:
  [0x0101] Service-Name: ""
  [0x0103] Host-Uniq: len = 1, value = 32
  [0x0105] Vendor-Specific: vendor-id = 0x0de9 (ADSL Forum)
    [0x01] Agent-Circuit-Id: "circuit11"
    [0x02] Agent-Remote-Id: "remote11"
"
```

Also a trap is generated, as follows:

```
*A:BSR-1# show log log-id 99

=====
Event Log 99
=====
Description : Default System Log
Memory Log contents [size=500  next event=88  (not wrapped)]

87 2017/05/16 16:03:22.15 CEST WARNING: PPPoE #2001 vprn1 PPPoE session failure
"PPPoE session failure on SAP 1/1/1:1 in service 1 - Reached the interface session
  limit (1) for "grp-int-1""

---snip---

*A:BSR-1#
```

Sap-session-limit

The maximum number of PPPoE sessions per SAP for an IES/VP RN group-interface is defined when enabling sap-session-limit. When omitted, a single PPPoE session per SAP is allowed:

```
configure
service
  vprn 1 customer 1 create
    subscriber-interface "sub-int-1" create
    group-interface "grp-int-1" create
      pppoe
        policy "ppp-policy-1"
        session-limit 10
        sap-session-limit 1
        user-db "ludb-1"
        no shutdown
      exit
    exit
  exit
exit
exit
exit
exit
exit
```

The debug log shows a message when the session-limit is reached, as follows:

```
201 2017/05/16 16:10:34.00 CEST MINOR: DEBUG #2001 vprn1 PPPoE
"PPPoE: Dropped Packet
  VPRN 1, SAP 1/1/1:1

  Problem: Reached the per-SAP session limit (1) for "grp-int-1"

  DMAC: ff:ff:ff:ff:ff:ff
  SMAC: 00:00:67:13:01:02
  Ether Type: 0x8863 (Discovery)

  PPPoE Header:
  Version: 1                      Type      : 1
  Code   : 0x09 (PADI)           Session-Id: 0x0000 (0)
  Length : 37

  PPPoE Tags:
  [0x0101] Service-Name: ""
  [0x0103] Host-Uniq: len = 1, value = 32
  [0x0105] Vendor-Specific: vendor-id = 0x0de9 (ADSL Forum)
    [0x01] Agent-Circuit-Id: "circuit11"
    [0x02] Agent-Remote-Id: "remote11"
"
```

A trap is generated when trying to instantiate a new PPPoE session while the configured number of the sessions per sap is reached.

```
*A:BSR-1# show log log-id 99

=====
Event Log 99
=====
Description : Default System Log
Memory Log contents [size=500  next event=89  (not wrapped)]

88 2017/05/16 16:10:34.00 CEST WARNING: PPPoE #2001 vprn1 PPPoE session failure
```

```
"PPPoE session failure on SAP 1/1/1:1 in service 1 - Reached the per-SAP session limit  
(1) for "grp-int-1"
```

```
---snip---
```

```
*A:BSR-1#
```

Max-sessions-per-mac

The BSR 7750 implementation defines a unique PPPoE session based on the PPPoE SESSION_ID and the client's MAC address.

The maximum number of PPPoE sessions per mac is defined when enabling max-sessions-per-mac. When omitted, a single PPPoE session per mac is allowed.

```
configure
  subscriber-mgmt
    ppp-policy "ppp-policy-1"
      disable-cookies
      max-sessions-per-mac 256
      reply-on-padt
    exit
  exit
exit
```

Although the command is max-session-per-mac, actually it means the maximum number of supported sessions-per-MAC-per-SAP especially in N: 1 VLAN model.

The debug log shows a message when the max-sessions-per-mac limit is reached, as follows:

```
226 2017/05/16 16:30:38.82 CEST MINOR: DEBUG #2001 vprn1 PPPoE
"PPPoE: Dropped Packet
  VPRN 1, SAP 1/1/1:1

  Problem: Reached the maximum number (1) of PPPoE sessions for MAC
  00:00:67:13:01:02

  DMAC: ff:ff:ff:ff:ff:ff
  SMAC: 00:00:67:13:01:02
  Ether Type: 0x8863 (Discovery)

  PPPoE Header:
  Version: 1                Type      : 1
  Code   : 0x09 (PADI)      Session-Id: 0x0000 (0)
  Length : 37
  ---snip---
"
```

A trap is generated when trying to instantiate a new PPPoE session using the same MAC while the configured number of max-sessions-per-mac is reached.

```
*A:BSR-1# show log log-id 99
```

```
=====
Event Log 99
=====
```

```
Description : Default System Log
Memory Log contents [size=500  next event=92  (not wrapped)]
```

```
91 2017/05/16 16:30:38.82 CEST WARNING: PPPoE #2001 vprn1 PPPoE session failure
"PPPoE session failure on SAP 1/1/1:1 in service 1 - Reached the maximum number (1)
of PPPoE sessions for MAC 00:00:67:13:01:02"

---snip---

*A:BSR-1#
```

Host-limit

The maximum number of PPPoE hosts is defined when enabling host-limit. When omitted, a single host is allowed.

```
configure
  subscriber-mgmt
    sla-profile "sla-profile-1"
      host-limits
        overall 10
      exit
    exit
  exit
exit
```

If the configured host-limit is reached for a subscriber, access is denied for a new host, and an event is generated.

```
*A:BSR-1# show log log-id 99

=====
Event Log 99
=====
Description : Default System Log
Memory Log contents [size=500  next event=110  (not wrapped)]

108 2017/05/16 16:38:30.68 CEST WARNING: PPPoE #2001 vprn1 PPPoE session failure
"PPPoE session failure on SAP 1/1/1:1 in service 1 -
[00:00:67:13:01:02,6,user3@domain1] sla-profile sla-profile-1 : host-limit overall
(1) exceeded for subscriber PPPoE-host-user3@domain1 on SAP 1/1/1:1 "

---snip---

*A:BSR-1#
```

An optional command **remove-oldest** can be specified. In this case, the new host is accepted and the old one will be removed.

```
configure
  subscriber-mgmt
    sla-profile "sla-profile-1"
      host-limits
        overall 10
        remove-oldest
      exit
    exit
  exit
exit
```

Multi-sub-sap

This parameter defines the maximum number of subscribers (dynamic and static) that can be simultaneously active on this SAP.

When omitted, a single PPPoE session per sap is allowed (no multi-sub-sap).

```
configure
service
  vprn 1
    subscriber-interface "sub-int-1"
      group-interface "grp-int-1"
        sap 1/1/1:1
          sub-sla-mgmt
            multi-sub-sap 100
          ---snip---
        exit
      exit
    exit
  exit
```

A trap is generated when trying to instantiate a new PPPoE session while the configured number of the multi-sub-sap is reached.

```
*A:BSR-1# show log log-id 99

=====
Event Log 99
=====
Description : Default System Log
Memory Log contents [size=500  next event=127  (not wrapped)]

125 2017/05/16 16:43:57.28 CEST WARNING: PPPOE #2001 vprn1 PPPoE session failure
"PPPoE session failure on SAP 1/1/1:1 in service 1 -
 [00:00:67:13:01:02,1,user3@domain1] Number of subscribers exceeds the configured
 multi-sub-sap limit (1)"

---snip---

*A:BSR-1#
```

Redundancy

Redundancy for PPPoE sessions can be used for load balancing sessions between the two BSRs. PADO-delays (which can come from RADIUS, LUDB, and policy) are used to achieve that.

The redundant BSRs need different IP subnets, and upon failure the PPP sessions will need to be re-established.


Because PADI messages are broadcast on a multi-access network, all BSRs on that network will reply with a PADO to the initiator.

The PADR and PADS are sent in unicast to the MAC address from the first received PADO message.

In order to allow control over the NAS/BSR selection for a PPPoE session, SR OS offers the ability to delay the PADO message. Due to the fact that PPPoE clients select the NAS/BSR for further communication based on the first PADO message that arrives, this functionality provides control over the NAS/BSR that gets selected for a PPPoE session.

On top if for some reason a NAS/BSR, without PADO delay configured in the PPPoE policy, does not reply on PADI messages to the client, another NAS/BSR with a PADO delay configured will reply based on the time configured and ultimately the PPPoE session will be established with the PADO delayed NAS/BSR.

Figure 62: Pado-Delay Scenario



Check the PADO delay value, as follows:

```
*A:BSR-1# show subscriber-mgmt ppp-policy "ppp-policy-1"

=====
PPP Policy "ppp-policy-1"
=====
Description           : (Not Specified)
Last Mgmt Change      : 05/16/2017 16:37:04
PPP-mtu               : N/A
Keepalive Interval    : 30s
Disable AC-Cookies    : Yes
Max Sessions-Per-Mac  : 2
Allow Same CID        : No
PPP-Authentication    : pref-CHAP
PPP-Init-Delay (ms)   : 0
Unique SIDs-Per-SAP   : disabled
Ignore-Magic-Num      : No
PADO AC-Name          : (Not Specified)
Default username      : (Not Specified)
Default password      : (Not Specified)

Force PPP-mtu >1492   : No
Keepalive Multiplier  : 3
PADO Delay            : 3000msec
Reply-On-PADT        : Yes
Re-establish Session  : Disabled
PPP-CHAP Challenge    : 32 - 64
IPCP negotiate subnet : No
Reject-Disabled-Ncp   : No
Session Timeout       : unlimited

-----
PPP Custom Options
-----
Protocol Number Value
-----
No options configured.

-----
MLPPP
-----
Accept MRRU           : false
Request short sequence nr. : false
Endpoint class        : null
Endpoint address       : (Not Specified)
-----
```

```
*A:BSR-1#
```

Another option to achieve redundancy is through Multi Chassis Synchronization (MCS), but that is beyond the scope of this chapter.

High availability

The PPPoE session state is HA: the session state is synchronized to the standby CPM. When the active CPM fails, all PPPoE hosts stay active without service interruption.

Conclusion

This chapter provides configuration and troubleshooting commands for PPPoE hosts in a Layer 3 Routed CO (IES/VPRN subscriber interface) context.

ESMv6: IPoE Dual Stack Hosts

This chapter describes IPoE dual stack hosts for ESMv6 configurations.

Topics in this chapter include:

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

Applicability

This chapter describes ESMv6: IPoE dual stack hosts and is based on SR OS Release 8.0.R4. The CLI is updated to Release 15.0.R1.

This chapter focuses on IPoE IPv6. IPv4 configuration is shown for completeness and is described in more detail in [IPv4 DHCP Hosts](#).

Prerequisites

Configuring IPoE dual stack hosts for ESMv6 are dependent on the following.

- Routed CO (IES/VPDN service) with Enhanced Subscriber Management (ESM)
- Routed Gateway (RG) in the home

Overview

In this chapter, the configuration, operation, and troubleshooting of IPoE dual stack hosts in a routed home gateway environment is described. The focus is on the Enhanced Subscriber Management for IPv6 (ESMv6) part where DHCPv6 is used for IPv6 address assignment. In the Broadband Network Gateway (BNG), authentication, authorization, and IPv6 prefix configuration for an IPoE IPv6 host can be done by a local user database (LUDB) or RADIUS.


IPoE Dual Stack Hosts

An IPoE dual stack subscriber may support both IPv4 and IPv6 simultaneously. The dual stack hosts share a common subscriber identification policy and have a common SLA- and Subscriber-profile.

IPoE IPv4 and IPv6 hosts operate independently because they are set up through different protocols, DHCPv4 and DHCPv6 respectively.

For a stateful IPoE dual stack subscriber, up to three different types of subscriber hosts can be instantiated.

Figure 63: Stateful IPoE Dual Stack Subscriber Hosts

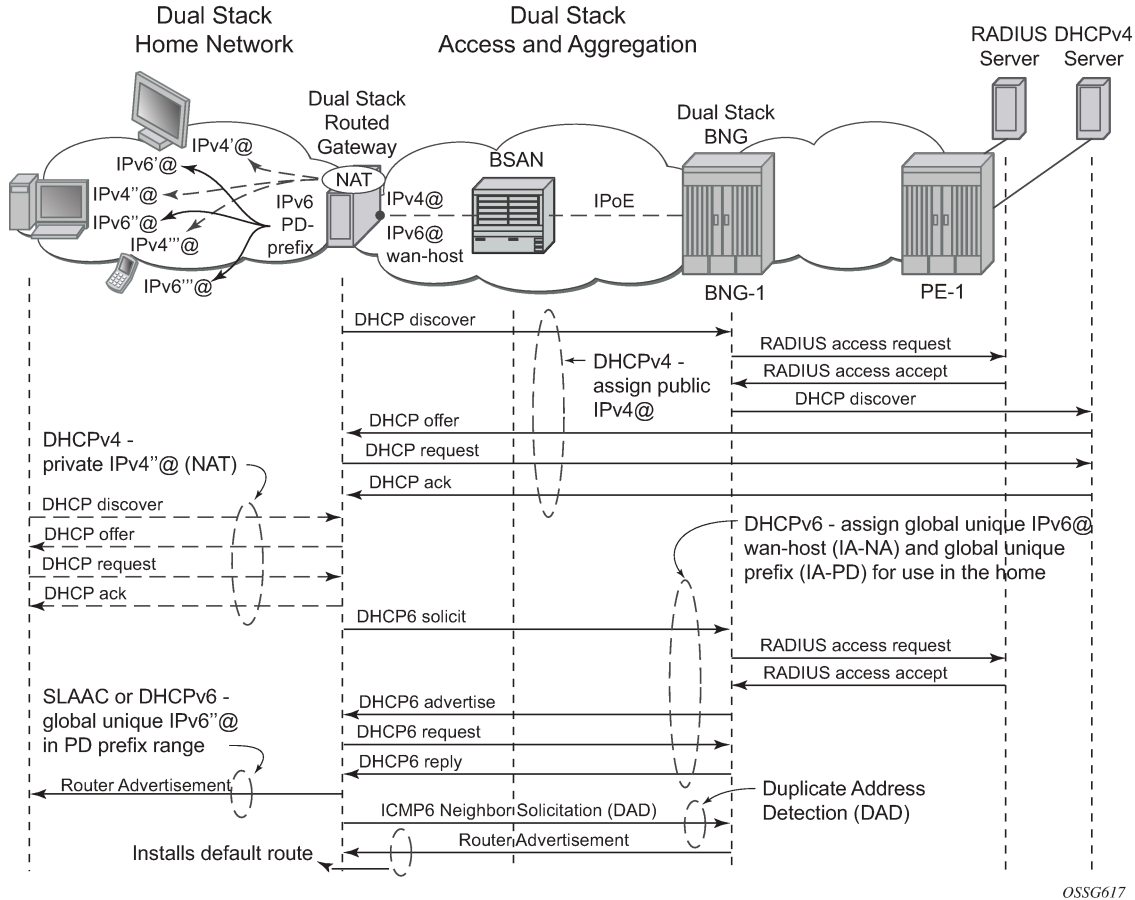


Dual Stack IPoE Routed Gateway

In services supporting dual stack IPoE Routed Gateways, the RG in the home network obtains an IPv4 address through the DHCPv4 protocol and an IPv6 Prefix Delegation (PD) prefix and/or wan-host IPv6 address through the DHCPv6 protocol. The Broadband Network Gateway (BNG) authenticates and authorizes both sessions independently.

In the home network, the dual stack RG performs Network Address Translation (NAT) for IPv4, using the assigned IPv4 address as outside address. A globally unique IPv6 prefix per subscriber is assigned and delegated by the BNG to the RG for use in the home network. The RG can use Stateless Address Auto Configuration (SLAAC) or DHCPv6 to allocate IPv6 addresses from this so called Prefix Delegation (PD) prefix to the devices in the home network. The wan-host IPv6 address is used by the RG on the WAN side (network facing). In case of an unnumbered RG, no wan-host address is obtained.

Figure 64: Dual Stack IPoE Routed Gateway Service



Recap of the DHCPv6 Protocol

The Dynamic Host Configuration Protocol for IPv6 (DHCPv6) is defined in RFC 3315, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*. The protocol enables DHCPv6 servers to pass configuration parameters such as IPv6 network addresses to IPv6 nodes.

DHCPv6 uses the Identity Association (IA) option to assign IPv6 addresses or prefixes. Two different IA types will be used in this section:

- Identity Association for Non-temporary Address (IA-NA) defined in RFC 3315. Used for wan-host IPv6 address assignment.

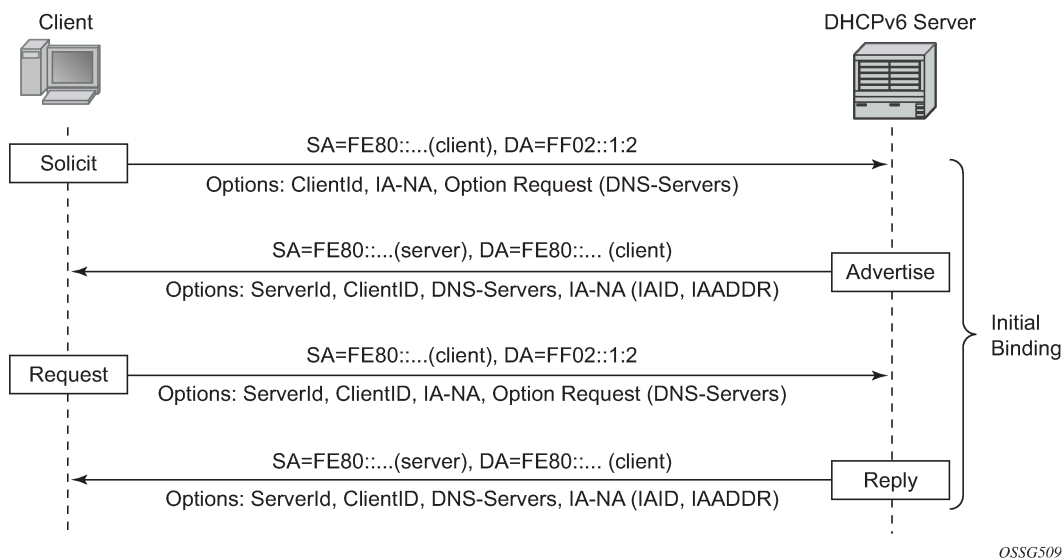
```
Option : IA_NA (3), Length : 40
  IAID : 1
  Time1: 1800 seconds
  Time2: 2880 seconds
Option : IAADDR (5), Length : 24
  Address : 2001:db8:b001:101::1
  Preferred Lifetime : 3600 seconds
  Valid Lifetime      : 86400 seconds
```

- Identity Association for Prefix Delegation (IA-PD), defined in RFC 3633. Used for prefix delegation assignment (for an explanation on prefix delegation, see [Prefix Delegation](#))

```
Option : IA_PD (25), Length : 41
  IAID : 1
  Time1: 1800 seconds
  Time2: 2880 seconds
Option : IAPREFIX (26), Length : 25
  Prefix : 2001:db8:a001:103::/56
  Preferred Lifetime : 3600 seconds
  Valid Lifetime    : 86400 seconds
```

The DHCPv6 lease process is outlined in [Figure 65: DHCPv6 Lease Process \(Part A\)](#) and [Figure 66: DHCPv6 Lease Process \(Part B\)](#).

Figure 65: DHCPv6 Lease Process (Part A)



A DHCPv6 client, sends a Solicit message to locate servers to the All DHCPv6 Relay Agents and Servers link-scoped multicast address (FF02::1:2), using its link-local address as source address. The DHCPv6 client includes in the Solicit message its ClientID, Identity Associations (IA) to request IPv6 address or prefix allocation and optionally an Option Request option.


Any on-link DHCPv6 server responds with a unicast Advertise message using the link local addresses. The server includes in the Advertise message the ClientID, its ServerID, IPv6 addresses and/or prefixes in Identity Associations (IA) and options containing the requested configuration parameters.

The DHCPv6 client selects an Advertise message and sends a Request message to the All DHCPv6 Relay Agents and Servers link-scoped multicast address. It includes its ClientID, the ServerID of the corresponding DHCPv6 server, Identity Associations (IA) to request IPv6 address or prefix allocation and optionally an Option Request option.

Upon receipt of a valid Request message, the DHCPv6 server with corresponding ServerID, sends a unicast Reply message using the link local addresses. The Reply contains the ClientID and ServerID, IPv6 addresses and/or prefixes in Identity Associations (IA) and options containing the requested configuration options.

The DHCPv6 client should perform Duplicate Address Detection (DAD) on the addresses in any IA it received in the REPLY before using that address for traffic.

Figure 66: DHCPv6 Lease Process (Part B)



Upon expiration of the renew timer T1 associated with the Identity Association option, the DHCPv6 client sends a Renew to the All DHCPv6 Relay Agents and Servers link-scoped multicast address to request an extension of the lifetime of an address. It includes its ClientID, the ServerID of the DHCPv6 server that originally provided the address, and Identity Associations (IA) containing the IPv6 address or prefix for which an extension of the lifetime is requested.

Upon expiration of the rebind timer T2 associated with the Identity Association option (no response received to the Renew), the DHCPv6 client sends a Rebind to the All DHCPv6 Relay Agents and Servers link-scoped multicast address to request an extension of the lifetime of an address. It includes its ClientID and Identity Associations (IA) containing the IPv6 address or prefix for which an extension of the lifetime is requested.

If a DHCPv6 client no longer uses one or more of the assigned addresses or prefixes, it sends a Release message to the server that assigned the address or prefix. The server acknowledges with a Reply message and includes a status code (for example, success).

If the DHCPv6 server sends a Server Unicast Option, then the DHCPv6 client should unicast the Request, Renew Release, and Decline messages to the server using the IPv6 address specified in the option. The 7750 SR DHCPv6 proxy server does not include the Server Unicast Option.


The DHCPv6 client should perform Duplicate Address Detection (DAD) on each of the addresses assigned through DHCPv6, before using that address for traffic. The DHCPv6 client uses Neighbor Solicitation for this purpose as described in RFC 4862, *IPv6 Stateless Address AutoConfiguration*.

Unlike DHCPv4, DHCPv6 does not provide a default route. In IPv6, default routes are learned via Router Advertisements (see [Enable Router Advertisements](#)).

Prefix Delegation

Prefix Delegation (PD) is a mechanism for automated delegation of IPv6 prefixes using DHCPv6. A delegating router delegates a long-lived IPv6 prefix to a requesting router. The delegating router does not require knowledge about the topology of the links in the network to which the prefixes will be assigned.

Figure 67: Prefix Delegation



OSSG511

In the context of ESM IPv6, the BNG is the delegating router (DHCPv6 server) and the Routed Gateway in the home is the requesting router (DHCPv6 client). The DHCPv6 option Identity Association for Prefix Delegation (IA-PD) (Figure 67: Prefix Delegation) is used to assign the IPv6 prefix.

Note that the mechanism through which a requesting router (routed gateway) assigns IPv6 addresses on its interfaces (home network) is arbitrary and can be based upon SLAAC (as shown in Figure 67: Prefix Delegation) or DHCPv6.

Configuration

ESMv6 for IPoE is applicable in a Routed CO environment. The two following scenarios show a minimal configuration to enable dual stack subscribers in a VPRN service context where the ESM IPv6 specific parts are highlighted. No subscriber QoS policies are defined because this is out of the scope for this chapter.

Scenario 1 - RADIUS

RADIUS is used for authentication and authorization (later referenced as RADIUS), and is configured as follows:

```
configure
router
  radius-server
    server "radius-172.16.1.2" address 172.16.1.2 secret vsecret1 create
    accept-coa
  exit
exit
exit
exit
configure
```

```
aaa
    radius-server-policy "rsp-1" create
        servers
            router "Base"
            source-address 192.0.2.1
            server 1 name "radius-172.16.1.2"
        exit
    exit
exit
configure
    subscriber-mgmt
        authentication-policy "auth-1" create
            description "RADIUS authentication policy"
            pppoe-access-method pap-chap
            radius-server-policy "rsp-1"
            password letmein
        exit
    exit
exit
```

The subscriber management profiles used in this chapter are defined as follows:

```
configure
    subscriber-mgmt
        sla-profile "sla-profile-1" create
        exit
        sub-profile "sub-profile-1" create
        exit
        sub-ident-policy "sub-ident-1" create
            sub-profile-map
                use-direct-map-as-default
            exit
            sla-profile-map
                use-direct-map-as-default
            exit
            strings-from-option 254
        exit
    exit
exit
```

Service VPRN-1 is defined as follows:

```
configure
    service
        vprn 1 customer 1 create
            dhcp
                local-dhcp-server "dhcp-s1" create
                    use-gi-address
                    pool "pool-1" create
                        subnet 10.1.0.0/16 create
                            options
                                subnet-mask 255.255.0.0
                                default-router 10.1.255.254
                            exit
                        address-range 10.1.0.1 10.1.0.255
                    exit
                exit
            no shutdown
        exit
    exit
    ---snip---
```

```
interface "system" create
  address 192.0.2.1/32
  local-dhcp-server "dhcp-s1"
  loopback
exit
subscriber-interface "sub-int-1" create
  address 10.1.255.254/16
  dhcp
    gi-address 10.1.255.254
  exit
  ipv6
    delegated-prefix-len 56
    subscriber-prefixes
      prefix 2001:db8:a001::/48 pd
      prefix 2001:db8:b001:100::/56 wan-host
    exit
  exit
  group-interface "grp-int-1" create
    description "radius authentication and authorization"
    ipv6
      router-advertisements
        managed-configuration
        no shutdown
      exit
      dhcp6
        proxy-server
        no shutdown
      exit
    exit
  exit
  dhcp
    proxy-server
    emulated-server 10.1.255.254
    no shutdown
  exit
  server 192.0.2.1
  trusted
  lease-populate 10
  no shutdown
  exit
  authentication-policy "auth-1"
  sap 1/1/1:1 create
    sub-sla-mgmt
    sub-ident-policy "sub-ident-1"
    multi-sub-sap 10
    no shutdown
  exit
  exit
  exit
  ---snip---
  exit
  service-name "dual-stack-service"
  no shutdown
  exit
exit
exit
exit
```

Scenario 2 - LUDB

The Local User Database used for authentication and authorization (later referenced as LUDB) is defined as follows:

```
configure
  subscriber-mgmt
    local-user-db ludb-1 create
      ipoe
        match-list mac
        host "host-3" create
          host-identification
            mac 00:0c:29:00:00:23
          exit
          address gi-address
          identification-strings 254 create
            subscriber-id "sub-3"
            sla-profile-string "sla-profile-1"
            sub-profile-string "sub-profile-1"
          exit
          options
            subnet-mask 255.255.0.0
            default-router 10.1.255.254
          exit
          ipv6-address 2001:db8:b001:103::3
          ipv6-delegated-prefix 2001:db8:a001:300::/56
          ipv6-delegated-prefix-len 56
          options6
            dns-server 2001:db8:dddd:1::1 2001:db8:dddd:2::1
          exit
          no shutdown
        exit
      ---snip---
    exit
  no shutdown
exit
exit
exit
```

Service VPRN-1 is extended as follows:

```
configure
  service
    vprn 1 customer 1 create
      ---snip---
    subscriber-interface "sub-int-1" create
      address 10.1.255.254/16
      dhcp
        gi-address 10.1.255.254
      exit
    ipv6
      delegated-prefix-len 56
      subscriber-prefixes
        prefix 2001:db8:a001::/48 pd
        prefix 2001:db8:b001:100::/56 wan-host
      exit
    exit
    group-interface "grp-int-2" create
      description "ludb authentication and authorization"
    ipv6
      router-advertisements
```



```
        prefix-options
        autonomous
        exit
        no shutdown
    exit
    dhcp6
        user-db "ludb-1"
        proxy-server
            client-applications ipoe
            no shutdown
        exit
    exit
exit
dhcp
    proxy-server
        emulated-server 10.1.255.254
        no shutdown
    exit
    server 192.0.2.1
    trusted
    lease-populate 10
    user-db "ludb-1"
    no shutdown
exit
sap 1/1/1:2 create
    sub-sla-mgmt
        def-sub-profile "sub-profile-1"
        def-sla-profile "sla-profile-1"
        sub-ident-policy "sub-ident-1"
        multi-sub-sap 10
        no shutdown
    exit
exit
exit
exit
service-name "dual-stack-service"
no shutdown
exit
exit
exit
```

Configuring IPv6 Subscriber Prefixes

Applies to both scenarios RADIUS and LUDB.

IPv6 subscriber prefixes must be defined at the **subscriber-interface>ipv6>subscriber-prefixes** context. Three types of prefixes can be configured:

- **wan-host** — Prefix from which the IPv6 addresses are assigned that are to be used on the Routed Gateway WAN interface (network facing).
- **pd** — Prefix from which the IPv6 Prefix Delegation prefixes are assigned that are to be used by the Routed Gateway for allocation in the home network (LAN interfaces).
- **pd wan-host** (both) — Prefix from which both IPv6 addresses (wan-host) and IPv6 Prefix Delegation prefixes (pd) can be assigned. This requires that the delegated prefix length is set to 64 bits.

A subscriber prefix length must be between /32 and /63.

Subscriber prefixes are subnetted in fixed length subnets that are assigned to subscriber hosts:

- /64 for **wan-host** subscriber prefixes

A /128 IPv6 address is assigned to the subscriber host. Broadband Forum standards require a /64 prefix per subscriber even when used for WAN interfaces and thus the full /64 subnet gets associated with the subscriber host [ref. WT-177 - IPv6 in the context of TR-101]. Two subscriber hosts cannot get an IPv6 address from the same /64 subnet.

- /delegated-prefix-len (/48..64) for **pd** subscriber prefixes

The delegated prefix length is configured in the **subscriber-interface>ipv6** context. The recommended value by Broadband Forum standards is /56 (default = /64) [ref. WT-177 - IPv6 in the context of TR-101]. The configured length applies to all **pd** subscriber prefixes on a subscriber-interface.

[Table 9: Applicable Subscriber-Prefix Parameters](#) provides an overview of the subscriber-prefix parameters that apply:

Table 9: Applicable Subscriber-Prefix Parameters

Subscriber prefix type	Subscriber prefix length	DHCPv6 option	Must be subnetted as
wan-host	/32..63	IA-NA	/64 (assigned as /128)
pd	/32..63 (*)	IA-PD	/delegated-prefix-len

(*) must be smaller than configured delegated prefix length

Enable DHCPv6 Proxy Server

Applies to RADIUS and LUDB scenarios.

An IPv6 IPoE subscriber host initiates a DHCPv6 session to request its configuration data (IPv6 addresses and/or IPv6 PD prefixes, DNS servers). Upon receipt of a DHCPv6 Solicit message, the BNG authenticates the IPv6 subscriber host and obtains its configuration information from a RADIUS server or local user database. A DHCPv6 proxy server in the BNG maintains the DHCPv6 session with the IPv6 IPoE subscriber host.

The DHCPv6 proxy server must be enabled in the **subscriber-interface>group-interface>ipv6>dhcp6>proxy-server** context. The default is **shutdown**.

```
configure
  service
    vprn 1 customer 1 create
      subscriber-interface "sub-int-1"
        group-interface "grp-int-1"
          ipv6
            dhcp6
              proxy-server
                server-id duid-ll
                renew-timer min 30
                rebind-timer min 48
                valid-lifetime days 1
                preferred-lifetime hrs 1
                client-applications dhcp
                no shutdown
              exit
            exit
          exit
        exit
      exit
    exit
```

```
exit
    exit
    exit
exit
```

When enabled, the DHCPv6 proxy server by default allows IPv6 IPoE hosts to authenticate (configured with client-applications dhcp). Additionally, you can enable support for IPv6 PPPoE hosts. See [ESMv6: PPPoE Dual Stack Hosts](#).

A number of timers associated with IPv6 addresses and IPv6 prefixes within DHCPv6 Identity Associations can be configured in the DHCPv6 proxy server.

RFC 4862 defines two timers associated with graceful degradation of address bindings:

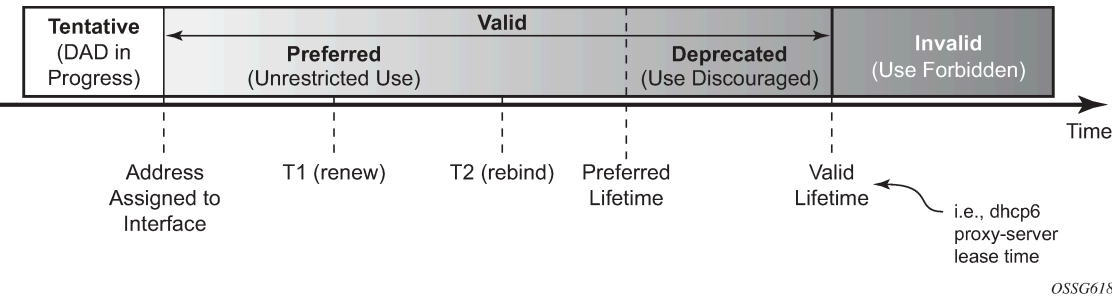
- Preferred lifetime — The length of time that a valid address is preferred (the time until deprecation). When the preferred lifetime expires, the address becomes deprecated and its use should be discouraged for new sessions.
- Valid lifetime — The length of time an address remains in the valid state (the time until invalidation). The valid lifetime must be greater than or equal to the preferred lifetime. When the valid lifetime expires, the address becomes invalid.

RFC 3315, *DHCPv6*, defines two timers associated with an Identity Association (IA) option that give the servers explicit control over when a client recontacts the server about a specific IA:

- T1 (renew) — The time at which the client contacts the server from which the addresses/prefix in the IA were obtained to extend the lifetimes of the addresses/prefix assigned to the IA
- T2 (rebind) — The time at which the client contacts any available server to extend the lifetimes of the addresses/prefixes assigned to the IA;

These timers are common for all DHCPv6 sessions in a group-interface and cannot be configured from RADIUS or local user database.

Figure 68: IPv6 Address/Prefix Timers



When violating the following rule, the default timers will be used:

Table 10: Timer Parameters

Timer	Use	Default	Range
T1	Renew timer	1800s (30 min)	0..604800s (7 days)
T2	Rebind timer	2880s (48 min)	0..1209600s (14 days)
preferred-lifetime		3600s (1hr)	300..4294967295s

Timer	Use	Default	Range
valid-lifetime	DHCPv6 lease time	86400s (24 hrs)	300..4294967295s

If the DHCPv6 lease is not renewed by the client before the DHCPv6 lease timer expires, then the subscriber host is deleted from the system. In other words, beyond the valid lifetime, subscriber traffic from/to the associated IPv6 addresses is dropped.

Enable Router Advertisements

Applies to both scenarios RADIUS and LUDB.

In IPv6, default routes are automatically installed via the router discovery mechanism. Unsolicited Router Advertisements (RA) must explicitly be enabled on a group interface. The default is **shutdown**.

```
configure
service
  vprn 1 customer 1 create
    subscriber-interface "sub-int-1"
    group-interface "grp-int-1"
      ipv6
        router-advertisements
          managed-configuration
          no shutdown
        exit
      exit
    exit
  exit
exit
exit
exit
exit
exit
```

The managed-configuration flag is set for consistency only. It tells the hosts that addresses can be requested using DHCPv6. However, as described in the Security section later (see [Security](#)), the host cannot rely on this flag because DHCPv6 must be initiated by the host before the BNG sends RAs.

Additional parameters that can be configured with respect to the router advertisements (defaults are shown):

```
configure
service
  vprn 1 customer 1 create
    subscriber-interface "sub-int-1"
    group-interface "grp-int-1"
      ipv6
        router-advertisements
          shutdown
          current-hop-limit 64
          dns-options
            no include-dns
            rdns-lifetime 3600
          exit
          no force-mcast
          no managed-configuration
          max-advertisement 1800
          min-advertisement 900
          no mtu
          no other-stateful-configuration
```

```

prefix-options
  no autonomous
  on-link
  preferred-lifetime 3600
  valid-lifetime 86400
exit
reachable-time 0
retransmit-time 0
router-lifetime 4500
exit
exit
exit
exit
exit
exit
exit

```

Table 11: Router Advertisements Parameters

Parameter	Description (RFC 4861, <i>Neighbor Discovery for IP version 6 (IPv6)</i>)	Value Range (default)
current-hop-limit	The default value that should be placed in the Hop Count field of the IP header for outgoing IP packets. A value of zero means unspecified (by this router); the RG picks its own value.	0..255 (64)
dns-options: include-dns	Indication to include the Recursive DNS Server (RDNSS) option as defined in RFC 6106 in IPv6 RAs for DNS name resolution of IPv6 SLAAC hosts	(no)
dns-options: rdNSS-lifetime	Indicates the maximum time that the RDNSS address may be used for name resolution	3600 (s)
force-mcast	Configures multicast router advertisements on this interface, either IP or MAC	(no)
managed-configuration	Managed address configuration flag. When set, it indicates that addresses are available through DHCPv6	(no)
max-advertisement	Unsolicited Router Advertisements are not strictly periodic: the interval between subsequent transmissions is randomized to reduce the probability of synchronization with the advertisements from other routers on the same link. Whenever a multicast advertisement is sent from an interface, the timer is reset to a uniformly distributed random value between the interface's configured MinRtrAdvInterval and MaxRtrAdv Interval.	900..1800 s (1800)
min-advertisement		900..1350 s (900)
mtu	Routers can advertise an MTU for hosts to use on the link.	1280..9212 bytes (no)
other-stateful-configuration (not applicable for IPoE)	Other configuration flag. When set, it indicates that other configuration information is available through DHCPv6. (DNS). Can be ignored if managed address configuration flag is enabled	(no)

Parameter	Description (RFC 4861, <i>Neighbor Discovery for IP version 6 (IPv6)</i>)	Value Range (default)
prefix-options: autonomous (not applicable for IPoE)	Autonomous address-configuration flag. When set indicates that this prefix can be used for stateless address autoconfiguration (SLAAC)	(no)
prefix-options: on-link	Indicates whether the prefix will be assigned to an interface on the specified link	(no)
prefix-options: preferred-lifetime (not applicable for IPoE)	The length of time in seconds that addresses generated from the prefix via stateless address autoconfiguration (SLAAC) remain preferred	0..4294967295 (3600)
prefix-options: valid-lifetime (not applicable for IPoE)	The length of time in seconds that the prefix is valid for the purpose of on-link determination. (also used by SLAAC)	0..4294967295 (86400)
reachable-time	The time that a node assumes a neighbor is reachable after having received a reachability confirmation. Used by the Neighbor Unreachability Detection algorithm. A value of zero means unspecified (by this router); the RG picks its own value.	0..3600000 ms (0)
retransmit-time	The time between retransmitted Neighbor Solicitation messages. Used by address resolution and the Neighbor Unreachability Detection algorithm. A value of zero means unspecified (by this router); the RG picks its own value.	0..1800000 ms (0)
router-lifetime	The lifetime associated with the default router in units of seconds.	2700..9000 s (4500)

RADIUS Authentication and Authorization

Applies to the RADIUS scenario only.

The RADIUS authentication and authorization configuration for IPoE IPv6 subscriber host is no different from an IPv4 subscriber host:

```

configure
  router
    radius-server
      server "radius-172.16.1.2" address 172.16.1.2 secret vsecret1 create
      accept-coa
    exit
  exit
exit

configure
  aaa
    radius-server-policy "rsp-1" create
    servers
      router "Base"
      source-address 192.0.2.1

```

```

        server 1 name "radius-172.16.1.2"
        exit
    exit
exit
configure
  subscriber-mgmt
    authentication-policy "auth-1" create
    description "RADIUS authentication policy"
    pppoe-access-method pap-chap
    radius-server-policy "rsp-1"
    password letmein
  exit
exit
exit

```

Additional RADIUS AVPs that are applicable to IPoE IPv6 subscriber hosts are listed in [Table 12: RADIUS AVPs](#).

Table 12: RADIUS AVPs

RADIUS AVP	Type	Purpose
Alc-IPv6-Address [26-6527-99]	ipv6addr	maps to IA_NA of DHCPv6 (RG WAN interface address)
Alc-Ipv6-Primary-Dns [26-6527-105]	ipv6addr	maps to DNS Recursive Name Server option (RFC 3646, <i>DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)</i>) in DHCPv6
Alc-Ipv6-Secondary-Dns [26-6527-106]	ipv6addr	maps to DNS Recursive Name Server option (RFC 3646) in DHCPv6
Delegated-IPv6-Prefix [123]	ipv6prefix	maps to IA_PD for prefix delegation (RFC 3633, <i>IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6</i>) in DHCPv6

A sample FreeRADIUS users record to authenticate a dual stack IPoE subscriber:

```

00:0c:20:00:00:21 Cleartext-Password := "letmein"
  Alc-Subsc-ID-Str = "sub-1",
  Alc-Subsc-Prof-Str = "sub-profile-1",
  Alc-SLA-Prof-Str = "sla-profile-1",
  Framed-IP-Address = 10.1.0.1,
  Framed-IP-Netmask = 255.255.0.0,
  Framed-Route = "172.16.11.0/24 0.0.0.0",
  Alc-IPv6-Address = 2001:db8:b001:101::1,
  Delegated-IPv6-Prefix = 2001:db8:a001:100::/56,
  Alc-IPv6-Primary-Dns = 2001:db8:dddd:1::1,
  Alc-IPv6-Secondary-Dns = 2001:db8:dddd:2::1,

```

The FreeRADIUS Server 2.0.0 and greater has full support for both IPv6 attributes and IPv6 network packets.

The IPv6 address/prefix related timers can be configured in the **dhcp6>proxy-server** context (see [Enable DHCPv6 Proxy Server](#)).

Local User Database Authentication and Authorization

Applies to the LUDB scenario only.

The configuration example below focuses on the IPv6 host configuration. The details for local user database host matching and IPv4 host specific parameters are out of scope for this section.

```
configure
  subscriber-mgmt
    local-user-db "ludb-1" create
      dhcp
        match-list mac
        host "host-1" create
          host-identification
            mac 00:0c:29:00:00:23
          exit
          address gi-address # IPv4 host
          identification-strings 254 create
            subscriber-id "sub-3"
            sla-profile-string "sla-profile-1"
            sub-profile-string "sub-profile-1"
          exit
          options
            subnet-mask 255.255.0.0 # IPv4 host
            default-router 10.1.255.254 # IPv4 host
          exit
          ipv6-address 2001:db8:b001:103::3 # IPv6 host
          ipv6-delegated-prefix 2001:db8:a001:300::/56 # IPv6 host
          options6
            dns-server 2001:db8:dddd:1::1 2001:db8:dddd:2::1
          exit
          no shutdown
        exit
      exit
    exit
  no shutdown
exit
exit
exit
exit
```

```
configure
  service
    vprn 1 customer 1 create
      subscriber-interface "sub-int-1" create
        group-interface "grp-int-2" create
          description "ludb authentication and authorization"
          ipv6
            ---snip---
            dhcp6
              user-db "ludb-1"
              proxy-server
                client-applications dhcp
                no shutdown
              exit
            exit
          exit
        exit
      dhcp
        ---snip---
```



```

server 192.0.2.1
trusted
lease-populate 10
user-db "ludb-1"
no shutdown
exit
exit
exit
exit
exit
exit

```

Besides the identification strings that are common to the IPv4 and IPv6 hosts, specific IPv6 host related parameters can be configured:

Table 13: Local User Database Parameters

local-user-db CLI parameter	Purpose
ipv6-address	Maps to IA_NA of DHCPv6 (RG WAN interface address)
ipv6-prefix	Maps to IA_PD for prefix delegation (RFC 3633) in DHCPv6
options6: dns-server	Defines the IPv6 DNS server address to be used for name resolution

The IPv6 address/prefix related timers can be configured in the **dhcp6>proxy-server** context (see [Enable DHCPv6 Proxy Server](#)).

DHCP and DHCP6 Lease State

Applies to both scenarios RADIUS and LUDB.

The DHCP lease state is an internal database structure that keeps track of the DHCP host states. The DHCP lease state enables subscriber management functions (for example, per subscriber QoS and accounting) and security functions (for example, dynamic anti-spoof filtering) on the DHCP host.

The DHCP lease information for a specific host is extracted from the DHCPv4 ack message in case of DHCPv4 and from the DHCPv6 reply message in case of DHCPv6

Typical information stored in the DHCP lease state includes (partial table; additional data can be stored for managed SAPs, wholesale-retail).

Table 14: DHCP Lease State Information

Parameter	Comment
Service ID	Service where the DHCP host is connected.
IP Address	IPv4 or IPv6 address of the DHCP host.
Client HW Address	Ethernet MAC address of the DHCP host.

Parameter	Comment
Subscriber-interface (Routed CO only)	Subscriber interface name where the DHCP host is instantiated.
Group-interface (Routed CO only)	Group interface name where the DHCP host is instantiated.
SAP	SAP where the DHCP hosts is connected.
Remaining Lifetime	The remaining time before the DHCP host is deleted from the system (updated each time a DHCP renew/rebind occurs).
Persistence Key	Lookup key for this host in the persistency file.
Sub-Ident	ESM: Subscriber ID of the DHCP host.
Sub-Profile-String	ESM: Subscriber profile string of the DHCP host.
SLA-Profile-String	ESM: SLA profile string of the DHCP host.
App-Profile-String	ESM: Application profile string of the DHCP host.
Lease ANCP-String	ESM: ANCP string for this DHCP host.
Lease Int Dest Id	ESM: Internal destination ID for this DHCP host.
Category-Map-Name	ESM: Volume and Time based accounting.
Dhcp6 ClientId (DUID)	DHCPv6 client unique identifier.
Dhcp6 IAID	Identity Association ID chosen by the client.
Dhcp6 IAID Type	Identity Association type: prefix (PD) or non-temporary (wan-host).
Dhcp6 Client Ip	Link local IPv6 address of the host.
Sub-Ident origin	ESM: Origin for the Subscriber ID for this host (None, DHCP, RADIUS).
Strings origin	ESM: Origin for the ESM strings for this host (None, DHCP, RADIUS).
Lease Info origin	ESM: Origin for the IP configuration for this host (None, DHCP, RADIUS).
Ip-Netmask	The IP netmask for this DHCP host.
Broadcast-Ip-Addr	The broadcast IP address for this host.
Default-Router	The default gateway for this host.
Primary-Dns	The primary DNS server for this host.
Secondary-Dns	The secondary DNS server for this host.

Parameter	Comment
Primary-Nbns	The primary NetBIOS name server for this host.
Secondary-Nbns	The secondary NetBIOS name server for this host.
ServerLeaseStart	Time and date that the lease for this host started (first DHCP ack received).
ServerLastRenew	Time and date that the lease for this host was last renewed.
ServerLeaseEnd	Time and date that the lease for this host will expire.
Session-Timeout	Lease time specified by the DHCP server.
DHCP Server Addr	IP address of the DHCP server that allocated the lease for this host.
Circuit Id	DHCP Relay Agent information option 82 Circuit ID content.
Remote Id	DHCP Relay Agent information option 82 Remote ID content.
RADIUS User-Name	ESM: Username used in the RADIUS authentication access request.

DHCPv4 lease state population is enabled by default on a group-interface with DHCP configured as **no shutdown**. The number of DHCPv4 leases allowed on each SAP of the group-interface must be configured with the **lease-populate** option (by default a single DHCPv4 host is allowed on each SAP of the group-interface).

DCHPv6 lease state population is enabled by default on a group-interface with DHCP6 proxy-server configured as **no shutdown**. The number of DHCPv6 leases (hosts) cannot be limited per group-interface.

```
configure
service
  vprn 1 customer 1 create
    subscriber-interface "sub-int-1" create
      group-interface "grp-int-1" create
        description "radius authentication and authorization"
        ipv6
          dhcp6
            proxy-server
              no shutdown
            exit
          exit
        exit
      dhcp
        proxy-server
          emulated-server 10.1.255.254
          no shutdown
        exit
        server 192.0.2.1
        trusted
        lease-populate 10
        no shutdown
      exit
    exit
  exit
exit
```

```
exit
exit
```

To check the DHCPv4 or DHCPv6 lease state for a particular service, use the following commands (detailed output as well as additional output filtering is available):

```
*A:BNG# show service id 1 dhcp | dhcp6 lease-state ?
```

```
- lease-state [wholesaler <service-id>] [sap <sap-id>|sdp <sdp-id:vc-id>|
interface <interface-name>|ip-address <ip-address[/mask]>|chaddr
<ieee-address>|mac <ieee-address>|{[port <port-id>][no-inter-dest-id |
inter-dest-id <inter-dest-id>]]} [session {none|ipoe}] [detail]
```

```
*A:BNG# show service id 1 dhcp lease-state detail
```

```
=====
DHCP lease states for service 1
=====
```

```
Service ID      : 1
IP Address      : 10.1.0.1
Client HW Address : 00:0c:29:00:00:21
Subscriber-interface : sub-int-1
Group-interface  : grp-int-1
SAP              : 1/1/1:1
Termination Type : local
Up Time          : 0d 00:58:34
Remaining Lease Time : 6d 23:01:26
Remaining SessionTime: N/A
Persistence Key   : 0x00000000

Sub-Ident       : "sub-1"
Sub-Profile-String : "sub-profile-1"
SLA-Profile-String : "sla-profile-1"
App-Profile-String : ""
Lease ANCP-String : ""
Lease Int Dest Id : ""
Category-Map-Name : ""

Lease Info origin : Radius

Ip-Netmask       : 255.255.0.0
Broadcast-Ip-Addr : 10.1.255.255
Default-Router    : N/A
Primary-Dns       : N/A
Secondary-Dns     : N/A
Primary-Nbns      : N/A
Secondary-Nbns    : N/A

ServerLeaseStart  : 04/20/2017 13:01:09
ServerLastRenew   : 04/20/2017 13:01:09
ServerLeaseEnd    : 04/27/2017 13:01:09
Session-Timeout   : N/A
IPoE|PPP session  : No
Lease-Time        : 7d 00:00:00
DHCP Server Addr  : N/A
Radius User-Name   : "00:0c:29:00:00:21"
```

```
-----
Number of lease states : 1
=====
```

*A:BNG#

*A:BNG# show service id 1 dhcp6 lease-state detail

=====

DHCP lease states for service 1

=====

Service ID : 1
IP Address : 2001:db8:a001:100::/56
Client HW Address : 00:0c:29:00:00:21
Subscriber-interface : sub-int-1
Group-interface : grp-int-1
SAP : 1/1/1:1
Termination Type : local
Up Time : 0d 00:55:11
Remaining Lease Time : 0d 23:34:49
Remaining SessionTime: N/A
Persistence Key : 0x0000000b

Sub-Ident : "sub-1"
Sub-Profile-String : "sub-profile-1"
SLA-Profile-String : "sla-profile-1"
App-Profile-String : ""
Lease ANCP-String : ""
Lease Int Dest Id : ""
Category-Map-Name : ""
Dhcp6 ClientId (DUID): 00010001208a25ac000c29000021
Dhcp6 IAID : 1
Dhcp6 IAID Type : prefix
Dhcp6 Client Ip : fe80::20c:29ff:fe00:21
Primary-Dns : 2001:db8:dddd:1::1
Secondary-Dns : 2001:db8:dddd:2::1
Pool Name : ""
Dhcp6 Server Addr : N/A
Dhcp6 ServerId (DUID): N/A
Dhcp6 InterfaceId : N/A
Dhcp6 RemoteId : N/A
Radius sub-if prefix : N/A
Router adv. policy : N/A

Lease Info origin : Radius

ServerLeaseStart : 04/20/2017 13:06:36
ServerLastRenew : 04/20/2017 13:36:36
ServerLeaseEnd : 04/21/2017 13:36:36
Session-Timeout : N/A
IPoE|PPP session : No
Radius User-Name : "00:0c:29:00:00:21"

Service ID : 1
IP Address : 2001:db8:b001:101::1/128
Client HW Address : 00:0c:29:00:00:21
Subscriber-interface : sub-int-1
Group-interface : grp-int-1
SAP : 1/1/1:1
Termination Type : local
Up Time : 0d 00:55:11
Remaining Lease Time : 0d 23:34:49
Remaining SessionTime: N/A
Persistence Key : 0x0000000a

Sub-Ident : "sub-1"
Sub-Profile-String : "sub-profile-1"

```
SLA-Profile-String      : "sla-profile-1"
App-Profile-String      : ""
Lease ANCP-String       : ""
Lease Int Dest Id       : ""
Category-Map-Name       : ""
Dhcp6 ClientId (DUID)   : 00010001208a25ac000c29000021
Dhcp6 IAID              : 2
Dhcp6 IAID Type         : non-temporary
Dhcp6 Client Ip         : fe80::20c:29ff:fe00:21
Primary-Dns             : 2001:db8:dddd:1::1
Secondary-Dns           : 2001:db8:dddd:2::1
Pool Name               : ""
Dhcp6 Server Addr       : N/A
Dhcp6 ServerId (DUID)   : N/A
Dhcp6 InterfaceId       : N/A
Dhcp6 RemoteId          : N/A
Radius sub-if prefix    : N/A
Router adv. policy      : N/A

Lease Info origin       : Radius

ServerLeaseStart        : 04/20/2017 13:06:36
ServerLastRenew         : 04/20/2017 13:36:36
ServerLeaseEnd          : 04/21/2017 13:36:36
Session-Timeout         : N/A
IPoE|PPP session        : No
Radius User-Name        : "00:0c:29:00:00:21"
-----
Number of lease states : 2
=====
*A:BNG#
```

Operation

An IPoE dual stack subscriber in a numbered Routed Gateway scenario consumes three subscriber host entries:

- IPv4 host — DHCPv4 session based
- IPv6 wan-host — DHCPv6 session based
- IPv6 Prefix Delegation host — DHCPv6 session based

```
*A:BNG# show service active-subscribers

=====
Active Subscribers
=====
-----
Subscriber sub-1 (sub-profile-1)
-----
(1) SLA Profile Instance sap:1/1/1:1 - sla:sla-profile-1
-----
IP Address      MAC Address      Session      Origin      Svc      Fwd
-----
10.1.0.1        00:0c:29:00:00:21  N/A         DHCP        1        Y
2001:db8:a001:100::/56
00:0c:29:00:00:21  N/A         DHCP6       1        Y
```

```

2001:db8:b001:101::1/128
      00:0c:29:00:00:21    N/A          DHCP6          1          Y
-----
Subscriber sub-3 (sub-profile-1)
-----
(1) SLA Profile Instance sap:1/1/1:2 - sla:sla-profile-1
-----
IP Address          MAC Address          Session          Origin          Svc          Fwd
-----
10.1.0.8            00:0c:29:00:00:23    N/A             DHCP            1            Y
2001:db8:a001:300::/56
      00:0c:29:00:00:23    N/A             DHCP6           1            Y
2001:db8:b001:103::3/128
      00:0c:29:00:00:23    N/A             DHCP6           1            Y
-----
Number of active subscribers : 2
=====
*A:BN#

```

The optional **hierarchy** parameter for the active-subscribers display provides a top-down level overview for this subscriber:

```

*A:BN# show service active-subscribers hierarchy

=====
Active Subscribers Hierarchy
=====
-- sub-1 (sub-profile-1)
  |
  +-- sap:1/1/1:1 - sla:sla-profile-1
    |
    |-- 10.1.0.1 - mac:00:0c:29:00:00:21 - DHCP - svc:1
    |-- 2001:db8:a001:100::/56 - mac:00:0c:29:00:00:21 - DHCP6 - svc:1
    +-- 2001:db8:b001:101::1/128 - mac:00:0c:29:00:00:21 - DHCP6 - svc:1
-- sub-3 (sub-profile-1)
  |
  +-- sap:1/1/1:2 - sla:sla-profile-1
    |
    |-- 10.1.0.8 - mac:00:0c:29:00:00:23 - DHCP - svc:1
    |-- 2001:db8:a001:300::/56 - mac:00:0c:29:00:00:23 - DHCP6 - svc:1
    +-- 2001:db8:b001:103::3/128 - mac:00:0c:29:00:00:23 - DHCP6 - svc:1

-----
Number of active subscribers : 2
Flags: (N) = the host or the managed route is in non-forwarding state
=====
*A:BN#

```

The total number (sum) of IPv4 and IPv6 hosts per subscriber can be limited in the corresponding sla-profile with the **host-limits** parameter:

```
configure
  subscr-mgmt
    sla-profile "sla-profile-1" create
    host-limits
      overall 3
    exit
  exit
exit
exit
```

To display the IPv4/IPv6 routing table for dual stack hosts:

```
*A:BNG# show router 1 route-table ipv4 protocol sub-mgmt

=====
Route Table (Service: 1)
=====
Dest Prefix[Flags]                Type   Proto   Age           Pref
  Next Hop[Interface Name]          Metric
-----
10.1.0.1/32                        Remote Sub Mgmt 01h05m03s  0
      [grp-int-1]                     0
10.1.0.8/32                        Remote Sub Mgmt 00h00m49s  0
      [grp-int-2]                     0
-----
No. of Routes: 2
Flags: n = Number of times nexthop is repeated
      B = BGP backup route available
      L = LFA nexthop available
      S = Sticky ECMP requested
=====
*A:BNG#
```

```
*A:BNG# show router 1 route-table ipv6 protocol sub-mgmt

=====
IPv6 Route Table (Service: 1)
=====
Dest Prefix[Flags]                Type   Proto   Age           Pref
  Next Hop[Interface Name]          Metric
-----
2001:db8:a001:100::/56            Remote Sub Mgmt 01h00m03s  0
      [grp-int-1]                     0
2001:db8:a001:300::/56            Remote Sub Mgmt 00h03m00s  0
      [grp-int-2]                     0
2001:db8:b001:101::/128           Remote Sub Mgmt 01h00m03s  0
      [grp-int-1]                     0
2001:db8:b001:103::/128           Remote Sub Mgmt 00h03m00s  0
      [grp-int-2]                     0
-----
No. of Routes: 4
Flags: n = Number of times nexthop is repeated
      B = BGP backup route available
      L = LFA nexthop available
      S = Sticky ECMP requested
=====
*A:BNG#
```


Troubleshooting

Apart from the show commands in this chapter, the following additional commands can be used for troubleshooting:

- Default system log
- Debug
- Statistics

The default system log can be shown as follows:

```
A:BNG-1# show log log-id 99
```

Use appropriate filtering to reduce the output if needed.

Debugging can be done with the definitions as follows:

```
debug
  router "Base"
    radius
      packet-type authentication accounting coa
      detail-level high
    exit
  exit
  router "1"
    ip
      dhcp
        detail-level high
        mode egr-ingr-and-dropped
      exit
      dhcp6
        mode egr-ingr-and-dropped
        detail-level high
      exit
      icmp6
    exit
    local-dhcp-server "dhcp-s1"
      detail-level medium
      mode egr-ingr-and-dropped
    exit
  exit
  subscriber-mgmt
    local-user-db "ludb-1"
      detail all
    exit
  exit
exit
```

Additional filtering (such as only DHCPv6 debug for a particular interface) may be needed to prevent a flood of debug messages.

DHCPv4 statistics can be shown as follows:

```
*A:BNG# show router 1 dhcp statistics
```

```
=====
DHCP Global Statistics (Service: 1)
=====
```

```
Rx Packets : 86
```

```

Tx Packets : 36
Rx Malformed Packets : 0
Rx Untrusted Packets : 0
Client Packets Discarded : 12
Client Packets Relayed : 46
Client Packets Snooped : 4
Client Packets Proxied (RADIUS) : 24
Client Packets Proxied (Diameter) : 0
Client Packets Proxied (User-Db) : 0
Client Packets Proxied (Lease-Split) : 0
Server Packets Discarded : 0
Server Packets Relayed : 25
Server Packets Snooped : 0
DHCP RELEASEs Spoofed : 0
DHCP FORCERENEWs Spoofed : 0
Client packets streamed : 0
=====
*A:BN#

```

DHCPv6 statistics can be shown as follows:

```

*A:BN# show router 1 dhcp6 statistics
=====
DHCP6 statistics (Router: 1)
=====

```

Msg-type	Rx	Tx	Dropped
1 SOLICIT	6	0	0
2 ADVERTISE	0	6	0
3 REQUEST	6	0	0
4 CONFIRM	0	0	0
5 RENEW	1	0	0
6 REBIND	0	0	0
7 REPLY	0	11	0
8 RELEASE	4	0	0
9 DECLINE	0	0	0
10 RECONFIGURE	0	0	0
11 INFO_REQUEST	0	0	0
12 RELAY_FORW	0	0	0
13 RELAY_REPLY	0	0	0
14 LEASEQUERY	0	0	0
15 LEASEQUERY_REPLY	0	0	0

```

-----
Dhcp6 Drop Reason Counters :
-----

```

1 Dhcp6 oper state is not Up on src itf	0
2 Dhcp6 oper state is not Up on dst itf	0
3 Relay Reply Msg on Client Itf	0
4 Hop Count Limit reached	0
5 Missing Relay Msg option, or illegal msg type	0
6 Unable to determine destination client Itf	0
7 Out of Memory	0
8 No global Pfx on Client Itf	0
9 Unable to determine src Ip Addr	0
10 No route to server	0
11 Subscr. Mgmt. Update failed	0
12 Received Relay Forw Message	0
13 Packet too small to contain valid dhcp6 msg	0
14 Server cannot respond to this message	0
15 No Server Id option in msg from server	0
16 Missing or illegal Client Id option in client msg	0

```

17 Server Id option in client msg                                0
18 Server DUID in client msg does not match our own            0
19 Client sent message to unicast while not allowed            0
20 Client sent message with illegal src Ip address              0
21 Client message type not supported in pfx delegation          0
22 Nbr of addrs or pfxs exceeds allowed max (128) in msg        0
23 Unable to resolve client's mac address                       0
24 The Client was assigned an illegal address                   0
25 Illegal msg encoding                                          0
26 Client message not supported                                 0
27 IA options in info request                                    0
28 No IA option in client msg                                    0
29 No addresses in confirm msg                                    0
30 No relay servers configured                                   0
31 Blocked by host lockout                                       0
32 No link address available                                     0
33 Dropped by Python                                             0
34 Invalid server                                                0
35 Packet dropped on SRRP backup interface                      0
36 DHCP transaction not found                                    0
37 Could not determine retail interface                         0
38 Packet dropped by DHCP filter                                 0
39 Packet dropped because authentication failed                  0
=====
*A:BNG#

```

RADIUS statistics can be shown as follow:

```

*A:BNG# show subscriber-mgmt authentication "auth-1" statistics

=====
Authentication Policy Statistics
=====
-----
Policy name                : auth-1
subscriber packets authenticated : 0
subscriber packets rejected   : 0
subscriber packets rejected send failed : 0
-----
=====
*A:BNG#

```

```

*A:BNG# show aaa radius-server-policy "rsp-1" statistics

=====
RADIUS server policy "rsp-1" statistics
=====
Tx transaction requests      : 24
Rx transaction responses     : 24
Transaction requests timed out : 0
Transaction requests send failed : 0
Packet retries               : 0
Transaction requests send rejected : 0
Authentication requests failed : 4
Accounting requests failed   : 0
Ratio of access-reject over auth responses : 16%
Transaction success ratio    : 100%
Transaction failure ratio    : 0%
Statistics last reset at     : n/a

Server 1 "radius-172.16.1.2" address 172.16.1.2 auth-port 1812 acct-port 1813
-----
Tx request packets          : 24

```

```
Rx response packets           : 24
Request packets timed out     : 0
Request packets send failed   : 0
Request packets send failed (overload) : 0
Request packets waiting for reply : 0
Response packets with invalid authenticator : 0
Response packets with invalid msg authenticator : 0
Authentication packets failed : 4
Accounting packets failed     : 0
Avg auth response delay (10 100 1K 10K) in ms : 1.25    168    168    168
Avg acct response delay (10 100 1K 10K) in ms : n/a
Statistics last reset at      : n/a

=====
*A: BNG#
```

Advanced Topics

Security

Downstream Router Advertisements

When a SAP is bound to a subscriber/group-interface which has IPv6 enabled, there will be no initial downstream Router Advertisement (RA) message sent. If a SAP is shared by multiple subscribers, it would be possible for an unauthenticated host to receive the RA.

Instead the RAs are sent in unicast to allow per-host IPv6 link configuration. This requires the host information (MAC address and link-local IPv6 address) to be known. Therefore, for IPoE, until a DHCPv6 session is bound, no unsolicited or solicited RAs are sent.

Processing Neighbor Discovery Messages

Processing Neighbor Discovery messages: Neighbor Advertisements (NA), Neighbor Solicitations (NS) and Router Solicitations (RS).

Neighbor discovery messages are not processed prior to IPoE IPv6 host authentication to avoid DoS attacks consuming CPU resources. This implies that an IPoE host should initiate the DHCPv6 session without link information and knowledge of routers on the link as required by the Broadband Forum standards (ref. TR-124 issue 2 — Functional Requirements for Broadband Residential Gateway Devices). This is not a problem as the DHCPv6 solicit/request messages are sent to a well-known multicast address with direct link-layer mapping.

After DHCP host authentication, Neighbor Discovery messages will not result in a neighbor cache entry. Instead a managed neighbor cache entry is created based on the DHCPv6 lease state. This managed neighbor cache entry cannot be displayed. The above mechanism prevents DoS attacks from poisoning the neighbor cache with bogus entries.

Router advertisements in response to a router solicitation are internally throttled so that they are not sent more often than once every three seconds.

Anti-spoof Filters

For each authenticated IPoE IPv6 host, an anti-spoof filter entry is created that allows upstream traffic with exact match on the tuple {masked source IP, source MAC} to pass. Traffic from unauthenticated hosts is silently dropped.

Managed SAPs

To allow the creation of managed SAPs in a dual stack environment, both DHCPv4 discover and DHCPv6 solicit messages received on a capture SAP should trigger RADIUS authentication:

```
configure
  service
    vpls 2 customer 1 create
      sap 1/1/2:* capture-sap create
        trigger-packet dhcp dhcp6
        authentication-policy "radius-1"
      exit
    no shutdown
  exit
exit
exit
exit
```

A full description of the managed SAP functionality is out of the scope of this chapter.

RADIUS Change of Authorization (CoA)

The only CoA action that is allowed for IPoE IPv6 hosts is a change of ESM strings (SLA-profile, subscriber-profile, application-profile, etc). Creation of a new IPv6 host or forcing a DHCPv6 renew is not supported.

Only a single address attribute (Framed-IP-Address, Delegated-IPv6-Prefix or Allocated-IPv6-Address) may be given in a single request. When host-accounting is enabled, only the host specific accounting session IDs (Acct-Session-Id) can be used. This means that to change for example the sla-profile for all three hosts of a dual stack subscriber, three CoA messages should be sent.

A full description of the RADIUS CoA functionality is out of the scope of this section.

Accounting

There are no separate accounting statistics available for IPv4 and IPv6 traffic unless they are mapped in a different Forwarding Class/queue.

In RADIUS accounting, host-accounting could be enabled to see the IPv4 and IPv6 host instantiations separately: an accounting start/stop is generated for each individual subscriber host. The actual accounting data is included in the interim updates and accounting stop message for the sla-profile instance.

A full description of the accounting functionality is out of the scope of this section.

Lease State Persistency

A DHCPv4/DHCPv6 (hereafter referred to as DHCP) session does not have a keep-alive mechanism to detect unavailability. A new DHCP session set-up is only attempted after expiration of the DHCP lease time. A node reboot causing the loss of DHCP lease state and the corresponding anti-spoof filters could therefore result in unacceptable long service outages.

The DHCP lease state can be made persistent across node reboots: DHCP lease state is restored from a persistency file stored on the compact flash file system. As a result, DHCP sessions will only lose connectivity during the time of reboot without being completely disconnected.

To activate the DHCP lease state persistency:

```
configure
  system
    persistence
      subscriber-mgmt
        description "DHCP lease state persistency"
        location cf1:
      exit
    exit
  exit
exit
```

A dedicated persistency file will be created on the specified compact flash file system. The file is initialized to store the maximum number of allowed hosts; its size is fixed to avoid file system space problems during operations.

```
*A:BNG# file dir cf1:

Volume in drive cf1 on slot A has no label.

Volume in drive cf1 on slot A is formatted as FAT32

Directory of cf1:\

09/19/2016  04:29p      <DIR>          .ssh/
04/20/2017  03:02p      536871424  submgmt.012
04/20/2017  02:59p      12583424  submgmt.i12
               2 File(s)          549454848 bytes.
               1 Dir(s)          7464747008 bytes free.

*A:BNG#
```

Each time the DHCP lease is renewed, the persistency file is updated together with the lease state. If the file update fails, an event is generated to indicate that persistency cannot be guaranteed.

The format of the persistency file may vary between different SR OS software releases. When upgrading, the persistency file is automatically upgraded to the new format. To downgrade the persistency file to a lower SR OS Release version, use the following command:

```
*A:BNG# tools perform persistence downgrade target-version ?
- downgrade target-version <target> [reboot]

<target>          : the version you want to downgrade to
                    submgmt
                      14.0 (current) - cf1:\submgmt.012
                      13.0           - cf1:\submgmt.011
                      12.0           - cf1:\submgmt.010
```

```

11.0      - cfl:\submgmt.009
10.0      - cfl:\submgmt.008
9.0       - cfl:\submgmt.007
8.0       - cfl:\submgmt.006
7.0       - cfl:\submgmt.005
6.0       - cfl:\submgmt.004
5.0       - cfl:\submgmt.003
4.0       - cfl:\submgmt.pst

<reboot>      : reboot system after successful conversion

*A:BNG#

```

The content of the persistency file can be looked at using the following commands:

```

*A:BNG# show service id 1 dhcp6 lease-state detail

=====
DHCP lease states for service 1
=====
Service ID      : 1
IP Address      : 2001:db8:a001:100::/56
Client HW Address : 00:0c:29:00:00:21
Subscriber-interface : sub-int-1
Group-interface  : grp-int-1
SAP              : 1/1/1:1
Termination Type : local
Up Time          : 0d 00:01:49
Remaining Lease Time : 0d 23:58:11
Remaining SessionTime: N/A
Persistence Key : 0x00000002

Sub-Ident       : "sub-1"
Sub-Profile-String : "sub-profile-1"
SLA-Profile-String : "sla-profile-1"
App-Profile-String : ""
Lease ANCP-String : ""
Lease Int Dest Id : ""
Category-Map-Name : ""
Dhcp6 ClientId (DUID): 00010001208a25ac000c29000021
Dhcp6 IAID      : 1
Dhcp6 IAID Type  : prefix
Dhcp6 Client Ip  : fe80::20c:29ff:fe00:21
Primary-Dns      : 2001:db8:dddd:1::1
Secondary-Dns    : 2001:db8:dddd:2::1
Pool Name        : ""
Dhcp6 Server Addr : N/A
Dhcp6 ServerId (DUID): N/A
Dhcp6 InterfaceId : N/A
Dhcp6 RemoteId    : N/A
Radius sub-if prefix : N/A
Router adv. policy : N/A

Lease Info origin : Radius

ServerLeaseStart : 04/20/2017 14:44:01
ServerLastRenew  : 04/20/2017 14:44:01
ServerLeaseEnd   : 04/21/2017 14:44:01
Session-Timeout  : N/A
IPoE|PPP session : No
Radius User-Name : "00:0c:29:00:00:21"

-----
Service ID      : 1
IP Address      : 2001:db8:a001:300::/56

```

```
Client HW Address : 00:0c:29:00:00:23
Subscriber-interface : sub-int-1
Group-interface : grp-int-2
SAP : 1/1/1:2
Termination Type : local
Up Time : 0d 00:01:36
Remaining Lease Time : 0d 23:58:24
Remaining SessionTime: N/A
Persistence Key : 0x00000005
```

```
Sub-Ident : "sub-3"
Sub-Profile-String : "sub-profile-1"
SLA-Profile-String : "sla-profile-1"
```

---snip---

Number of lease states : 4
=====

*A:BNG#

*A:BNG# tools dump persistence submgt record 0x2

Persistence Record

```
Client : submgt
Persist-Key : 0x00000002
Filename : cfl:\submgt.012
Entries : Index FedHandle Last Update Action Valid
          000002 0x00000002 2017/04/20 12:45:24 (UTC) ADD Yes
Data : 366 bytes
```

```
Host Type : IPv6 node address
Service ID : 1
SAP ID : 1/1/1:1
NH MAC : 00:0c:29:00:00:21
Created : 2017/04/20 12:44:01 (UTC)
IP : 2001:db8:a001:100::/56
Srvr Last Renew: 2017/04/20 12:44:01 (UTC)
Srvr Lse End : 2017/04/21 12:44:01 (UTC)
Dhcp6 Pfx len : 56
Dhcp6 Iaid : 1
Dhcp6 Iaid Typ : 25
Dhcp6 Client Mg: fe80::20c:29ff:fe00:21
Dhcp6 Client Id: 00010001208a25ac000c29000021
RADIUS Fallback: NO
Acct-Sess-Id : 14F2FF0000003658F8AD11
Multi-Sess-Id : 14F2FF0000003458F8AD0A
Class Attr : 0 bytes
User-Name : "00:0c:29:00:00:21"
host is authenticated by radius: true
Subscriber-Id : "sub-1"
Sub-Profile-Str: "sub-profile-1"
SLA-Profile-Str: "sla-profile-1"
Ipv6 Primary Dns: 2001:db8:dddd:1::1
Ipv6 Secondary Dns: 2001:db8:dddd:2::1
Ipv6 Delegated Prefix Origin: Radius
PD Server validLifeTime: 86400
PD Server preferredLifeTime: 3600
```

*A:BNG#

Conclusion

This chapter provides configuration, operation, and troubleshooting commands for dual stack IPoE subscribers on Routed Gateways. Focus is on the ESMv6 part where DHCPv6 is used for IPv6 address assignment on the RG network interface (wan host) and for allocation of an IPv6 prefix delegation prefix for use in the home network (pd host). In the BNG, authentication, authorization and IPv6 prefix configuration for an IPoE IPv6 host is done by a local user database or RADIUS.

ESMv6: PPPoE Dual Stack Hosts

This chapter describes ESMv6 PPPoE dual stack host configurations.

Topics in this chapter include:

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

Applicability

This chapter is applicable to SR OS routers and was initially based on Release 8.0.R4. The CLI is updated to Release 15.0.R1.

Prerequisites:

- Routed CO (IES/VRN service) with Enhanced Subscriber Management (ESM)
- Bridged or routed home gateway



Note:

The focus of this chapter is on PPPoE IPv6. IPv4 configuration is shown for completeness.


Overview

PPPoE Dual Stack

A PPPoE dual stack subscriber may support both IPv4 and IPv6 simultaneously. The dual stack hosts share a common subscriber identification policy and have a common sla-profile and subscriber-profile and are linked together through one PPPoE session.

For PPPoE dual stack hosts, one subscriber host is created for IPv4 and another one for the IPv6 address family.

Figure 69: PPPoE Dual Stack Hosts



ESM for IPv6 is supported through RADIUS and local user database (LUDB) for authentication, address assignment and authorization.

PPPoE dual stack subscriber-hosts are supported for bridged and routed home gateways.

Dual Stack PPPoE Bridged Gateway Service

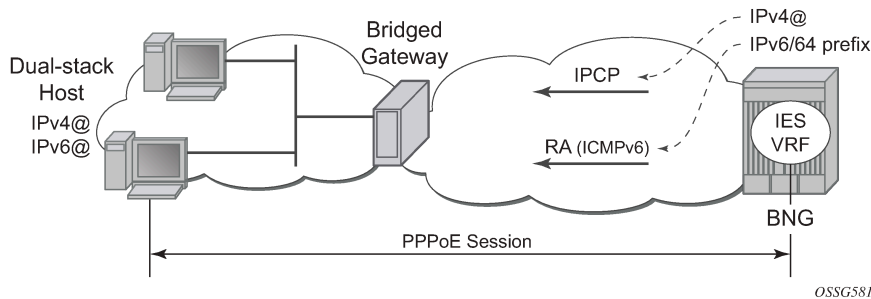
In the dual stack PPPoE host service, the PPPoE session is initiated directly from a dual stack device in the home network. PPPoE is used to carry IPv6 and (optionally) IPv4 traffic from the device to the broadband remote access server (BRAS), also called broadband network gateway (BNG).

Unlike the routed gateway application examples (see later), no IPv6 prefix delegation occurs in the bridged gateway service. Instead, a global unicast address prefix (/64) is advertised using Router Advertisements (RAs) directly to the PPPoE interface on the host.

The device addresses are self-assigned through stateless auto configuration (SLAAC), where SLAAC makes use of ICMPv6 router-advertisements to announce these IPv6 prefixes. The SLAAC prefixes have a mandatory length of /64.

This application is targeted at operators who currently use a bridging modem in the customer premises and who want to incrementally add IPv6 capability without a change of the modem on the customer site.

Figure 70: Dual Stack PPPoE Bridged Gateway Service Example

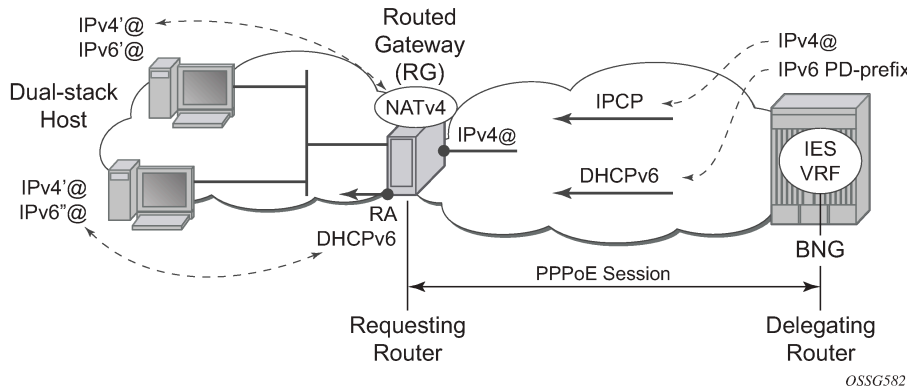


Dual Stack PPPoE Routed Gateway Service

The dual stack PPPoE routed gateway service runs over a dual stack PPPoE session between a dual stack router and BNG. It allows operators using PPPoE in their networks (with either PPPoE to the RG or PPPoA with translation to PPPoE in the DSLAM) to deploy IPv6 services in conjunction with existing IPv4 services.

Because a routed RG is used, a unique subscriber IPv6 prefix is delegated to the dual stack router for use within the home network. DHCPv6 is used to provide prefix delegation (PD). No WAN IPv6 address assignment is supported in this model. The dual stack router does not perform any NAT for IPv6 traffic.

Figure 71: Dual Stack PPPoE Routed Gateway Service Example



SLAAC

The IPv6 stateless auto configuration (SLAAC) mechanism requires no manual configuration of hosts, minimal configuration of routers, and no additional servers (such as DHCP). The stateless mechanism allows a host to generate its own address using a combination of locally available information and information advertised by routers. Routers advertise /64 prefixes, by an ICMPv6 router advertisement, that identify the subnet(s) associated with a link, while hosts generate a 64-bit "interface identifier" that uniquely identifies an interface on a subnet. An address is formed by combining the two.

DHCPv6

The Dynamic Host Configuration Protocol for IPv6 (DHCPv6) is defined in RFC 3315. The protocol enables DHCPv6 servers to pass configuration parameters such as IPv6 network addresses or DNSv6 addresses to IPv6 nodes.

For further information on DHCPv6, see [ESMv6: IPoE Dual Stack Hosts](#).

Prefix Delegation

Prefix Delegation (PD) is a mechanism for automated delegation of IPv6 prefixes using DHCPv6. A delegating router delegates a long-lived IPv6 prefix to a requesting router. The delegating router does not require knowledge about the topology of the links in the network to which the prefixes will be assigned.

For further information on Prefix Delegation, see [ESMv6: IPoE Dual Stack Hosts](#).

Configuration

ESMv6 for PPPoE is applicable in a routed CO environment. Details of non-specific dual stack configurations like authentication-policies, sla-profile, subscriber-profile, accounting-policies and QoS policies are out of scope for this chapter.

The minimal RADIUS authentication configuration and ESM string configuration is added for completeness.

```
configure
  router
    radius-server
      server "radius-172.16.1.2" address 172.16.1.2 secret vsecret1 create
      accept-coa
    exit
  exit
exit
```

```
configure
  aaa
    radius-server-policy "rsp-1" create
    servers
      router "Base"
      source-address 192.0.2.1
      server 1 name "radius-172.16.1.2"
    exit
  exit
exit
```

```
configure
  subscriber-mgmt
    authentication-policy "auth-1" create
    description "RADIUS authentication policy"
    pppoe-access-method pap-chap
    radius-server-policy "rsp-1"
  exit
exit
```

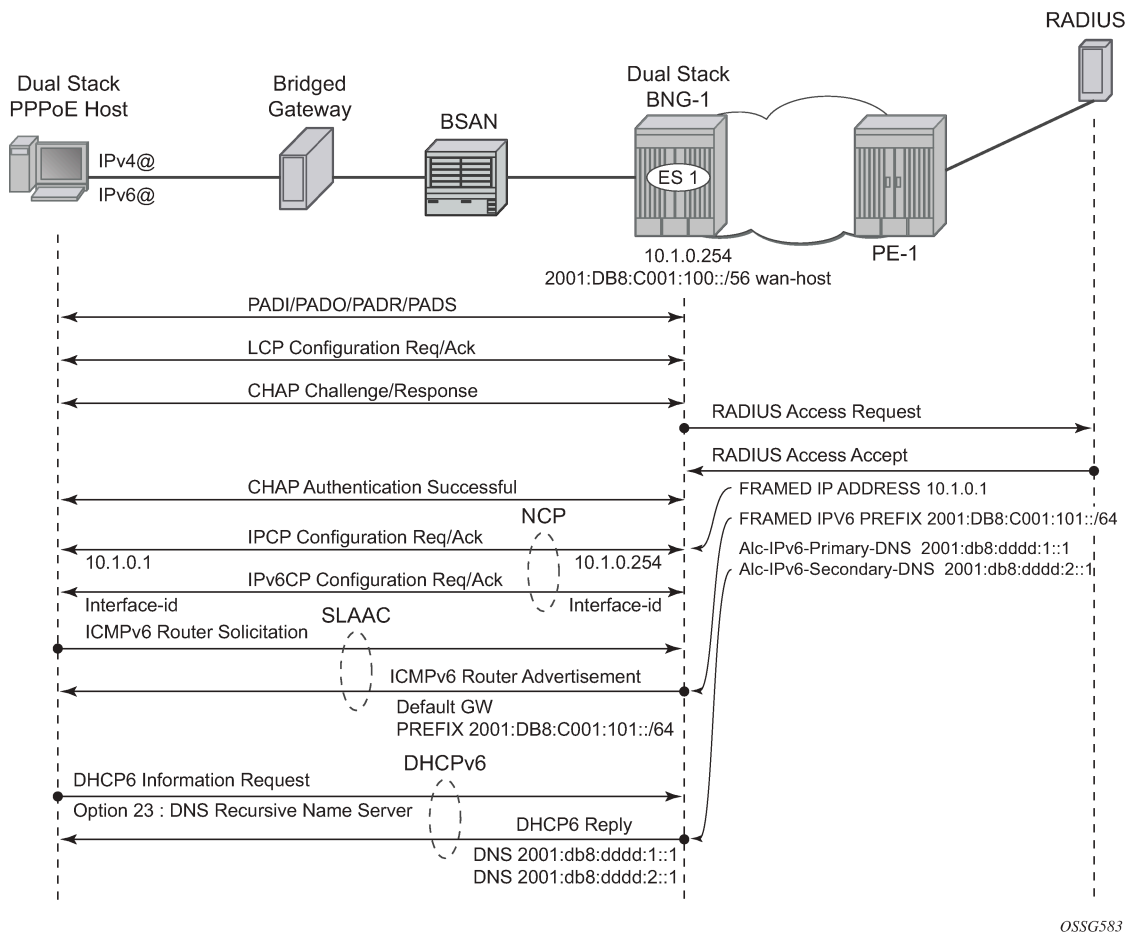
```
configure
  subscriber-mgmt
    sla-profile "sla-profile-1" create
    exit
    sub-profile "sub-profile-1" create
    exit
    sub-ident-policy "sub-ident-1" create
    sub-profile-map
      use-direct-map-as-default
    exit
    sla-profile-map
      use-direct-map-as-default
    exit
  exit
exit
```

Service

Dual Stack PPPoE for Bridged Gateway

Figure 72: Message Flow for a Dual Stack PPPoE Host shows the message flow for a dual stack PPPoE host behind a bridged gateway corresponding with the configured service.

Figure 72: Message Flow for a Dual Stack PPPoE Host



For dual stack PPPoE, the BNG initiates the IPv6 control protocol (IPv6CP) protocol to the client during the session setup phase if the appropriate attributes have been returned by the RADIUS server on authentication. The RADIUS attribute that triggers the setup of a dual stack PPPoE host in bridged mode is **framed-ipv6-prefix** which should contain a /64 prefix for the client. When a PPPoE host has successfully completed the IPv6CP negotiation, the BNG will transmit an RA to the PPPoE host containing the prefix and any other option that is configured. The host can request optional IPv6 DNS server information from the BNG by sending a DHCPv6 information-request.

The following example shows a minimal configuration to enable dual stack subscribers in an IES service context with the ESM IPv6-specific parts in bold.

```
configure
service
  ies 1 customer 1 create
    subscriber-interface "sub-int-1" create
      address 10.1.0.254/16
      ipv6
        subscriber-prefixes
          prefix 2001:db8:c001:100::/56 wan-host
        exit
      exit
    group-interface "group-int-1" create
      ipv6
        router-advertisements
          prefix-options
            autonomous
          exit
          no shutdown
        exit
        dhcp6
          proxy-server
            client-applications ppp
            no shutdown
          exit
        exit
      exit
    authentication-policy "auth-1"
    sap 1/1/1:1 create
      sub-sla-mgmt
        sub-ident-policy "sub-ident-1"
        multi-sub-sap 10
        no shutdown
      exit
    exit
  pppoe
    session-limit 10
    sap-session-limit 10
    no shutdown
  exit
exit
no shutdown
exit
exit
```

IPv6 subscriber prefixes must be defined in the **subscriber-interface>ipv6>subscriber-prefixes** context.

Three types of prefixes can be configured where **wan-host** is required for the bridged gateway scenario and **pd** is used for the dual stack PPPoE routed gateway scenario.

- **wan-host** — Prefix from which the IPv6 addresses are assigned (by DHCPv6 IA_NA) for the IPoEv6 routed gateway WAN interface (network facing) or a prefix from which /64 prefixes are assigned for the PPPoE (by RA SLAAC) hosts in the bridged gateway model.
- **pd** — Prefix from which the IPv6 prefix delegation prefixes are assigned that are to be used by the IPoEv6 or PPPoEv6 routed gateway for allocation in the home network (LAN interfaces).
- **pd wan-host (both)** — Prefix from which both IPv6 addresses (**wan-host**) and IPv6 prefix delegation prefixes (**pd**) can be assigned. This requires that the delegated prefix length is set to 64 bits.

[Table 15: Subscriber Prefix Parameters](#) and [Table 16: Subscriber Prefix Subnetting for SLAAC](#) provide an overview of the subscriber-prefix parameters that apply and an example of subscriber prefix subnetting for SLAAC.

Table 15: Subscriber Prefix Parameters

Subscriber Prefix Type	Prefix Length	DHCPv6 Option	SLAAC	RADIUS AVP	Must be subnetted as
wan-host	/32..63	N/A	yes	[97]Framed-IPv6-Prefix	/64


Table 16: Subscriber Prefix Subnetting for SLAAC

Subscriber prefix	Framed-IPv6-Prefix	Hosts	
2001:db8:c001:100::/56	2001:db8:c001:101::/64	pppoev6-host-1	
	2001:db8:c001:102::/64	pppoev6-host-2	
	2001:db8:c001:103::/64	pppoev6-host-3	
	<snip>	<snip>	
	2001:db8:c001:1FF::/64	pppoev6-host-256	
2001:db8:c001:200::/56	2001:db8:c001:201::/64	PPPoEv6-host-257	
	2001:db8:c001:202::/64	PPPoEv6-host-258	
	2001:db8:c001:203::/64	PPPoEv6-host-259	
	<snip>	<snip>	
	2001:db8:c001:2FF::/64	PPPoEv6-host-512	

Dual Stack PPPoE for Routed Gateway

[Figure 73: Dual Stack PPPoE for Routed Gateway](#) shows the message flow for a dual stack PPPoE host located behind a routed gateway corresponding with the configured service.

Figure 73: Dual Stack PPPoE for Routed Gateway



OSSG677

Initially, a PPPoE routed gateway follows the same steps as a dual stack PPPoE host. The BNG receives a prefix from RADIUS (in this case through a **delegated-ipv6-prefix** attribute), which is used as a trigger to initiate the IPv6CP protocol to the client. The prefix that is offered to the client should have the same prefix length as the one configured under the subscriber interface ipv6 context (delegated-prefix length). This length should be between 48 and 64 bits, inclusive.

After the IPv6CP protocol has completed, the client must run the DHCPv6 protocol over its PPPoE tunnel to receive a delegated prefix (IA_PD) and optionally IPv6 DNS server information.

This delegated prefix can then be subdivided by the client and distributed over its own downstream interfaces. During the DHCPv6 message exchange, no extra RADIUS request will be made; the information is stored during the initial PPPoE authentication until the client starts DHCPv6. Only after DHCPv6 has completed, the IPv6 subscriber host will be instantiated, and the BNG will start sending RAs if configured. (It is a mandatory requirement for the BNG to send RAs which makes enabling router-advertisements under the group-level mandatory). The router advertisements do not contain any prefix information, which has already been provided by DHCPv6, but it is used as an indication to the client that its default gateway should be the BNG.

```
configure
service
  ies 1 customer 1 create
    subscriber-interface "sub-int-1" create
      address 10.1.0.254/16
```

```

        ipv6
        delegated-prefix-len 56
        subscriber-prefixes
            prefix 2001:db8:d001::/48 pd
        exit
    exit
    group-interface "group-int-1" create
    ipv6
        router-advertisements
            prefix-options
                autonomous
            exit
        no shutdown
    exit
    dhcp6
        proxy-server
            client-applications ppp
            no shutdown
        exit
    exit
    authentication-policy "auth-1"
    sap 1/1/1:1 create
        sub-sla-mgmt
            sub-ident-policy "sub-ident-1"
            multi-sub-sap 10
            no shutdown
        exit
    exit
    pppoe
        session-limit 10
        sap-session-limit 10
        no shutdown
    exit
exit
exit
no shutdown
exit
exit
exit
exit

```

IPv6 subscriber prefixes must be defined at the **subscriber-interface>ipv6>subscriber-prefixes** context, see [Dual Stack PPPoE Bridged Gateway Service](#).

Subscriber prefixes are subnetted in fixed length subnets that are assigned to subscriber hosts:

- /delegated-prefix-len (/48..64) for p subscriber prefixes

The delegated prefix length is configured in the **subscriber-interface>ipv6** context. The recommended value is /56 (default = /64). The configured length applies to all pd subscriber prefixes on a subscriber-interface.

[Table 17: Subscriber-Prefix Parameters](#) and [Table 18: Prefix Subnetting for delegated-prefix-length /56](#) provide an overview of the subscriber-prefix parameters that apply and an example of prefix subnetting for delegated-prefix-length /56.

Table 17: Subscriber-Prefix Parameters

Subscriber Prefix Type	Prefix Length	DHCPv6 Option	SLAAC	RADIUS AVP	Must be sub netted as
pd	/48..64 *	IA-PD	N/A	[123] Delegated-IPv6-Prefix	/delegated-prefix-len

*Must be smaller than configured delegated prefix length.

Table 18: Prefix Subnetting for delegated-prefix-length /56

Subscriber Prefix and /56 delegated-prefix-len	Framed-IPv6-Prefix	Hosts
2001:db8:d001::/48	2001:db8:d001:100::/56	Responsibility Home Gateway (HGW)
		Responsibility HGW
		Responsibility HGW
	2001:db8:d001:200::/56	Responsibility HGW
		Responsibility HGW
		Responsibility HGW
	<snip>	--
	2001:db8:d001:FF00::/56	Responsibility HGW
		Responsibility HGW
		Responsibility HGW
	<snip>	--
	2001:db8:d002:100::/56	Responsibility HGW
		Responsibility HGW
		Responsibility HGW
	2001:db8:d002:200::/56	Responsibility HGW
		Responsibility HGW
		Responsibility HGW
	<snip>	--
	2001:db8:d002:FF00::/56	Responsibility HGW

Subscriber Prefix and /56 delegated-prefix-len	Framed-IPv6-Prefix	Hosts
		Responsibility HGW
		Responsibility HGW

RADIUS

The RADIUS authentication policy shown at the beginning of the [Configuration](#) section must be applied to the group-interface, and is used for both IPv4 and IPv6.

```
configure
  service
    ies 1 customer 1 create
      subscriber-interface "sub-int-1" create
      group-interface "group-int-1" create
      authentication-policy "auth-1"
    exit
  exit
exit
exit
exit
exit
```

IPv4 and IPv6 configuration information can come from LUDB or AAA/RADIUS.

Commonly used RADIUS Attribute Value pairs (AVPs) that are applicable for PPPoE IPv6 subscriber hosts are listed in [Table 19: RADIUS AVPs](#).

Table 19: RADIUS AVPs

RADIUS AVP	Type	Purpose
Framed-IPv6-Prefix [97]	ipv6prefix	Maps to SLAAC (RFC 4862) /64 Prefix-information in ICMPv6 RA.
Delegated-IPv6-Prefix [123]	ipv6prefix	Maps to IA_PD for prefix delegation (RFC 3633) in DHCPv6
Alc-IPv6-Primary-Dns [26-6527-105]	ipv6addr	Maps to DNS recursive name server option (RFC 3646) in DHCPv6
Alc-IPv6-Secondary-Dns [26-6527-106]	ipv6addr	Maps to DNS recursive name server option (RFC 3646) in DHCPv6

Dual Stack PPPoE for Bridged Gateway

The following shows a sample of a FreeRADIUS user record to authenticate a dual stack PPPoE subscriber for a bridged gateway:

```
bridged@domain1 Cleartext-Password := "letmein"
  Framed-IP-Address = 10.1.0.1,
  Framed-IP-Netmask = 255.255.255.0,
```

```
Alc-Subsc-ID-Str = "%{User-name}",  
Alc-Subsc-Prof-Str = "sub-profile-1",  
Alc-SLA-Prof-Str = "sla-profile-1",  
Framed-IPv6-Prefix = "2001:db8:c001:0101::/64",  
Alc-IPv6-Primary-DNS = "2001:db8:dddd:1::1",  
Alc-IPv6-Secondary-DNS = "2001:db8:dddd:2::1",
```

Dual Stack PPPoE for Routed Gateway

The following shows a sample of a FreeRADIUS user record to authenticate a dual stack PPPoE subscriber for a routed gateway:

```
routed@domain1 Cleartext-Password := "letmein"  
Framed-IP-Address = 10.1.0.2,  
Framed-IP-Netmask = 255.255.255.0,  
Alc-Subsc-ID-Str = "%{User-name}",  
Alc-Subsc-Prof-Str = "sub-profile-1",  
Alc-SLA-Prof-Str = "sla-profile-1",  
Delegated-IPv6-Prefix = "2001:db8:d001:0100::/56",  
Alc-IPv6-Primary-DNS = "2001:db8:dddd:1::1",  
Alc-IPv6-Secondary-DNS = "2001:db8:dddd:2::1",
```

A RADIUS user's configuration with multiple delegated-ipv6-prefixes for the same dual stack PPPoE host will result in a single DHCPv6 advertise message sent by the BNG with a single IA_PD option and single IA-Prefix. The other RADIUS configured delegated-IPv6-prefixes are silently dropped by the BNG.

Router Advertisements

ICMPv6 router advertisements have two major functions.

- Default router function for hosts
- Address auto-configuration for hosts aka SLAAC

Unsolicited RA must explicitly be enabled on a group interface (default shutdown) and are refreshed with a pseudo random timer. The boundaries of this random timer are configurable with the min-advertisement parameter (minimum with default set to 900s) and max-advertisement (maximum with default set to 1800s).

```
configure  
  service  
    ies 1 customer 1 create  
      subscriber-interface "sub-int-1" create  
        group-interface "group-int-1" create  
          ipv6  
            router-advertisements  
              max-advertisement 1800 # default 30 min  
              min-advertisement 900 # default 15 min  
              no shutdown  
            exit  
          exit  
        exit  
      exit  
    exit  
  exit
```

The **router-advertisements router-lifetime** parameter (default 4500 sec) specifies how long the host is allowed to use the originator of the RA as default gateway. This timer is configurable between 2700 and 9000 seconds.

Configuring a **router-advertisements router-lifetime** timer smaller than the **router-advertisements min-advertisement** timer results in a dual stack PPPoE host without a default gateway.

```
configure
service
  ies 1 customer 1 create
    subscriber-interface "sub-int-1" create
    group-interface "group-int-1" create
    ipv6
      router-advertisements
        router-lifetime 4500
        no shutdown
      exit
    exit
  exit
exit
exit
exit
exit
```

The following **prefix-options autonomous** parameter specifies whether or not offered RADIUS IPv6 prefix can be used for stateless address configuration (SLAAC). The **prefix-options lifetime** parameter defines how long the host is allowed to use this prefix. Configuring a **prefix-option valid-lifetime** smaller than the **router-advertisements min-advertisement** timer results in host traffic being sourced with the link-local address instead of global unique IPv6 address.

```
configure
service
  ies 1 customer 1 create
    subscriber-interface "sub-int-1" create
    group-interface "group-int-1" create
    ipv6
      router-advertisements
        prefix-options
          autonomous          # required for SLAAC
          on-link
          preferred-lifetime 3600 # default 1 hour
          valid-lifetime 86400  # default 24 hours
        exit
      no shutdown
    exit
  exit
exit
exit
exit
exit
exit
exit
```

The following is a snapshot from an ICMPv6 RA message with default timer settings with a focus on the SLAAC function.

```
Internet Control Message Protocol v6
Type: 134 (Router advertisement)

---snip---

ICMPv6 Option (Prefix information)
```

```
Type: Prefix information (3)
Length: 32
Prefix length: 64
Flags: 0x40
    1... .... = on link
    .1.. .... = Auto                # Auto-Configuration flag
    ..0. .... = Not router address
    ...0 .... = Not site prefix
Valid lifetime: 86400              # Default value 24 hour
Preferred lifetime: 3600           # Default value 1 hour
Prefix: 2001:DB8:C001:101::       # SLAAC prefix
```

SLAAC-related parameters are listed in [Table 20: SLAAC-Related Parameters](#) .

Table 20: SLAAC-Related Parameters

Parameter	Description (RFC-4861)	Value Range (Default)
prefix-options: autonomous	Autonomous address-configuration flag. When set indicates that this prefix can be used for stateless address auto configuration (SLAAC)	(no)
prefix-options: preferred-lifetime	The length of time in seconds that the addresses generated from the prefix through stateless address auto configuration (SLAAC) remains preferred.	0..4294967295 s (3600s) 1hour
prefix-options: valid-lifetime	The length of time in seconds that the prefix is valid for the purpose of on-link determination.	0..4294967295 s (86400s) 24hours

Router advertisements parameters common to PPPoEv6 and IPoEv6 are listed and explained in [ESMv6: IPoE Dual Stack Hosts](#).

For dual stack PPPoE hosts, the default values, as shown in the following output, can be used. Timer values equal to zero (reachable-time and retransmit-time) causes the host to use its own timers for that function. The reachable-time is used by the host for Neighbor_Unreachable_Detection (NUD) whereas the retransmit-time is used by the host for Duplicate_Address_Detection (DAD). DAD is normally only performed by dual stack IPoE hosts and not by dual stack PPPoE hosts.

```
configure
service
  ies 1 customer 1 create
    subscriber-interface "sub-int-1" create
    group-interface "group-int-1" create
    ipv6
      router-advertisements
        current-hop-limit 64
        dns-options
          no include-dns
          rdns-lifetime 3600
        exit
        no force-mcast
        no managed-configuration
        no mtu
        no other-stateful-configuration
        reachable-time 0
        retransmit-time 0
```

```
no shutdown
exit
exit
exit
exit
exit
```

DHCPv6 Proxy Server

Dual Stack PPPoE for Bridged Gateway

Dual stack PPPoE hosts using SLAAC for address assignment do not require DHCPv6. SR OS supports DNSv6 information through the RA DNS Option (RFC 5006, *IPv6 Router Advertisement Option for DNS Configuration*), and can also be configured to use DHCPv6 information requests and replies to retrieve the DNSv6 information. This requires the DHCPv6 proxy server to be enabled (the default is **shutdown**) and PPPoE defined as client-application (default=dhcp only). No lease state is kept for this DNSv6 information and therefore it is known as stateless DHCPv6.

```
configure
service
  ies 1 customer 1 create
    subscriber-interface "sub-int-1" create
    group-interface "group-int-1" create
    ipv6
      dhcp6
        proxy-server
        client-applications ppp
        no shutdown
      exit
    exit
  exit
exit
exit
exit
exit
exit
exit
```

Dual Stack PPPoE for Routed Gateway

An IPv6 PPPoE routed gateway initiates, after successful IPv6CP negotiation, a DHCPv6 session to request its configuration data (IPv6 PD prefixes, DNS servers). A DHCPv6 proxy server in the BNG maintains the DHCPv6 session with the IPv6 PPPoE subscriber host. The DHCPv6 proxy server must be enabled (the default is **shutdown**) and PPPoE defined as client-application (default=dhcp only).

```
configure
service
  ies 1 customer 1 create
    subscriber-interface "sub-int-1" create
    group-interface "group-int-1" create
    ipv6
      dhcp6
        proxy-server
        server-id duid-ll
        renew-timer min 30          # default
```



```


rebind-timer min 48      # default
valid-lifetime days 1    # default
preferred-lifetime hrs 1 # default
client-applications ppp
no shutdown
exit
exit

```

A number of timers associated with IPv6 addresses and IPv6 prefixes within DHCPv6 identity associations can be configured in the DHCPv6 proxy server context. These timers are valid for IPoEv6 and PPPoEv6 sessions and are listed and further explained in [ESMv6: IPoE Dual Stack Hosts](#).

There is never RADIUS re-authentication for dual stack PPPoE routed gateways on DHCPv6 renewals as indicated in [Figure 74: DHCPv6 Renewals](#).

Figure 74: DHCPv6 Renewals



DHCPv6 Lease State

The DHCPv6 lease state table keeps track of the DHCPv6 host states. The DHCP lease information for a specific host is extracted from the DHCPv6 reply message in case of DHCPv6. Stateful (with lease state) DHCPv6 is applicable for dual stack PPPoE on routed gateway where Stateless (without lease state) DHCPv6 is optional and applicable for dual stack PPPoE on bridged gateways.

For more information on DHCPv6 lease states, see [ESMv6: IPoE Dual Stack Hosts](#).

Operation

Dual Stack PPPoE for Bridged Gateway

A PPPoEv6 dual stack subscriber scenario for a bridged home gateway consumes two subscriber host entries shared by a common subscriber.

- IPv4 host-addressing by IPCP
- IPv6 wan-host addressing by SLAAC

```
*A:BNG# show service active-subscribers subscriber bridged@domain1"bridged@domain1"

=====
Active Subscribers
=====
Subscriber bridged@domain1 (sub-profile-1)
-----
(1) SLA Profile Instance sap:1/1/1:1 - sla:sla-profile-1
-----
IP Address          MAC Address          Session      Origin      Svc      Fwd
-----
10.1.0.1            00:0c:29:00:00:11    PPP 1        IPCP         1         Y
2001:db8:c001:101::/64 00:0c:29:00:00:11    PPP 1        SLAAC        1         Y
-----

=====
*A:BNG#
```

The **hierarchy** parameter for active-subscribers gives a top level down overview for this subscriber.

```
*A:BNG# show service active-subscribers hierarchy subscriber "bridged@domain1"

=====
Active Subscribers Hierarchy
=====
-- bridged@domain1 (sub-profile-1)
|
+-- sap:1/1/1:1 - sla:sla-profile-1
|
+-- PPP-session - mac:00:0c:29:00:00:11 - sid:1 - svc:1
|
|   |-- 10.1.0.1 - IPCP
|   |
|   +-- 2001:db8:c001:101::/64 - SLAAC
|
=====
*A:BNG#
```

IPCP and IPv6CP are in an opened state for the dual stack PPPoE session and their origin is RADIUS, as shown below.

```
*A:BNG# show service id 1 pppoe session ip-address 10.1.0.1 detail

=====
PPPoE sessions for svc-id 1
=====
Sap Id          Mac Address      Sid    Up Time          Type
  IP/L2TP-Id/Interface-Id          MC-Stdbby
-----
1/1/1:1         00:0c:29:00:00:11 1       0d 00:01:27      local
  10.1.0.1
  02:0C:29:FF:FE:00:00:11

LCP State       : Opened
IPCP State      : Opened
IPv6CP State    : Opened
PPP MTU         : 1492
PPP Auth-Protocol : CHAP
PPP User-Name   : bridged@domain1

Subscriber-interface : sub-int-1
Group-interface     : group-int-1

IP Origin        : radius
DNS Origin       : none
NBNS Origin      : none

Subscriber       : "bridged@domain1"
Sub-Profile-String : "sub-profile-1"
SLA-Profile-String : "sla-profile-1"

---snip---

IPv6 Prefix      : 2001:db8:c001:101::/64
IPv6 Prefix Origin : radius
IPv6 Prefix Pool  : ""
IPv6 Del.Pfx.    : N/A
IPv6 Del.Pfx. Origin : none
IPv6 Del.Pfx. Pool : ""
IPv6 Address     : N/A
IPv6 Address Origin : none
IPv6 Address Pool : ""
Primary IPv6 DNS  : 2001:db8:dddd:1::1
Secondary IPv6 DNS : 2001:db8:dddd:2::1

---snip---

-----
Number of sessions : 1
=====
*A:BNG#
```

The IPv6 routing table for dual stack hosts is displayed using the **protocol** keyword **sub-mgmt**.

```
*A:BNG# show router route-table ipv6 protocol sub-mgmt

=====
IPv6 Route Table (Router: Base)
=====
Dest Prefix[Flags]          Type    Proto    Age          Pref
  Next Hop[Interface Name]          Metric
```

```
-----
2001:db8:c001:101::/64          Remote  Sub Mgmt  00h01m37s  0
      [group-int-1]                      0
-----
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
      B = BGP backup route available
      L = LFA nexthop available
      S = Sticky ECMP requested
=====
*A:BNG#
```

DNSv6

DNSv6 information, in a dual stack PPPoE bridged gateway model, is optionally retrieved through stateless DHCPv6 information requests. Debugging is done through debug commands or/and observation by statistics counters.

```
100 2017/04/19 14:16:40.01 CEST MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
  Incoming DHCP6 Msg : INFO_REQUEST (11)
  on itf group-int-1
    Trans Id : 0xcef3f0
    Option : CLIENTID (1), Length : 14
      LLT : HwTyp=0001,T=322878930,LL=000c29c851ca
      00010001133ebdd2000c29c851ca
    Option : ELAPSED_TIME (8), Length : 2
      Time : 100 seconds
    Option : OR0 (6), Length : 4
      Requested Option : DNS_NAME_SRVR (23)

101 2017/04/19 14:16:40.02 CEST MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
  Outgoing DHCP6 Msg : REPLY (7)
  to itf group-int-1
    Trans Id : 0xcef3f0
    Option : SERVERID (2), Length : 10
      LL : HwTyp=0001,LL=24b1ff000000
      0003000124b1ff000000
    Option : CLIENTID (1), Length : 14
      LLT : HwTyp=0001,T=322878930,LL=000c29c851ca
      00010001133ebdd2000c29c851ca
    Option : DNS_NAME_SRVR (23), Length : 32
      Server : 2001:db8:dddd:1::1
      Server : 2001:db8:dddd:2::1
```

DHCPv6 statistics can be shown as follows:

```
*A:BNG# show router dhcp6 statistics

=====
DHCP6 statistics (Router: Base)
=====
Msg-type          Rx          Tx          Dropped
-----
1 SOLICIT         0           0           0
2 ADVERTISE       0           0           0
3 REQUEST         0           0           0
4 CONFIRM         0           0           0
5 RENEW           0           0           0
```

```

6 REBIND                0                0                0
7 REPLY                 0                0                0
8 RELEASE               0                0                0
9 DECLINE               0                0                0
10 RECONFIGURE          0                0                0
11 INFO_REQUEST         0                0                0
12 RELAY_FORW          0                0                0
13 RELAY_REPLY         0                0                0
14 LEASEQUERY           0                0                0
15 LEASEQUERY_REPLY     0                0                0

-----
Dhcp6 Drop Reason Counters :
-----
1 Dhcp6 oper state is not Up on src itf                0
2 Dhcp6 oper state is not Up on dst itf                0
3 Relay Reply Msg on Client Itf                        0

---snip---

38 Packet dropped by DHCP filter                        0
39 Packet dropped because authentication failed          0
=====
*A:BNG#

```

To clear the statistics use following command:

```
*A:BNG# clear router dhcp6 statistics
```

Entries, for dual stack PPPoE subscribers, in the IPv4 ARP and/or IPv6 neighbor cache table are counted as internal entries and are shown from the **summary** parameter.

```

*A:BNG# show router arp summary

=====
ARP Table Summary (Router: Base)
=====
Local ARP Entries   : 3
Static ARP Entries  : 0
Dynamic ARP Entries : 1
Managed ARP Entries : 0
Internal ARP Entries : 1
BGP-EVPN ARP Entries : 0
-----
No. of ARP Entries   : 5
=====
*A:BNG#

```

```

*A:BNG# show router neighbor summary

=====
Neighbor Table Summary (Router: Base)
=====
Static Nbr Entries   : 0
Dynamic Nbr Entries  : 0
Managed Nbr Entries : 0
Internal Nbr Entries : 1
Evpn Nbr Entries     : 0
-----
No. of Neighbor Entries : 1
=====

```

```
*A:BNG#
```

Dual Stack PPPoE for Routed Gateway

A PPPoEv6 dual stack subscriber scenario for a routed CPE consumes two subscriber host entries sharing a common subscriber.

- IPv4 host addressing through IPCP
- IPv6 pd addressing through DHCPv6

```
*A:BNG# show service active-subscribers subscriber routed@domain1"routed@domain1"

=====
Active Subscribers
=====
-----
Subscriber routed@domain1 (sub-profile-1)
-----
(1) SLA Profile Instance sap:1/1/1:1 - sla:sla-profile-1
-----
IP Address
-----
MAC Address      Session      Origin      Svc      Fwd
-----
10.1.0.2          00:0c:29:00:00:12  PPP 1       IPCP      1        Y
2001:db8:d001:100::/56
00:0c:29:00:00:12  PPP 1       DHCP6-PD    1        Y
-----
=====
*A:BNG#
```

The hierarchy parameter for active-subscribers gives a top level down overview for this subscriber.

```
*A:BNG# show service active-subscribers hierarchy subscriber "routed@domain1"

=====
Active Subscribers Hierarchy
=====
-- routed@domain1 (sub-profile-1)
|
+-- sap:1/1/1:1 - sla:sla-profile-1
|
+-- PPP-session - mac:00:0c:29:00:00:12 - sid:1 - svc:1
|
|-- 10.1.0.2 - IPCP
|
+-- 2001:db8:d001:100::/56 - DHCP6-PD
=====
*A:BNG#
```

IPCP and IPv6CP are in an opened state for the dual stack PPPoE session and their origin is RADIUS, as shown below.

```
*A:BNG# show service id 1 pppoe session ip-address 10.1.0.2 detail
```

```

PPPoE sessions for svc-id 1
=====
Sap Id          Mac Address      Sid   Up Time      Type
IP/L2TP-Id/Interface-Id      MC-Stdby
-----
1/1/1:1        00:0c:29:00:00:12 1      0d 00:00:55   local
10.1.0.2
02:0C:29:FF:FE:00:00:12

LCP State       : Opened
IPCP State      : Opened
IPv6CP State    : Opened
PPP MTU         : 1492
PPP Auth-Protocol : CHAP
PPP User-Name   : routed@domain1

Subscriber-interface : sub-int-1
Group-interface     : group-int-1

IP Origin        : radius
DNS Origin       : none
NBNS Origin      : none

Subscriber       : "routed@domain1"
Sub-Profile-String : "sub-profile-1"
SLA-Profile-String : "sla-profile-1"

---snip---

IPv6 Prefix      : N/A
IPv6 Prefix Origin : none
IPv6 Prefix Pool  : ""
IPv6 Del.Pfx.    : 2001:db8:d001:100::/56
IPv6 Del.Pfx. Origin : radius
IPv6 Del.Pfx. Pool : ""
IPv6 Address     : N/A
IPv6 Address Origin : none
IPv6 Address Pool : ""
Primary IPv6 DNS  : 2001:db8:dddd:1::1
Secondary IPv6 DNS : 2001:db8:dddd:2::1

---snip---

-----
Number of sessions : 1
=====
*A:BN#

```

The IPv6 routing table for dual stack hosts is displayed using the **protocol** keyword **sub-mgmt**.

```

*A:BN# show router route-table ipv6 protocol sub-mgmt

=====
IPv6 Route Table (Router: Base)
=====
Dest Prefix[Flags]      Type   Proto   Age      Pref
Next Hop[Interface Name]      Metric
-----
2001:db8:d001:100::/56    Remote Sub Mgmt 00h01m01s 0
[group-int-1]              0
-----
No. of Routes: 1
Flags: n = Number of times nexthop is repeated

```

```
B = BGP backup route available
L = LFA nexthop available
S = Sticky ECMP requested
=====
*A:BN#
```

DNSv6

DNSv6 information, in a dual stack PPPoE routed gateway model, is optionally retrieved by stateful DHCPv6 information requests. Troubleshooting is done through debug commands or/and observation by statistics counters.

```
16 2017/04/19 09:38:54.26 CEST MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
  Incoming DHCP6 Msg : SOLICIT (1)
  on itf group-int-1
  Trans Id : 0xe3f882
  Option : CLIENTID (1), Length : 14
    LLT : HwTyp=0001,T=545839492,LL=000c29000012
    000100012088d984000c29000012
  Option : IA_PD (25), Length : 12
    IAID : 1
    Time1: 0 seconds
    Time2: 0 seconds
  Option : OR0 (6), Length : 2
    Requested Option : DNS_NAME_SRVR (23)
"
```

```
17 2017/04/19 09:38:54.26 CEST MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
  Outgoing DHCP6 Msg : ADVERTISE (2)
  to itf group-int-1
  Trans Id : 0xe3f882
  Option : SERVERID (2), Length : 10
    LL : HwTyp=0001,LL=14f2ff000000
    0003000114f2ff000000
  Option : CLIENTID (1), Length : 14
    LLT : HwTyp=0001,T=545839492,LL=000c29000012
    000100012088d984000c29000012
  Option : DNS_NAME_SRVR (23), Length : 32
    Server : 2001:db8:dddd:1::1
    Server : 2001:db8:dddd:2::1
  Option : IA_PD (25), Length : 41
    IAID : 1
    Time1: 1800 seconds
    Time2: 2880 seconds
  Option : IAPREFIX (26), Length : 25
    Prefix : 2001:db8:d001:100::/56
    Preferred Lifetime : 3600 seconds
    Valid Lifetime : 86400 seconds
"
```

```
18 2017/04/19 09:38:54.27 CEST MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
  Incoming DHCP6 Msg : REQUEST (3)
  on itf group-int-1
  Trans Id : 0x1971dc
  Option : CLIENTID (1), Length : 14
    LLT : HwTyp=0001,T=545839492,LL=000c29000012
    000100012088d984000c29000012
  Option : SERVERID (2), Length : 10
```



```

    LL : HwTyp=0001,LL=14f2ff000000
    0003000114f2ff000000
    Option : IA_PD (25), Length : 12
    IAID : 1
    Time1: 0 seconds
    Time2: 0 seconds
    Option : OR0 (6), Length : 2
    Requested Option : DNS_NAME_SRV (23)
"

19 2017/04/19 09:38:54.27 CEST MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
  Outgoing DHCP6 Msg : REPLY (7)
  to itf group-int-1
  Trans Id : 0x1971dc
  Option : SERVERID (2), Length : 10
  LL : HwTyp=0001,LL=14f2ff000000
  0003000114f2ff000000
  Option : CLIENTID (1), Length : 14
  LLT : HwTyp=0001,T=545839492,LL=000c29000012
  000100012088d984000c29000012
  Option : DNS_NAME_SRV (23), Length : 32
  Server : 2001:db8:dddd:1::1
  Server : 2001:db8:dddd:2::1
  Option : IA_PD (25), Length : 41
  IAID : 1
  Time1: 1800 seconds
  Time2: 2880 seconds
  Option : IAPREFIX (26), Length : 25
  Prefix : 2001:db8:d001:100::/56
  Preferred Lifetime : 3600 seconds
  Valid Lifetime : 86400 seconds
"

```

Use the following command to display the DHCPv6 statistics.

```
*A:BNG# show router dhcp6 statistics
```

```
=====
DHCP6 statistics (Router: Base)
=====
```

Msg-type	Rx	Tx	Dropped
1 SOLICIT	1	0	0
2 ADVERTISE	0	1	0
3 REQUEST	1	0	0
4 CONFIRM	0	0	0
5 RENEW	0	0	0
6 REBIND	0	0	0
7 REPLY	0	1	0
8 RELEASE	0	0	0
9 DECLINE	0	0	0
10 RECONFIGURE	0	0	0
11 INFO_REQUEST	0	0	0
12 RELAY_FORW	0	0	0
13 RELAY_REPLY	0	0	0
14 LEASEQUERY	0	0	0
15 LEASEQUERY_REPLY	0	0	0

```
-----
Dhcp6 Drop Reason Counters :
```

```
-----
 1 Dhcp6 oper state is not Up on src itf          0
 2 Dhcp6 oper state is not Up on dst itf          0
---snip---
38 Packet dropped by DHCP filter                  0
39 Packet dropped because authentication failed     0
=====
*A:BNG#
```

Use the following command to clear the DHCPv6 statistics.

```
A:BNG-1# clear router dhcp6 statistics
```

Troubleshooting

Following tools are available for troubleshooting PPPoE dual stack scenarios.

- system log (log-id 99)
- debugging aids
- protocol statistics

Log-id 99 is the default system log. Use appropriate filtering to reduce the output if needed.

```
*A:BNG# show log log-id 99
```

Following debug configuration is useful for troubleshooting PPPoE, RADIUS, DHCPv6 and ICMPv6.

```
debug
  service
    id 1
      ppp
        packet
          mode egr-ingr-and-dropped
          detail-level high
          discovery
          ppp
          dhcp-client
        exit
      exit
    exit
  exit
exit
```

```
debug
  router
    radius
      packet-type authentication accounting coa
      detail-level high
    exit
  exit
exit
```

```
debug
  router "Base"
```

```
ip
  dhcp6
    mode egr-ingr-and-dropped
    detail-level high
  exit
  icmp6
exit
exit
exit
```

Use the following commands for showing protocol related statistics.

```
show router dhcp6 statistics
show service id 1 pppoe session statistics
```

Advanced Topics

RADIUS COA

For dual stack PPPoE subscriber hosts, RADIUS-triggered mid-session change or/and session terminations identify the subscriber host to be changed by the same prefix that was originally returned from RADIUS or by the host-session-id (If RADIUS accounting host-accounting is enabled and the accounting session-id format equals number). Changing either the IPv4 or IPv6 information will result in both the v4 and v6 subscriber hosts being modified. Further elaboration on accounting is out of scope in this document.

IPv6CP Interface ID

IPv6CP negotiates, unlike ipv4-addresses in IPv4CP, only interface-ids (interface-id: the last 64 bits of an IPv6 address is the interface identifier that is unique to the 64-bit prefix of the IPv6 address and is usually derived from the link-layer or MAC address).

Dual stack PPPoE subscribers and the BNG exchange their interface-ids during the NCP phase. For ESM subscriber-interfaces on the BNG the interface-id is derived from the chassis-mac address.

- The BNG will nack the PPPoE host’s IPv6CP configuration request if the dual stack PPPoE host negotiates an interface-id equal zero or an interface-id equal to the BNG interface ID. In that scenario, the BNG offers in the IPv6CP nack message a suitable interface ID, see [Table 21: IPv6CP Nack Message Format](#).
- The BNG terminates the session if the dual stack PPPoE hosts nacks its IPv6CP configuration request and offers something else to the BNG.

Table 21: IPv6CP Nack Message Format

1	2	3	4	5	6	7	8
SAP ID				Last 2 bytes MAC host		Session ID	

Conclusion

This chapter provides configuration and troubleshooting commands for dual stack PPPoE subscribers on bridged or routed gateways. SLAAC is used as IPv6 address assignment for bridged gateway scenarios and stateful DHCPv6 prefix-delegation is used for address assignment for routed gateway scenarios. No RG WAN IPv6 address assignment is supported in this latter model.

DNSv6 addressing on a bridged gateway is retrieved by stateless DHCPv6 (information request and reply).

Flexible Authentication Model in ESM

This chapter provides information about Flexible Authentication Models in ESM.

Topics in this chapter include:

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

Applicability

This chapter is applicable to SR OS routers in the Routed Central Office (RCO) model and was initially written for SR OS Release 11.0.R2, but the CLI in the current edition is based on SR OS Release 15.0.R2.

Overview

The flexible authentication model for IPoE and PPPoE subscribers allows for mixing of configuration parameters obtained during the authentication phase from different sources: Local User Database (LUDB), RADIUS, or DHCP options that can be populated via a custom Python script. In case the same parameter is available from multiple sources, a priority mechanism is enforced whereby the parameter received from a higher priority source overrides the parameters received from the lower priority source in the following priority: LUDB to RADIUS to Python.

In this chapter we will configure a dual-stack IPoE and a dual stack PPPoE host using four different methods to obtain their configuration parameters. The setup will utilize a single BNG node with a locally configured DHCP server and LUDB as well as an external RADIUS server. Subscriber hosts are instantiated on managed (dynamic) SAPs.

The subscriber configuration parameters are in general divided into two categories:

- IP addressing parameters of the host — IPv4/v6 address/prefix, DNS servers, IPv4 default-gateway, IPv4 subnet-mask, IPv4/v6 address pool name, DHCPv4/v6 lease times, etc.
- Non IP addressing parameters of the host — Subscriber hosts strings are used to associate the subscriber-host with the desired level of service (sub/sla-profiles, inter-dest-id string, etc); managed routes are used for routing purposes to/from the host; etc.

The following four scenarios will be examined:

1. DHCP relay case (IP address is assigned via local DHCP server) with NO authentication. See [DHCP Relay Case with No Authentication](#).
2. DHCP relay case (IP address is assigned via local DHCP server) with LUDB + RADIUS authentication. See [DHCP Relay Case with LUDB + RADIUS Authentication](#).

RADIUS provides: sub/sla-profile strings and a framed IPv4 route.

LUDB provides: IP address pool, inter-dest-id string for Vport assignment, msap-defaults (routing context parameters and msap-policy).

3. IP proxy case (IP address is assigned via RADIUS) with LUDB + RADIUS authentication. [IP Proxy Case with LUDB + RADIUS Authentication](#)

RADIUS provides: IP addresses and related parameters (DNS server, IPv4 default-gateway, etc), inter-dest-id string for Vport assignment and a framed route.

IPv6 lease-times are provided under the group-interface.

LUDB provides: sub/sla-profile strings and msap-defaults (routing context parameters and msap-policy).

4. IP proxy case (IP address is assigned via LUDB) with LUDB + RADIUS authentication. [IP Proxy Case with LUDB + RADIUS Authentication](#)

RADIUS provides: sub/sla-profile strings and a framed IPv4 route.


LUDB provides IP addresses and related parameters (DNS server, IPv4 default-gateway, etc), inter-dest-id string for Vport assignment and msap-defaults (routing context parameters and msap-policy).

In cases 2-4, the domain-name *domain1* is appended to the IPoE and PPPoE username in LUDB, before RADIUS authentication takes place.

Configuration

The topology is shown in [Figure 75: Topology](#).

Figure 75: Topology



There is a common part of the configuration that applies uniformly across all four examined scenarios. This common part is outlined below and will not be repeated again when we describe more specific cases. It is assumed that the more specific cases also contain this common part of the configuration.

Common Configuration Part

Access Ethernet Port with QinQ Encapsulation

The following output displays a configuration example.

```
configure
  port 1/1/1
    ethernet
      mode access
      encap-type dot1q
    exit
  exit
exit
```

Capture SAP

A capture SAP is used to dynamically detect the VLAN id(s) in incoming DHCP/PPPoE packets (triggering packets) and conditionally instantiate the managed (dynamic) SAP. LUDB must be configured under the capture SAP to authorize the user accessing the capture SAP. The LUDB may contain additional parameters needed to set up the subscriber, it can point the subscriber to the RADIUS server for additional parameters or it may contain a default subscriber-host entry without any configuration parameters.

In this case, the **msap-defaults** under the capture SAP is used to select the routing context where the msap is created. **msap-defaults** can be also configured in the LUDB or be supplied via RADIUS.

PPPoE policy and *msap policy* are used to define PPPoE and SAP level parameters. Because the (dynamic) SAP does not exist at the time when the initial DHCP/PPPoE packets are received, the PPPoE/SAP level parameters are taken from the PPPoE/msap policy under the capture SAP. For example, those parameters are used in the PPP PADx/LCP/Authentication setup phase, they define default subscriber host strings, maximum number of subscriber hosts per SAP, the anti-spoofing mode, etc.

The X in the LUDB name (ludb-X) has to be replaced by a number depending on the scenario.

```
configure
  service
    vpls 2 customer 1 create
      stp
        shutdown
      exit
      sap 1/1/1:* capture-sap create
        description "open DHCP model testing"
        trigger-packet dhcp dhcp6 pppoe
        dhcp-user-db "ludb-X"
        dhcp6-user-db "ludb-X"
        pppoe-policy "pppoe-pol-1"
        pppoe-user-db "ludb-X"
        msap-defaults
          group-interface "grp-int-1"
          policy "msap-pol-1"
          service 1
        exit
      exit
    no shutdown
  exit
```

```
    exit
exit
```

auto-sub-id

The **auto-sub-id-key** command can be used in situations where the more specific **subscriber-id** string is not returned from LUDB or RADIUS. In this case, the auto subscriber-id for IPE hosts is set to the circuit-id while for PPPoE hosts the auto subscriber-id is set to the circuit-id plus session-id separated by the "|" delimiter which is inserted by default.

```
configure
  subscriber-mgmt
    auto-sub-id-key
      ipoe-sub-id-key circuit-id
      ppp-sub-id-key circuit-id session-id
    exit
  exit
exit
```

PPPoE Policy

There is a maximum of PPPoE sessions per MAC on a managed SAP. The default is 1 but is increased here to 10.

```
configure
  subscriber-mgmt
    ppp-policy "pppoe-pol-1" create
    ppp-mtu 1400
    max-sessions-per-mac 10
  exit
exit
exit
```

MSAP Policy

The MSAP policy defines the anti-spoofing mode which is in this particular example set to next-hop MAC (nh-mac). It also defines the default subscriber management parameters in case they are not supplied via LUDB or RADIUS.

```
configure
  subscriber-mgmt
    msap-policy "msap-pol-1" create
    sub-sla-mgmt
      def-sub-id use-auto-id
      def-sub-profile "sub-profile-1"
      def-sla-profile "sla-profile-1"
      sub-ident-policy "sub-ident-1"
      multi-sub-sap limit 500
    exit
    ies-vprn-only-sap-parameters
      anti-spoof nh-mac
    exit
  exit
exit
```



```
exit
```

subscriber-interface Configuration

The following output displays a subscriber interface configuration.

```
configure
  service
    vprn 1 customer 1 create
      subscriber-interface "sub-int-1" create
        # support for un-numbered IPv4 clients
        allow-unmatching-subnets
        # default gateway for IPv4 numbered clients
        address 10.12.0.1/24
        ipv6
          # fixed delegated prefix length for IA-PD
          delegated-prefix-len 56
          # support for un-numbered IPv6 clients
          allow-unmatching-prefixes
        exit
      group-interface "grp-int-1" create
        ipv6
          router-advertisements
            # hint to the client to use DHCPv6
            managed-configuration
            # enabling router-advertisements
            no shutdown
          exit
        dhcp6
          # must be the same as under the capture-SAP
          user-db "ludb-1"
        exit
      exit
    # ARP table is populated based on the lease state table
    arp-populate
    dhcp
      server 10.10.1.1
      # accept DHCP packets on this group interface
      trusted
      # max number of DHCPv4 clients on each
      # SAP of this group-interface
      lease-populate 100
      # must be the same as under the capture-SAP
      user-db "ludb-1"
      no shutdown
    exit
  pppoe
    policy "pppoe-pol-1"
    session-limit 1000
    sap-session-limit 1000
    # must be the same as under the capture-SAP
    user-db "ludb-1"
    no shutdown
  exit
exit
exit
no shutdown
exit
exit
exit
exit
```

For numbered/unnumbered subscriber-hosts also take a look at the DHCP/PPPoE clients whose assigned IP address is outside of any IP subnet/prefix configured under the subscriber-interface.

Specific Configuration Parts

DHCP Relay Case with No Authentication

The IP address is assigned via local DHCP server. The LUDB is accessed even in the scenario without authentication. There must be a default host LUDB entry present that will match on any value specified in the match-list criteria. The LUDB is accessed from the capture SAP (part of the common configuration).

```
configure
  subscriber-mgmt
    local-user-db "ludb-1" create
      description "for CASE-1"
      ipoe
        # host matching is based on circuit-ID in DHCP packets
        match-list circuit-id
        host "default" create
          no shutdown
        exit
      exit
    exit
  ppp
    # host matching is base on PPPoE username
    match-list username
    host "default" create
      # explicitly enable IPCPv6
      force-ipv6cp
      no shutdown
    exit
  exit
  no shutdown
exit
exit
exit
```

Once the routing context (service id and group-interface) is determined as defined under the capture SAP defaults (part of the common configuration), the DHCP/PPPoE requests are served according to the group-interface configuration. The IP address request is relayed to the DHCPv4/v6 server. Since the LUDB does not provide a pool name, the **gi-address** and the **link-address** is used by the DHCP relay/server to select the pool from which the IP address will be assigned.

```
configure
  service
    vprn 1 customer 1 create
      subscriber-interface "sub-int-1" create
        group-interface "grp-int-1" create
          ipv6
            dhcp6
              # DHCPv6 relay configuration
              relay
                link-address 2001:DB8:30::
                # DHCPv6 server IPv6 address
                server 2001:DB8::1
                client-applications dhcp ppp
                no shutdown
              exit
            exit
          exit
        exit
      exit
    exit
```

```

        exit
    exit
    arp-populate
    dhcp
        # DHCPv4 server IP address
        server 10.10.1.1
        client-applications dhcp ppp
        gi-address 10.12.0.1
        no shutdown
    exit
exit
exit
no shutdown
exit
exit
exit
exit

```

DHCPv4/v6 servers are locally configured in the node and attached to a loopback interface.

```

configure
service
    vprn 1 customer 1 create
        interface "int-DHCP" create
            # IPv4 address which the DHCPv4 is listening on
            address 10.10.1.1/24
            ipv6
                # IPv6 address which the DHCPv6 server is listening on
                address 2001:DB8::1/128
                local-dhcp-server "svc-1-dhcp6"
            exit
            # attaching the DHCPv4 server to the loopback interface
            local-dhcp-server "svc-1-dhcp4"
            loopback
        exit
    exit
exit
exit
exit

```

In the local DHCP servers two pools are defined:

- LUDb — To be used for IP address assignment when LUDb returns the pool name.
- Gi-addr — To be used when gi-address/link-address are used to select the pool for IP address assignment.

Lease times for IPv4 and IPv6 are configured in the local DHCP server which is used only in the relay case (when the IP address is supplied via DHCP server and not through RADIUS or the LUDb).

```

configure
service
    vprn 1 customer 1 create
        dhcp
            local-dhcp-server "svc-1-dhcp4" create
                # the gi-address can be used to select the pool
                use-gi-address
                # the pool name can be explicitly provided
                use-pool-from-client
                # the pool used when the LUDb provides the pool name
                pool "ludb" create
                    options
                        dns-server 172.16.16.16 172.16.16.17
                        # DHCPv4 lease time
                        lease-time hrs 1
            exit
        exit
    exit

```

```
        exit
        subnet 10.10.0.0/24 create
        options
            subnet-mask 255.255.255.0
            default-router 10.10.0.1
        exit
        address-range 10.10.0.100 10.10.0.200
    exit
exit
# pool selected based on the gi-address
pool "gi-addr" create
options
    dns-server 172.16.16.16 172.16.16.17
    # DHCPv4 lease time
    lease-time hrs 1
exit
subnet 10.12.0.0/24 create
options
    subnet-mask 255.255.255.0
    default-router 10.12.0.1
exit
address-range 10.12.0.100 10.12.0.200
exit
exit
no shutdown
exit
dhcp6
    local-dhcp-server "svc-1-dhcp6" create
    use-link-address
    use-pool-from-client
    pool "ludb" create
        prefix 2001:DB8:10::/48 pd wan-host create
            preferred-lifetime min 30
            rebind-timer min 20
            renew-timer min 15
            # DHCPv6 lease time
            valid-lifetime hrs 1
            options
                dns-server 2001:DB8::1000 2001:DB8::1001
            exit
        exit
    exit
    pool "gi-addr" create
        prefix 2001:DB8:30::/48 pd wan-host create
            preferred-lifetime min 30
            rebind-timer min 20
            renew-timer min 15
            # DHCPv6 lease time
            valid-lifetime hrs 1
            options
                dns-server 2001:DB8::1000 2001:DB8::1001
            exit
        exit
    exit
    no shutdown
exit
exit
exit
exit
exit
exit
```

Default sub/sla-profiles, from the msap-policy, are used (part of the common configuration).

```
configure
  subscriber-mgmt
    sla-profile "sla-profile-1" create
      description "default SLA profile"
      host-limits
        overall 3
      exit
    exit
  sub-profile "sub-profile-1" create
    description "default SUB profile"
    egress
      agg-rate-limit 1000
    exit
  exit
exit
exit
```

Show Commands

The following command shows that the default sub/sla-profiles are in use, that the IP addresses are selected from the gi-addr pool in local DHCP server and that the subscriber-id is set to circuit-id for the IPoE subscriber-host and to **username|session-id** combination for the PPPoE subscriber-host.

```
*A:BNG-1# show service active-subscribers

=====
Active Subscribers
=====
Subscriber open-dhcp-1 (sub-profile-1)
-----
(1) SLA Profile Instance sap:[1/1/1:11] - sla:sla-profile-1
-----
IP Address
-----
MAC Address      Session      Origin      Svc      Fwd
-----
10.12.0.107      00:0c:29:00:00:11  N/A         DHCP      1         Y
2001:db8:30:103::1/128  00:0c:29:00:00:11  N/A         DHCP6     1         Y
2001:db8:30:400::/56    00:0c:29:00:00:11  N/A         DHCP6     1         Y
-----

Subscriber open-pppoe-1|1 (sub-profile-1)
-----
(1) SLA Profile Instance sap:[1/1/1:21] - sla:sla-profile-1
-----
IP Address
-----
MAC Address      Session      Origin      Svc      Fwd
-----
10.12.0.108      00:0c:29:00:00:21  PPP 1       IPCP      1         Y
2001:db8:30:104::1/128  00:0c:29:00:00:21  PPP 1       DHCP6     1         Y
```

```
2001:db8:30:500::/56
00:0c:29:00:00:21    PPP 1          DHCP6-PD    1          Y
-----
-----
Number of active subscribers : 2
=====
*A:BNG-1#
```

The following command shows more details about the subscriber-host, such as the group-interface, address origin, acct-session-id, etc. Even though there are only two dual-stack hosts (one IPoE and one PPPoE), each of them has three IP addresses that show up as different hosts.

For the purpose of brevity, the output for only two IP hosts are shown, one with an IPv4 address and one with an IPv6 address. The remaining IP addresses/prefixes are not shown because the output follows the same logic.

```
*A:BNG-1# show service id 1 subscriber-hosts detail

=====
Subscriber Host table
=====
Sap          Subscriber
  IP Address
  MAC Address      PPPoE-SID Origin      Fwding State
-----
[1/1/1:11]      open-dhcp-1
  10.12.0.107
  00:0c:29:00:00:11    N/A      DHCP      Fwding
-----
Subscriber-interface : sub-int-1
Group-interface      : grp-int-1
Sub Profile          : sub-profile-1
SLA Profile          : sla-profile-1
App Profile          : N/A
Egress Q-Group       : N/A
Egress Vport         : N/A
Acct-Session-Id      : 14F2FF00000026591168E2
Acct-Q-Inst-Session-Id: 14F2FF00000027591168E2
Address Origin        : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-----
[1/1/1:11]      open-dhcp-1
  2001:db8:30:103::1/128
  00:0c:29:00:00:11    N/A      IPoE-DHCP6  Fwding
-----

---snip---

-----
[1/1/1:11]      open-dhcp-1
  2001:db8:30:400::/56
  00:0c:29:00:00:11    N/A      IPoE-DHCP6  Fwding
-----

---snip---
-----
```

```
[1/1/1:21]          open-pppoe-1|1
10.12.0.108
00:0c:29:00:00:21   1          IPCP          Fwding
-----
---snip---
-----
[1/1/1:21]          open-pppoe-1|1
2001:db8:30:104::1/128
00:0c:29:00:00:21   1          PPP-DHCP6     Fwding
-----
---snip---
-----
[1/1/1:21]          open-pppoe-1|1
2001:db8:30:500::/56
00:0c:29:00:00:21   1          PPP-DHCP6     Fwding
-----
Subscriber-interface : sub-int-1
Group-interface      : grp-int-1
Sub Profile          : sub-profile-1
SLA Profile          : sla-profile-1
App Profile          : N/A
Egress Q-Group       : N/A
Egress Vport         : N/A
Acct-Session-Id      : 14F2FF0000002E5911690E
Acct-Q-Inst-Session-Id: 14F2FF0000002C5911690E
Address Origin       : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status   : N/A
OT HTTP Rdr Fltr Src : N/A
HTTP Rdr URL Override : N/A
GTP local break-out  : No
DIAMETER session ID Gx: N/A
-----
Number of subscriber hosts : 6
=====
*A:BNG-1#
```

The following command shows that there are no sub/sla-profile strings assigned to the subscriber. Instead the default sub/sla-profiles from the msap-policy are used.

The IP address is assigned by the DHCP server which also supplied the def-gw information, DNS servers, the net-mask and the lease time.

The circuit-id and the subscriber-id are set to the same value.

```
*A:BNG-1# show service id 1 dhcp lease-state detail

=====
DHCP lease states for service 1
=====
Service ID          : 1
IP Address          : 10.12.0.107
Client HW Address   : 00:0c:29:00:00:11
Subscriber-interface : sub-int-1
Group-interface     : grp-int-1
SAP                 : [1/1/1:11]
Termination Type    : local
Up Time             : 0d 00:02:11
Remaining Lease Time : 0d 00:57:49
```

```

Remaining SessionTime: N/A
Persistence Key       : N/A

Sub-Ident            : "open-dhcp-1"
Sub-Profile-String   : ""
SLA-Profile-String   : ""
App-Profile-String   : ""
Lease ANCP-String    : ""
Lease Int Dest Id    : ""
Category-Map-Name    : ""

Lease Info origin    : DHCP

Ip-Netmask           : 255.255.255.0
Broadcast-Ip-Addr    : N/A
Default-Router       : 10.12.0.1
Primary-Dns          : 172.16.16.16
Secondary-Dns        : 172.16.16.17
Primary-Nbns         : N/A
Secondary-Nbns       : N/A

ServerLeaseStart     : 05/09/2017 08:59:46
ServerLastRenew      : 05/09/2017 08:59:46
ServerLeaseEnd       : 05/09/2017 09:59:46
Session-Timeout      : N/A
IPoE|PPP session     : No
Lease-Time           : 0d 01:00:00
DHCP Server Addr     : 10.10.1.1

```

Relay Agent Information

```

Circuit Id          : open-dhcp-1
Radius User-Name    : ""

```

```

-----
Number of lease states : 1
=====

```

```

*A:BNG-1#

```

Then there is a similar command used for DHCPv6 lease-state details.

For the purpose of brevity, the output for only two IPv6 leases is shown. The remaining two IPv6 leases are not shown because the output follows the same logic.

```

*A:BNG-1# show service id 1 dhcp6 lease-state detail

```

```

=====
DHCP lease states for service 1
=====

```

```

Service ID          : 1
IP Address          : 2001:db8:30:103::1/128
Client HW Address   : 00:0c:29:00:00:11
Subscriber-interface : sub-int-1
Group-interface     : grp-int-1
SAP                 : [1/1/1:11]
Termination Type    : local
Up Time             : 0d 00:02:34
Remaining Lease Time : 0d 00:57:26
Remaining SessionTime: N/A
Persistence Key      : N/A

```

```

Sub-Ident          : "open-dhcp-1"
Sub-Profile-String : ""
SLA-Profile-String : ""
App-Profile-String : ""

```



```

Lease ANCP-String      : ""
Lease Int Dest Id     : ""
Category-Map-Name     : ""
Dhcp6 ClientId (DUID) : 0001000120a31b12000c29000011
Dhcp6 IAID            : 2
Dhcp6 IAID Type       : non-temporary
Dhcp6 Client Ip       : fe80::20c:29ff:fe00:11
Primary-Dns           : N/A
Secondary-Dns         : N/A
Pool Name             : ""
Dhcp6 Server Addr     : 2001:db8::1
Dhcp6 ServerId (DUID) : 0003000114f2ff000000
Dhcp6 InterfaceId     : open-dhcp-1
Dhcp6 RemoteId        : N/A
Radius sub-if prefix  : N/A
Router adv. policy    : N/A

```

```

Lease Info origin      : DHCP

```

```

ServerLeaseStart       : 05/09/2017 09:00:00
ServerLastRenew       : 05/09/2017 09:00:00
ServerLeaseEnd        : 05/09/2017 10:00:00
Session-Timeout       : N/A
IPoE|PPP session      : No
Radius User-Name      : ""

```

```

-----
Service ID             : 1
IP Address             : 2001:db8:30:104::1/128
Client HW Address      : 00:0c:29:00:00:21

```

```

---snip---

```

```

-----
Service ID             : 1
IP Address             : 2001:db8:30:400::/56
Client HW Address      : 00:0c:29:00:00:11

```

```

---snip---

```

```

-----
Service ID             : 1
IP Address             : 2001:db8:30:500::/56
Client HW Address      : 00:0c:29:00:00:21
Subscriber-interface   : sub-int-1
Group-interface        : grp-int-1
SAP                    : [1/1/1:21]
Termination Type       : local
Up Time                : 0d 00:02:06
Remaining Lease Time   : 0d 00:57:54
Remaining SessionTime  : N/A
Persistence Key        : N/A

```

```

Sub-Ident              : "open-pppoe-1|1"
Sub-Profile-String     : ""
SLA-Profile-String     : ""
App-Profile-String     : ""
Lease ANCP-String      : ""
Lease Int Dest Id     : ""
Category-Map-Name     : ""
Dhcp6 ClientId (DUID) : 0001000120a33d28000c29000021
Dhcp6 IAID            : 1
Dhcp6 IAID Type       : prefix
Dhcp6 Client Ip       : fe80::20c:29ff:fe00:21
Primary-Dns           : N/A

```

```

Secondary-Dns      : N/A
Pool Name          : ""
Dhcp6 Server Addr  : 2001:db8::1
Dhcp6 ServerId (DUID): 0003000114f2ff000000
Dhcp6 InterfaceId  : open-pppoe-1
Dhcp6 RemoteId     : N/A
Radius sub-if prefix : N/A
Router adv. policy  : N/A

```

```

Lease Info origin   : DHCP

```

```

ServerLeaseStart    : 05/09/2017 09:00:30
ServerLastRenew     : 05/09/2017 09:00:30
ServerLeaseEnd      : 05/09/2017 10:00:30
Session-Timeout     : N/A
IPoE|PPP session    : PPP
Radius User-Name     : "open-pppoe-1"

```

```

-----
Number of lease states : 4
=====

```

```

*A:BNG-1#

```

DHCP Relay Case with LUDB + RADIUS Authentication

IP address is assigned via local DHCP server.

- RADIUS provides sub/sla-profile strings and a framed IPv4 route.
- LUDB provides IP address pool, inter-dest-id string for Vport assignment, msap-defaults (routing context parameters and msap-policy).

Vport aggregate rate limit and the port scheduler are now added to the physical port. The Vport is associated with the subscriber through the inter-dest-id string obtained via LUDB.

```

configure
  port 1/1/1
    ethernet
      mode access
      encap-type dot1q
      egress-scheduler-policy "port"
      access
        egress
          vport "open-dhcp" create
            agg-rate
              rate 500
            exit
            host-match dest "open-auth-vport" create
          exit
        exit
      exit
    exit
  no shutdown
exit

```

The LUDB is used to assign the IP pool name (pool-name = ludb) and the inter-dest-id string (inter-dest-id = open-auth-vport) to the subscriber. The pool name is carried to the DHCP server via custom DHCP options [(82,9,13) in DHCPv4 and (17,1->wan_pool and 2->pfx_pool) in DHCPv6].

The domain name *domain1* is appended to the username (circuit-id = open-dhcp-2 or username = open-pppoe-2) before an Access-Request message is sent to the RADIUS server which is configured in the authentication policy *auth-pol-1*.

The inter-dest-id string taken from the LUDB is passed to the subscriber management module in the node via DHCP option 254 in DHCP ACK/Reply.

```
configure
  subscriber-mgmt
    local-user-db "ludb-2" create
      description "for CASE-2"
      ipoe
        match-list circuit-id
        host "open-dhcp-2" create
          host-identification
            circuit-id string "open-dhcp-2"
          exit
          address pool "ludb"
          auth-policy "auth-pol-1"
          auth-domain-name "domain1"
          identification-strings 254 create
            inter-dest-id "open-auth-vport"
          exit
        msap-defaults
          group-interface "grp-int-1"
          policy "msap-pol-1"
          service 1
        exit
        ipv6-wan-address-pool "ludb"
        ipv6-delegated-prefix-pool "ludb"
        no shutdown
      exit
    exit
  ppp
    match-list circuit-id mac username
    host "open-ppp-2" create
      host-identification
        username "open-pppoe-2"
      exit
      auth-policy "auth-pol-1"
      address pool "ludb"
      password chap "letmein"
      identification-strings 254 create
        inter-dest-id "open-auth-vport"
      exit
      msap-defaults
        group-interface "grp-int-1"
        policy "msap-pol-1"
        service 1
      exit
      ipv6-delegated-prefix-pool "ludb"
      ipv6-wan-address-pool "ludb"
      no shutdown
    exit
  exit
  no shutdown
exit
exit
exit
```

The **inter-dest-id** string taken from the LUDB is passed to the subscriber management module in the node via DHCPv4/v6 option 254 that is specified in the subscriber identification policy.

```
configure
  subscriber-mgmt
    sub-ident-policy "sub-ident-1" create
      strings-from-option 254
    exit
  exit
exit
```

The RADIUS server is defined via the authentication policy. The domain name can be appended to the PPPoE subscriber host directly via the authentication-policy while for IPoE subscribers, the domain name is appended via the authentication-policy in conjunction with the LUDB. This can be verified in the output (shown later) of the **show service id 1 dhcp lease-state detail** and **show service id 1 dhcp6 lease-state detail** commands (on the "radius user-name" line).

```
configure
  subscriber-mgmt
    authentication-policy "auth-pol-1" create
      description "RADIUS authentication policy"
      password "letmein"
      ppp-user-name append "domain1"
      user-name-format circuit-id append
      accept-authorization-change
      pppoe-access-method pap-chap
      radius-server-policy "rad-serv-pol-1"
    exit
  exit
exit
```

The RADIUS user configuration file uses the domain-name extension, as inserted by the BNG, to authenticate the user:

```
open-dhcp-2@domain1 Cleartext-Password := "letmein"
  Alc-Subsc-Prof-Str = rad-sub,
  Alc-SLA-Prof-Str = rad-sla,
  Framed-Route = "192.168.1.0/24 0.0.0.0",

open-pppoe-2@domain1 Cleartext-Password := "letmein"
  Alc-Subsc-Prof-Str = rad-sub,
  Alc-SLA-Prof-Str = rad-sla,
  Framed-Route = "192.168.2.0/24 0.0.0.0",
```

DHCPv4/v6 servers are locally configured in the SR OS and attached to a loopback interface:

```
configure
  service
    vprn 1 customer 1 create
      interface "int-DHCP" create
        # IPv4 address which the DHCPv4 server is listening on
        address 10.10.1.1/24
        ipv6
          # IPv6 address which the DHCPv6 server is listening on
          address 2001:DB8::1/128
          # attach the DHCPv6 server to this loopback interface
          local-dhcp-server "svc-1-dhcp6"
        exit
        # attach the DHCPv4 server to this loopback interface
        local-dhcp-server "svc-1-dhcp4"
```

```
        loopback
      exit
    exit
  exit
exit
```

Group-interface configuration. Note that common parts of the configuration as defined earlier, still apply:

```
configure
  service
    vprn 1 customer 1 create
      subscriber-interface "sub-int-1" create
        group-interface "grp-int-1" create
          ipv6
            dhcp6
              user-db "ludb-2"
              # DHCPv6 relay configuration
              relay
                link-address 2001:DB8:30::
                server 2001:DB8::1
                client-applications dhcp ppp
                no shutdown
              exit
            exit
          exit
        arp-populate
        # DHCPv6 relay configuration
        dhcp
          proxy-server
            emulated-server 10.12.0.1
            no shutdown
          exit
          option
            # SR OS will not insert its own circuit-ID
            no circuit-id
            # SR OS will not insert its own remote-ID
            no remote-id
            vendor-specific-option
              pool-name
            exit
          exit
          server 10.10.1.1
          client-applications dhcp ppp
          user-db "ludb-2"
          no shutdown
        exit
      pppoe
        policy "pppoe-pol-1"
        session-limit 1000
        sap-session-limit 1000
        user-db "ludb-2"
        no shutdown
      exit
    exit
  exit
  no shutdown
exit
exit
exit
exit
```

Lease times for IPv4 and IPv6 are configured in the local DHCP server. Lease times under the local DHCP server are used only in the relay case (when IP address is supplied via DHCP server and **not** RADIUS or LUDB). In the proxy case, the lease times can be obtained via LUDB, RADIUS, or group-interface.

```
configure
  service
    vprn 1 customer 1 create
      dhcp
        local-dhcp-server "svc-1-dhcp4" create
          # gi-address can be used to select the pool
          use-gi-address
          # pool name can be explicitly provided
          use-pool-from-client
          # pool used when LUDB provides the pool name
          pool "ludb" create
            options
              dns-server 172.16.16.16 172.16.16.17
              lease-time hrs 1
            exit
          subnet 10.10.0.0/24 create
            options
              subnet-mask 255.255.255.0
              default-router 10.10.0.1
            exit
          address-range 10.10.0.100 10.10.0.200
          exit
        exit
        # pool selected based on the gi-address
        pool "gi-addr" create
          options
            dns-server 172.16.16.16 172.16.16.17
            lease-time hrs 1
          exit
        subnet 10.12.0.0/24 create
          options
            subnet-mask 255.255.255.0
            default-router 10.12.0.1
          exit
        address-range 10.12.0.100 10.12.0.200
        exit
      exit
    no shutdown
  exit
exit
dhcp6
  local-dhcp-server "svc-1-dhcp6" create
    use-link-address
    use-pool-from-client
    pool "ludb" create
      prefix 2001:DB8:10::/48 pd wan-host create
        preferred-lifetime min 30
        rebind-timer min 20
        renew-timer min 15
        valid-lifetime hrs 1
      options
        dns-server 2001:DB8::1000 2001:DB8::1001
      exit
    exit
  exit
  pool "gi-addr" create
    prefix 2001:DB8:30::/48 pd wan-host create
      preferred-lifetime min 30
      rebind-timer min 20
```

```

renew-timer min 15
valid-lifetime hrs 1
options
    dns-server 2001:DB8::1000 2001:DB8::1001
exit
exit
    exit
        no shutdown
            exit
                exit
                    exit
                        exit
                            exit

```

RADIUS sub/sla-profiles supplied via RADIUS are used:

```

configure
    subscriber-mgmt
        sla-profile "rad-sla" create
            description "sla-profile obtained from RADIUS"
            host-limits
                overall 100
            exit
            egress
                qos 1 vport-scheduler
            exit
            ip-filter 1
        exit
    exit
    sub-profile "rad-sub" create
        description "sub-profile obtained from RADIUS"
        egress
            agg-rate-limit 15000
        exit
    exit
exit
exit
exit

```

Show Commands

The following command shows that the rad-sub/sla-profiles, as supplied via RADIUS, are in use.

The IP addresses are selected from the pool-name LUDB in the local DHCP server. The subscriber-id is **circuit-id** for IPoE subscriber-host and the **username|session-id** combination for PPPoE subscriber host.

```

*A:BNG-1# show service active-subscribers

=====
Active Subscribers
=====
-----
Subscriber  open-dhcp-2 (rad-sub)
-----
(1) SLA Profile Instance sap:[1/1/1:12] - sla:rad-sla
-----
IP Address
-----
MAC Address      Session      Origin      Svc      Fwd
-----
10.10.0.102

```

```

00:0c:29:00:00:12 N/A DHCP 1 Y
2001:db8:10:101::1/128
00:0c:29:00:00:12 N/A DHCP6 1 Y
2001:db8:10:200::/56
00:0c:29:00:00:12 N/A DHCP6 1 Y
-----
Subscriber open-pppoe-2|1 (rad-sub)
-----
(1) SLA Profile Instance sap:[1/1/1:22] - sla:rad-sla
-----
IP Address
MAC Address Session Origin Svc Fwd
-----
10.10.0.103
00:0c:29:00:00:22 PPP 1 IPCP 1 Y
2001:db8:10:102::1/128
00:0c:29:00:00:22 PPP 1 DHCP6 1 Y
2001:db8:10:300::/56
00:0c:29:00:00:22 PPP 1 DHCP6-PD 1 Y
-----
Number of active subscribers : 2
=====
*A:BNG-1#

```

The following command shows more details about the subscriber-host, such as the group-interface, vport, address origin, acct-session-id, etc. Vport is selected based on the **inter-dest-id** string supplied via the LUDB.

For the purpose of brevity, the output for only two IP addresses *hosts* is shown, one with an IPv4 address and one with an IPv6 address. The remaining IP addresses/prefixes are not shown because the output follows the same logic.

```

*A:BNG-1# show service id 1 subscriber-hosts detail
=====
Subscriber Host table
=====
Sap          Subscriber
IP Address
MAC Address  PPPoE-SID Origin  Fwding State
-----
[1/1/1:12]   open-dhcp-2
10.10.0.102
00:0c:29:00:00:12 N/A DHCP Fwding
-----
Subscriber-interface : sub-int-1
Group-interface     : grp-int-1
Sub Profile         : rad-sub
SLA Profile         : rad-sla
App Profile         : N/A
Egress Q-Group      : N/A
Egress Vport        : open-dhcp
Acct-Session-Id     : 14F2FF0000002F59116C4A
Acct-Q-Inst-Session-Id: 14F2FF0000003059116C4A
Address Origin      : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status   : N/A
OT HTTP Rdr Fltr Src : N/A
HTTP Rdr URL Override : N/A

```



```
GTP local break-out : No
DIAMETER session ID Gx: N/A
-----
[1/1/1:12]          open-dhcp-2
2001:db8:10:101::1/128
00:0c:29:00:00:12   N/A      IPoE-DHCP6   Fwding
-----

---snip---

-----
[1/1/1:12]          open-dhcp-2
2001:db8:10:200::/56
00:0c:29:00:00:12   N/A      IPoE-DHCP6   Fwding
-----

---snip---

-----
[1/1/1:22]          open-pppoe-2|1
10.10.0.103
00:0c:29:00:00:22   1       IPCP       Fwding
-----

---snip---

-----
[1/1/1:22]          open-pppoe-2|1
2001:db8:10:102::1/128
00:0c:29:00:00:22   1       PPP-DHCP6   Fwding
-----

---snip---

-----
[1/1/1:22]          open-pppoe-2|1
2001:db8:10:300::/56
00:0c:29:00:00:22   1       PPP-DHCP6   Fwding
-----

Subscriber-interface : sub-int-1
Group-interface      : grp-int-1
Sub Profile          : rad-sub
SLA Profile          : rad-sla
App Profile          : N/A
Egress Q-Group       : N/A
Egress Vport         : open-dhcp
Acct-Session-Id      : 14F2FF0000003759116C68
Acct-Q-Inst-Session-Id: 14F2FF0000003559116C68
Address Origin       : Dynamic
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status   : N/A
OT HTTP Rdr Fltr Src : N/A
HTTP Rdr URL Override : N/A
GTP local break-out  : No
DIAMETER session ID Gx: N/A
-----
Number of subscriber hosts : 6
=====
*A:BNG-1#
```

The following command shows that the subscriber identity is set to **circuit-id** (plus session-id) as instructed by **auto-sub-id-key** command (subscriber-id string is not returned via the LUDB or RADIUS). The lease

times are set to 1h as defined in the DHCP server. The username passed to RADIUS is a **circuit-id** or a **username** appended with the *domain1* domain name.

```
*A:BNG-1# show service id 1 dhcp lease-state detail

=====
DHCP lease states for service 1
=====
Service ID           : 1
IP Address           : 10.10.0.102
Client HW Address    : 00:0c:29:00:00:12
Subscriber-interface : sub-int-1
Group-interface      : grp-int-1
SAP                  : [1/1/1:12]
Termination Type     : local
Up Time              : 0d 00:01:36
Remaining Lease Time : 0d 00:58:25
Remaining SessionTime: N/A
Persistence Key       : N/A

Sub-Ident            : "open-dhcp-2"
Sub-Profile-String    : "rad-sub"
SLA-Profile-String    : "rad-sla"
App-Profile-String    : ""
Lease ANCP-String     : ""
Lease Int Dest Id     : "open-auth-vport"
Category-Map-Name     : ""

Lease Info origin     : DHCP

Ip-Netmask            : 255.255.255.0
Broadcast-Ip-Addr     : N/A
Default-Router        : 10.10.0.1
Primary-Dns           : 172.16.16.16
Secondary-Dns         : 172.16.16.17
Primary-Nbns          : N/A
Secondary-Nbns        : N/A

ServerLeaseStart      : 05/09/2017 09:14:18
ServerLastRenew       : 05/09/2017 09:14:18
ServerLeaseEnd        : 05/09/2017 10:14:18
Session-Timeout       : N/A
IPoE|PPP session      : No
Lease-Time            : 0d 01:00:00
DHCP Server Addr      : 10.10.1.1

Relay Agent Information
  Circuit Id          : open-dhcp-2
  Radius User-Name    : "open-dhcp-2@domain1"

-----
Managed Routes
-----
IP Address              Status      Metric Tag      Pref
-----
192.168.1.0/24          installed    0      none      0
-----

Number of lease states : 1
=====
*A:BNG-1#
```

For the purpose of brevity the output for only two IPv6 leases is shown. The remaining two IPv6 leases are not shown because the output follows the same logic.

```
*A:BNG-1# show service id 1 dhcp6 lease-state detail

=====
DHCP lease states for service 1
=====
Service ID      : 1
IP Address      : 2001:db8:10:101::1/128
Client HW Address : 00:0c:29:00:00:12
Subscriber-interface : sub-int-1
Group-interface  : grp-int-1
SAP              : [1/1/1:12]
Termination Type : local
Up Time          : 0d 00:01:59
Remaining Lease Time : 0d 00:58:02
Remaining SessionTime: N/A
Persistence Key   : N/A

Sub-Ident       : "open-dhcp-2"
Sub-Profile-String : "rad-sub"
SLA-Profile-String : "rad-sla"
App-Profile-String : ""
Lease ANCP-String : ""
Lease Int Dest Id : "open-auth-vport"
Category-Map-Name : ""
Dhcp6 ClientId (DUID): 0001000120a33b22000c29000012
Dhcp6 IAID      : 2
Dhcp6 IAID Type  : non-temporary
Dhcp6 Client Ip   : fe80::20c:29ff:fe00:12
Primary-Dns       : N/A
Secondary-Dns     : N/A
Pool Name         : "ludb"
Dhcp6 Server Addr : 2001:db8::1
Dhcp6 ServerId (DUID): 0003000114f2ff000000
Dhcp6 InterfaceId : open-dhcp-2
Dhcp6 RemoteId    : N/A
Radius sub-if prefix : N/A
Router adv. policy : N/A

Lease Info origin : DHCP

ServerLeaseStart   : 05/09/2017 09:14:27
ServerLastRenew    : 05/09/2017 09:14:27
ServerLeaseEnd     : 05/09/2017 10:14:27
Session-Timeout    : N/A
IPoE|PPP session   : No
Radius User-Name    : "open-dhcp-2@domain1"
-----
Service ID      : 1
IP Address      : 2001:db8:10:102::1/128
Client HW Address : 00:0c:29:00:00:22

---snip---

-----
Service ID      : 1
IP Address      : 2001:db8:10:200::/56
Client HW Address : 00:0c:29:00:00:12

---snip---
```

```
-----
Service ID      : 1
IP Address      : 2001:db8:10:300::/56
Client HW Address : 00:0c:29:00:00:22
Subscriber-interface : sub-int-1
Group-interface  : grp-int-1
SAP             : [1/1/1:22]
Termination Type : local
Up Time         : 0d 00:01:39
Remaining Lease Time : 0d 00:58:21
Remaining SessionTime: N/A
Persistence Key  : N/A
Sub-Ident       : "open-pppoe-2|1"
Sub-Profile-String : "rad-sub"
SLA-Profile-String : "rad-sla"
App-Profile-String : ""
Lease ANCP-String : ""
Lease Int Dest Id : "open-auth-vport"
Category-Map-Name : ""
Dhcp6 ClientId (DUID): 0001000120a3175e000c29000022
Dhcp6 IAID      : 1
Dhcp6 IAID Type  : prefix
Dhcp6 Client Ip   : fe80::20c:29ff:fe00:22
Primary-Dns       : N/A
Secondary-Dns     : N/A
Pool Name         : "ludb"
Dhcp6 Server Addr : 2001:db8::1
Dhcp6 ServerId (DUID): 0003000114f2ff000000
Dhcp6 InterfaceId : open-pppoe-2
Dhcp6 RemoteId    : N/A
Radius sub-if prefix : N/A
Router adv. policy : N/A

Lease Info origin : DHCP

ServerLeaseStart   : 05/09/2017 09:14:48
ServerLastRenew    : 05/09/2017 09:14:48
ServerLeaseEnd     : 05/09/2017 10:14:48
Session-Timeout    : N/A
IPoE|PPP session   : PPP
Radius User-Name    : "open-pppoe-2@domain1"
-----
Number of lease states : 4
=====
*A: BNG-1#
```

IP Proxy Case with LUDB + RADIUS Authentication

IP address is assigned via RADIUS.

- RADIUS provides IP addresses (IPv6 lease-times are provided under the group-interface) and related parameters (DNS server, IPv4 default-gateway, etc), inter-dest-id string for Vport assignment and a framed route.
- LUDB provides sub/sla-profile strings and msap-defaults (routing context parameters and msap-policy).

Vport aggregate rate limit and the port scheduler are now added to the physical port. The Vport is associated with the subscriber through the inter-dest-id string obtained via the LUDB.

```
configure
port 1/1/1
```

```

ethernet
  mode access
  encaps-type dot1q
  egress-scheduler-policy "port"
  access
    egress
      vport "open-dhcp" create
      agg-rate
        rate 500
      exit
      host-match dest "open-auth-vport" create
    exit
  exit
exit
exit
no shutdown
exit
exit

```

The LUDB is used to assign the sub/sla-profile strings.

The domain name *domain1* is appended to the username (circuit-id = open-dhcp-3 or username = open-pppoe-3) before an Access-Request is sent to the RADIUS server that is configured in the authentication policy *auth-pol-1*.

```

configure
  subscriber-mgmt
    local-user-db "ludb-3" create
    description "for CASE-3"
    ipoe
      match-list circuit-id
      host "open-dhcp-3" create
      host-identification
        circuit-id string "open-dhcp-3"
      exit
      auth-policy "auth-pol-1"
      auth-domain-name "domain1"
      identification-strings 254 create
        sla-profile-string "ludb-sla"
        sub-profile-string "ludb-sub"
      exit
      msap-defaults
        group-interface "grp-int-1"
        policy "msap-pol-1"
        service 1
      exit
      no shutdown
    exit
  exit
  ppp
    match-list circuit-id mac username
    host "open-ppp-3" create
    host-identification
      username "open-pppoe-3"
    exit
    auth-policy "auth-pol-1"
    password chap "letmein"
    identification-strings 254 create
      sla-profile-string "ludb-sla"
      sub-profile-string "ludb-sub"
    exit
    msap-defaults
      group-interface "grp-int-1"

```

```

        policy "msap-pol-1"
        service 1
        exit
        no shutdown
    exit
exit
no shutdown
exit
exit
exit

```

RADIUS is defined via the **authentication-policy**. The domain name can be appended to the PPPoE subscriber host directly via authentication-policy, while for IPoE subscribers the domain name is appended via authentication-policy in conjunction with LUDB.

```

configure
  subscriber-mgmt
    authentication-policy "auth-pol-1" create
    description "RADIUS authentication policy"
    password "letmein"
    ppp-user-name append "domain1"
    user-name-format circuit-id append
    accept-authorization-change
    pppoe-access-method pap-chap
    radius-server-policy "rad-serv-pol-1"
  exit
exit
exit

```

The RADIUS user configuration file uses the domain extension as inserted by the BNG node to authenticate the user. The **inter-dest-id** string and the host IP address are provided by the RADIUS server (proxy case) along with other IP addressing parameters.

The IPv4 lease time (30 minutes) for IPv4 addresses are provided by the RADIUS server, while the lease time (30 minutes) for IPv6 addresses/prefixes are configured under the **group-interface**.

```

open-dhcp-3@domain1  Cleartext-Password := "letmein"
    Alc-Int-Dest-Id-Str = open-auth-vport,
    Framed-IP-Address = 10.10.0.230,
    Framed-IP-Netmask = 255.255.255.0,
    Alc-Default-Router = 10.10.0.1,
    Alc-Lease-Time = 1800,
    Client-DNS-Pri = 172.16.20.20,
    Client-DNS-Sec = 172.16.20.21,
    Alc-IPv6-Address = 2001:db8::100,
    Delegated-IPv6-Prefix = 2001:DB8:40:100::/56,
    Alc-IPv6-Primary-Dns = 2001:DB8::2000,
    Alc-Ipv6-Secondary-Dns = 2001:DB8::2001,
    Framed-Route = "192.168.1.0/24 0.0.0.0",

open-pppoe-3@domain1  Cleartext-Password := "letmein"
    Alc-Int-Dest-Id-Str = open-auth-vport,
    Framed-IP-Address = 10.10.0.231,
    Framed-IP-Netmask = 255.255.255.255,
    Client-DNS-Pri = 172.16.20.20,
    Client-DNS-Sec = 172.16.20.21,
    Alc-IPv6-Address = 2001:db8:0:1::100,
    Delegated-IPv6-Prefix = 2001:DB8:40:200::/56,
    Alc-IPv6-Primary-Dns = 2001:DB8::2000,
    Alc-Ipv6-Secondary-Dns = 2001:DB8::2001,
    Framed-Route = "192.168.2.0/24 0.0.0.0",

```

The group-interface configuration is as follows. Note that common parts of the configuration as defined earlier still apply.

```
configure
  service
    vprn 1 customer 1 create
      subscriber-interface "sub-int-1" create
        group-interface "grp-int-1" create
          ipv6
            dhcp6
              proxy-server
                renew-timer min 7
                rebind-timer min 10
                valid-lifetime min 30
                preferred-lifetime min 15
                client-applications dhcp ppp
                no shutdown
              exit
            exit
          exit
        dhcp
          proxy-server
            emulated-server 10.12.0.1
            no shutdown
          exit
        exit
      exit
    exit
  exit
exit
```

RADIUS sub/sla-profiles supplied via the LUDB are used:

```
configure
  subscriber-mgmt
    sla-profile "ludb-sla" create
      description "sla-profile obtained via LUDB"
      host-limits
        overall 100
      exit
    egress
      qos 1 vport-scheduler
      exit
    ip-filter 1
    exit
  exit
  sub-profile "ludb-sub" create
    description "sub-profile obtained via LUDB"
    egress
      agg-rate-limit 15000
    exit
  exit
exit
exit
exit
```

Show Commands

The following command shows that the LUDB-sub/sla-profiles, as supplied via LUDB, are in use. The IP addresses are supplied via the RADIUS server. The subscriber-id is auto-generated (not returned via

LUDB or RADIUS) and it is set to circuit-id for the IPoE subscriber-host, and to the **username|session-id** combination for PPPoE subscriber host.

```
*A:BNG-1# show service active-subscribers

=====
Active Subscribers
=====
Subscriber open-dhcp-3 (ludb-sub)
-----
(1) SLA Profile Instance sap:[1/1/1:13] - sla:ludb-sla
-----
IP Address      MAC Address      Session      Origin      Svc      Fwd
-----
10.10.0.230     00:0c:29:00:00:13 N/A          DHCP        1         Y
2001:db8::100/128 00:0c:29:00:00:13 N/A          DHCP6       1         Y
2001:db8:40:100::/56 00:0c:29:00:00:13 N/A          DHCP6       1         Y
-----

Subscriber open-pppoe-3|1 (ludb-sub)
-----
(1) SLA Profile Instance sap:[1/1/1:23] - sla:ludb-sla
-----
IP Address      MAC Address      Session      Origin      Svc      Fwd
-----
10.10.0.231     00:0c:29:00:00:23 PPP 1        IPCP        1         Y
2001:db8:0:1::100/128 00:0c:29:00:00:23 PPP 1        DHCP6       1         Y
2001:db8:40:200::/56 00:0c:29:00:00:23 PPP 1        DHCP6-PD    1         Y
-----

Number of active subscribers : 2
=====
*A:BNG-1#
```

The following command shows more details about the subscriber-host, such as the group-interface, vport, address origin, acct-session-id, etc. Vport is selected based on the **inter-dest-id** string supplied via RADIUS.

For the purpose of brevity, the output for only two hosts is shown, one with IPv4 address and one with IPv6 prefix. The remaining IP addresses/prefixes are not shown because the output follows the same logic.

```
*A:BNG-1# show service id 1 subscriber-hosts detail

=====
Subscriber Host table
=====
Sap      Subscriber
IP Address
MAC Address  PPPoE-SID Origin      Fwding State
```



```

-----
[1/1/1:13]          open-dhcp-3
 10.10.0.230
 00:0c:29:00:00:13   N/A      DHCP      Fwding
-----
Subscriber-interface : sub-int-1
Group-interface      : grp-int-1
Sub Profile          : ludb-sub
SLA Profile          : ludb-sla
App Profile          : N/A
Egress Q-Group       : N/A
Egress Vport         : N/A
Acct-Session-Id      : 14F2FF0000003A59119C39
Acct-Q-Inst-Session-Id: 14F2FF0000003B59119C39
Address Origin       : AAA
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status   : N/A
OT HTTP Rdr Fltr Src : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-----
[1/1/1:13]          open-dhcp-3
 2001:db8::100/128
 00:0c:29:00:00:12   N/A      IPoE-DHCP6  Fwding
-----

---snip---

-----
[1/1/1:13]          open-dhcp-3
 2001:db8:40:100::/56
 00:0c:29:00:00:12   N/A      IPoE-DHCP6  Fwding
-----

---snip---

-----
[1/1/1:23]          open-pppoe-3|1
 10.10.0.231
 00:0c:29:00:00:23    1      IPCP      Fwding
-----

---snip---

-----
[1/1/1:23]          open-pppoe-3|1
 2001:db8:0:1::100/128
 00:0c:29:00:00:23    1      PPP-DHCP6  Fwding
-----

---snip---

-----
[1/1/1:23]          open-pppoe-3|1
 2001:db8:40:200::/56
 00:0c:29:00:00:23    1      PPP-DHCP6  Fwding
-----
Subscriber-interface : sub-int-1
Group-interface      : grp-int-1
Sub Profile          : ludb-sub
SLA Profile          : ludb-sla
App Profile          : N/A
Egress Q-Group       : N/A

```

```
Egress Vport      : N/A
Acct-Session-Id   : 14F2FF0000004259119D4E
Acct-Q-Inst-Session-Id: 14F2FF0000004059119D4E
Address Origin    : AAA
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status : N/A
OT HTTP Rdr Fltr Src : N/A
HTTP Rdr URL Override : N/A
GTP local break-out : No
DIAMETER session ID Gx: N/A
-----
Number of subscriber hosts : 6
=====
*A:BNG-1#
```

The following command shows that the subscriber identity is set to **circuit-id** (plus **session-id**) as instructed by the **auto-sub-id-key** command (the **subscriber-id** string is not returned via LUDB or RADIUS). The lease times are set to 30 minutes as defined by RADIUS for IPv4 addresses and by the group-interface for IPv6 addresses/prefixes (proxy-case). The username passed to RADIUS is the circuit-id or username appended with the *domain1* domain name. The origin of the lease is RADIUS.

```
*A:BNG-1# show service id 1 dhcp lease-state detail

=====
DHCP lease states for service 1
=====
Service ID       : 1
IP Address       : 10.10.0.230
Client HW Address : 00:0c:29:00:00:13
Subscriber-interface : sub-int-1
Group-interface  : grp-int-1
SAP              : [1/1/1:13]
Termination Type : local
Up Time          : 0d 00:06:17
Remaining Lease Time : 0d 00:23:43
Remaining SessionTime: N/A
Persistence Key   : N/A

Sub-Ident        : "open-dhcp-3"
Sub-Profile-String : "ludb-sub"
SLA-Profile-String : "ludb-sla"
App-Profile-String : ""
Lease ANCP-String : ""
Lease Int Dest Id : "open-auth-vport"
Category-Map-Name : ""

Lease Info origin : Radius

Ip-Netmask       : 255.255.255.0
Broadcast-Ip-Addr : 10.10.0.255
Default-Router    : 10.10.0.1
Primary-Dns       : 172.16.20.20
Secondary-Dns     : 172.16.20.21
Primary-Nbns      : N/A
Secondary-Nbns    : N/A

ServerLeaseStart  : 05/09/2017 12:38:49
ServerLastRenew   : 05/09/2017 12:38:49
ServerLeaseEnd    : 05/09/2017 13:08:49
Session-Timeout   : N/A
IPoE|PPP session  : No
Lease-Time        : 0d 00:30:00
```

```
DHCP Server Addr      : N/A

Relay Agent Information
  Circuit Id          : open-dhcp-3
  Radius User-Name    : "open-dhcp-3@domain1"

-----
Managed Routes
-----


| IP Address     | Status    | Metric | Tag  | Pref |
|----------------|-----------|--------|------|------|
| 192.168.1.0/24 | installed | 0      | none | 0    |


-----
Number of lease states : 1
=====
*A:BNG-1#
```

For the purpose of brevity, the details for only two IPv6 prefixes are shown. The remaining two IPv6 leases are not shown because the output follows the same logic.

```
*A:BNG-1# show service id 1 dhcp6 lease-state detail

=====
DHCP lease states for service 1
=====
Service ID          : 1
IP Address          : 2001:db8::100/128
Client HW Address   : 00:0c:29:00:00:12
Subscriber-interface : sub-int-1
Group-interface     : grp-int-1
SAP                 : [1/1/1:13]
Termination Type    : local
Up Time             : 0d 00:06:24
Remaining Lease Time : 0d 00:23:37
Remaining SessionTime : N/A
Persistence Key      : N/A

Sub-Ident           : "open-dhcp-3"
Sub-Profile-String   : "ludb-sub"
SLA-Profile-String   : "ludb-sla"
App-Profile-String   : ""
Lease ANCP-String    : ""
Lease Int Dest Id    : "open-auth-vport"
Category-Map-Name    : ""
Dhcp6 ClientId (DUID) : 0001000120a33b41000c29000012
Dhcp6 IAID          : 2
Dhcp6 IAID Type      : non-temporary
Dhcp6 Client Ip      : fe80::20c:29ff:fe00:12
Primary-Dns          : 2001:db8::2000
Secondary-Dns        : 2001:db8::2001
Pool Name            : ""
Dhcp6 Server Addr    : N/A
Dhcp6 ServerId (DUID) : N/A
Dhcp6 InterfaceId    : open-dhcp-3
Dhcp6 RemoteId       : N/A
Radius sub-if prefix : N/A
Router adv. policy   : N/A

Lease Info origin    : Radius

ServerLeaseStart      : 05/09/2017 12:38:59
ServerLastRenew       : 05/09/2017 12:38:59
```

```

ServerLeaseEnd      : 05/09/2017 13:08:59
Session-Timeout    : N/A
IPoE|PPP session   : No
Radius User-Name    : "open-dhcp-3@domain1"
-----
Service ID         : 1
IP Address         : 2001:db8:0:1::100/128
Client HW Address   : 00:0c:29:00:00:23
Subscriber-interface : sub-int-1
Group-interface     : grp-int-1
SAP                 : [1/1/1:23]

---snip---

-----
Service ID         : 1
IP Address         : 2001:db8:40:100::/56
Client HW Address   : 00:0c:29:00:00:12
Subscriber-interface : sub-int-1
Group-interface     : grp-int-1
SAP                 : [1/1/1:13]

---snip---

-----
Service ID         : 1
IP Address         : 2001:db8:40:200::/56
Client HW Address   : 00:0c:29:00:00:23
Subscriber-interface : sub-int-1
Group-interface     : grp-int-1
SAP                 : [1/1/1:23]
Termination Type    : local
Up Time             : 0d 00:01:58
Remaining Lease Time : 0d 00:28:02
Remaining SessionTime : N/A
Persistence Key      : N/A

Sub-Ident          : "open-pppoe-3|1"
Sub-Profile-String  : "ludb-sub"
SLA-Profile-String  : "ludb-sla"
App-Profile-String  : ""
Lease ANCP-String   : ""
Lease Int Dest Id    : "open-auth-vport"
Category-Map-Name    : ""
Dhcp6 ClientId (DUID) : 0001000120a45903000c29000023
Dhcp6 IAID          : 1
Dhcp6 IAID Type      : prefix
Dhcp6 Client Ip      : fe80::20c:29ff:fe00:23
Primary-Dns          : 2001:db8::2000
Secondary-Dns        : 2001:db8::2001
Pool Name            : ""
Dhcp6 Server Addr    : N/A
Dhcp6 ServerId (DUID) : N/A
Dhcp6 InterfaceId    : open-pppoe-3
Dhcp6 RemoteId       : N/A
Radius sub-if prefix : N/A
Router adv. policy   : N/A

Lease Info origin    : Radius

ServerLeaseStart     : 05/09/2017 12:43:26
ServerLastRenew      : 05/09/2017 12:43:26
ServerLeaseEnd       : 05/09/2017 13:13:26
Session-Timeout      : N/A

```

```
IPoE|PPP session      : PPP
Radius User-Name      : "open-pppoe-3@domain1"
-----
Number of lease states : 4
=====
*A:BNG-1#
```

IP Proxy Case with LUDB + RADIUS Authentication

The IP address is assigned via LUDB.

- RADIUS provides sub/sla-profile strings and a framed IPv4 route.
- LUDB provides IP addresses (IPv6 lease-times are provided under the group-interface) and related parameters (DNS server, IPv4 default-gateway, etc), inter-dest-id string for Vport assignment and msap-defaults (routing context parameters and msap-policy).

Vport aggregate rate limit and the port scheduler are now added to the physical port. The Vport is associated with the subscriber through the inter-dest-id string obtained via the LUDB.

```
configure
  port 1/1/1
    ethernet
      mode access
      encap-type dot1q
      egress-scheduler-policy "port"
      access
        egress
          vport "open-dhcp" create
            agg-rate
              rate 500
            exit
            host-match dest "open-auth-vport" create
          exit
        exit
      exit
    exit
  no shutdown
exit
```

The LUDB is used to assign the inter-dest-id string, host IP addresses and IP addressing parameters. The DHCP lease time for IPv4 addresses is set to 15 minutes in the LUDB while lease times for IPv6 addresses/prefixes is set under the group-interface (set to 30 minutes).

The domain name *domain1* is appended to the username (circuit-id = open-dhcp-4 or username = open-pppoe-4) before an Access-Request is sent to the RADIUS server that is configured in the authentication-policy *auth-pol-1*.

```
configure
  subscriber-mgmt
    local-user-db "ludb-4" create
      description "for CASE-4"
      ipoe
        match-list circuit-id
        host "open-dhcp-4" create
          host-identification
            circuit-id string "open-dhcp-4"
          exit
        address 10.10.0.230
```

```

        auth-policy "auth-pol-1"
        auth-domain-name "domain1"
        identification-strings 254 create
            inter-dest-id "open-auth-vport"
        exit
        msap-defaults
            group-interface "grp-int-1"
            policy "msap-pol-1"
            service 1
        exit
        options
            subnet-mask 255.255.255.0
            default-router 10.10.0.254
            dns-server 172.16.20.20 172.16.20.21
            lease-time min 15
        exit
        options6
            dns-server 2001:DB8::2000 2001:DB8::2001
        exit
        ipv6-address 2001:DB8::100
        ipv6-delegated-prefix 2001:DB8:40:100::/56
        no shutdown
    exit
exit
ppp
match-list circuit-id mac username
host "open-ppp-4" create
    host-identification
        username "open-pppoe-4"
    exit
    auth-policy "auth-pol-1"
    address 10.10.0.231/32
    password chap "letmein"
    identification-strings 254 create
        inter-dest-id "open-auth-vport"
    exit
    msap-defaults
        group-interface "grp-int-1"
        policy "msap-pol-1"
        service 1
    exit
    options
        dns-server 172.16.20.20 172.16.20.21
    exit
    options6
        dns-server 2001:DB8::2000 2001:DB8::2001
    exit
    ipv6-address 2001:DB8::1:0:0:0:100
    ipv6-delegated-prefix 2001:DB8:40:200::/56
    no shutdown
    exit
exit
no shutdown
exit
exit
exit
exit

```

RADIUS is defined via the authentication-policy. The domain name can be appended to the PPPoE subscriber host directly via authentication-policy while for IPOE subscribers, the domain name is appended via authentication policy in conjunction with LUDB.

```

configure
subscriber-mgmt

```

```
authentication-policy "auth-pol-1" create
description "RADIUS authentication policy"
password "letmein"
ppp-user-name append "domain1"
user-name-format circuit-id append
accept-authorization-change
pppoe-access-method pap-chap
radius-server-policy "rad-serv-pol-1"
exit
exit
exit
```

The RADIUS user configuration file uses the domain extension as inserted by the SR OS to authenticate the user.

```
open-dhcp-4@domain1 Cleartext-Password := "letmein"
Alc-Subsc-Prof-Str = rad-sub,
Alc-SLA-Prof-Str = rad-sla,
Framed-Route = "192.168.1.0/24 0.0.0.0",

open-pppoe-4@domain1 Cleartext-Password := "letmein"
Alc-Subsc-Prof-Str = rad-sub,
Alc-SLA-Prof-Str = rad-sla,
Framed-Route = "192.168.2.0/24 0.0.0.0",
```

The group interface configuration is as follows. Common parts of the configuration as defined earlier still apply.

```
configure
service
  vprn 1 customer 1 create
    subscriber-interface "sub-int-1" create
    group-interface "grp-int-1" create
    ipv6
      user-db "ludb-4"
      dhcp6
        proxy-server
          renew-timer min 7
          rebind-timer min 10
          valid-lifetime min 30
          preferred-lifetime min 15
          client-applications dhcp ppp
          no shutdown
        exit
      exit
    exit
  exit
  dhcp
    proxy-server
      emulated-server 10.12.0.1
      no shutdown
    exit
    user-db "ludb-4"
  exit
exit
exit
exit
exit
exit
```

RADIUS sub/sla-profiles supplied by RADIUS are defined as:

```
configure
  subscriber-mgmt
    sla-profile "ludb-sla" create
      description "sla-profile obtained via LUDB"
      host-limits
        overall 3
      exit
      egress
        qos 1 vport-scheduler
      exit
      ip-filter 1
    exit
  exit
  sub-profile "ludb-sub" create
    description "sub-profile obtained via LUDB"
    egress
      agg-rate-limit 15000
    exit
  exit
exit
exit
```

Show Commands

The following command shows that the rad-sub/sla-profiles, as provided by RADIUS, are in use. The IP addresses are provided by LUDB. The **subscriber-id** is auto-generated (not returned via the LUDB or RADIUS) and it is set to **circuit-id** for IPoE subscriber-host(s) and to **username|session-id** combination for PPPoE subscriber host(s).

```
*A:BNG-1# show service active-subscribers

=====
Active Subscribers
=====
Subscriber open-dhcp-4 (rad-sub)
-----
(1) SLA Profile Instance sap:[1/1/1:14] - sla:rad-sla
-----
IP Address
-----
MAC Address      Session      Origin      Svc      Fwd
-----
10.10.0.230      00:0c:29:00:00:14  N/A         DHCP      1         Y
2001:db8::100/128 00:0c:29:00:00:14  N/A         DHCP6     1         Y
2001:db8:40:100::/56 00:0c:29:00:00:14  N/A         DHCP6     1         Y
-----

Subscriber open-pppoe-4|1 (rad-sub)
-----
(1) SLA Profile Instance sap:[1/1/1:24] - sla:rad-sla
-----
IP Address
```


MAC Address	Session	Origin	Svc	Fwd
10.10.0.231				
00:0c:29:00:00:24	PPP 1	IPCP	1	Y
2001:db8:0:1::100/128				
00:0c:29:00:00:24	PPP 1	DHCP6	1	Y
2001:db8:40:200::/56				
00:0c:29:00:00:24	PPP 1	DHCP6-PD	1	Y

Number of active subscribers : 2				
=====				
*A:BNG-1#				

The following command shows more details about the subscriber-host, such as the group-interface, vport, address origin, acct-session-id, etc. Vport is selected based on the **inter-dest-id** string as supplied via RADIUS.

For the purpose of brevity, the details for only two hosts is shown, one with IPv4 address and one with IPv6 prefix. The remaining IP addresses/prefixes are not shown because the output follows the same logic.

```
*A:BNG-1# show service id 1 subscriber-hosts detail

=====
Subscriber Host table
=====
Sap
  IP Address
  MAC Address
  Subscriber
  PPPoE-SID Origin
  Fwding State
-----
[1/1/1:14]      open-dhcp-4
  10.10.0.230
  00:0c:29:00:00:14  N/A      DHCP      Fwding
-----
Subscriber-interface : sub-int-1
Group-interface      : grp-int-1
Sub Profile          : rad-sub
SLA Profile          : rad-sla
App Profile          : N/A
Egress Q-Group       : N/A
Egress Vport         : open-dhcp
Acct-Session-Id      : 14F2FF000000475911D18B
Acct-Q-Inst-Session-Id: 14F2FF000000445911D087
Address Origin       : Static
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status   : N/A
OT HTTP Rdr Fltr Src : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
DIAMETER session ID Gx: N/A
-----
[1/1/1:14]      open-dhcp-4
  2001:db8::100/128
  00:0c:29:00:00:14  N/A      IPoE-DHCP6  Fwding
-----
Subscriber-interface : sub-int-1
Group-interface      : grp-int-1
Sub Profile          : rad-sub
SLA Profile          : rad-sla

---snip---
```

```

-----
[1/1/1:14]          open-dhcp-4
2001:db8:40:100::/56
00:0c:29:00:00:14   N/A      IPoE-DHCP6   Fwding
-----
Subscriber-interface : sub-int-1
Group-interface      : grp-int-1
Sub Profile          : rad-sub
SLA Profile          : rad-sla
App Profile          : N/A
---snip---
-----
[1/1/1:24]          open-pppoe-4|1
10.10.0.231
00:0c:29:00:00:24   1        IPCP          Fwding
-----
Subscriber-interface : sub-int-1
Group-interface      : grp-int-1
Sub Profile          : rad-sub
SLA Profile          : rad-sla
---snip---
-----
[1/1/1:24]          open-pppoe-4|1
2001:db8:0:1::100/128
00:0c:29:00:00:24   1        PPP-DHCP6   Fwding
-----
Subscriber-interface : sub-int-1
Group-interface      : grp-int-1
Sub Profile          : rad-sub
SLA Profile          : rad-sla
---snip---
-----
[1/1/1:24]          open-pppoe-4|1
2001:db8:40:200::/56
00:0c:29:00:00:24   1        PPP-DHCP6   Fwding
-----
Subscriber-interface : sub-int-1
Group-interface      : grp-int-1
Sub Profile          : rad-sub
SLA Profile          : rad-sla
App Profile          : N/A
Egress Q-Group       : N/A
Egress Vport         : open-dhcp
Acct-Session-Id      : 14F2FF0000004C5911D288
Acct-Q-Inst-Session-Id: 14F2FF0000004A5911D288
Address Origin       : Static
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status   : N/A
OT HTTP Rdr Fltr Src : N/A
HTTP Rdr URL Override : N/A
GTP local break-out  : No
DIAMETER session ID Gx: N/A
-----
Number of subscriber hosts : 6
=====
*A:BNG-1#

```

The following command shows that the subscriber identity is set to circuit-id (plus session-id) as instructed by the **auto-sub-id-key** command (the **subscriber-id** string is not returned via the LUDB or RADIUS). The DHCPv4 lease time is set to 15 minutes as defined by the LUDB. The DHCPv6 lease times are set to 30 minutes as configured under the group-interface. The username passed to RADIUS is the circuit-id or username appended with the *domain1* domain name. The origin of the lease is RADIUS.

```
*A:BNG-1# show service id 1 dhcp lease-state detail

=====
DHCP lease states for service 1
=====
Service ID           : 1
IP Address           : 10.10.0.230
Client HW Address    : 00:0c:29:00:00:14
Subscriber-interface : sub-int-1
Group-interface      : grp-int-1
SAP                  : [1/1/1:14]
Termination Type     : local
Up Time              : 0d 00:09:25
Remaining Lease Time : 0d 00:13:07
Remaining SessionTime: N/A
Persistence Key      : N/A

Sub-Ident            : "open-dhcp-4"
Sub-Profile-String   : "rad-sub"
SLA-Profile-String   : "rad-sla"
App-Profile-String   : ""
Lease ANCP-String    : ""
Lease Int Dest Id    : "open-auth-vport"
Category-Map-Name    : ""

Lease Info origin    : UserDb

Ip-Netmask           : 255.255.255.0
Broadcast-Ip-Addr    : 10.10.0.255
Default-Router       : 10.10.0.254
Primary-Dns          : 172.16.20.20
Secondary-Dns        : 172.16.20.21
Primary-Nbns         : N/A
Secondary-Nbns       : N/A

ServerLeaseStart      : 05/09/2017 16:26:19
ServerLastRenew       : 05/09/2017 16:33:50
ServerLeaseEnd        : 05/09/2017 16:48:50
Session-Timeout       : N/A
IPoE|PPP session     : No
Lease-Time            : 0d 00:15:00
DHCP Server Addr      : N/A

Relay Agent Information
  Circuit Id          : open-dhcp-4
  Radius User-Name    : "open-dhcp-4@domain1"

-----
Managed Routes
-----
IP Address           Status      Metric Tag      Pref
-----
192.168.1.0/24       installed    0      none      0
-----

Number of lease states : 1
=====
```

*A:BNG-1#

For the purpose of brevity, the details for only two IPv6 leases is shown. The remaining two IPv6 leases are not shown because the output follows the same logic.

*A:BNG-1# show service id 1 dhcp6 lease-state detail

```
=====
DHCP lease states for service 1
=====
```

```
Service ID      : 1
IP Address      : 2001:db8::100/128
Client HW Address : 00:0c:29:00:00:14
Subscriber-interface : sub-int-1
Group-interface  : grp-int-1
SAP              : [1/1/1:14]
Termination Type : local
Up Time          : 0d 00:13:48
Remaining Lease Time : 0d 00:23:11
Remaining SessionTime: N/A
Persistence Key   : N/A

Sub-Ident       : "open-dhcp-4"
Sub-Profile-String : "rad-sub"
SLA-Profile-String : "rad-sla"
App-Profile-String : ""
Lease ANCP-String : ""
Lease Int Dest Id : "open-auth-vport"
Category-Map-Name : ""
Dhcp6 ClientId (DUID): 0001000120a487cf000c29000014
Dhcp6 IAID      : 2
Dhcp6 IAID Type  : non-temporary
Dhcp6 Client Ip   : fe80::20c:29ff:fe00:14
Primary-Dns       : 2001:db8::2000
Secondary-Dns     : 2001:db8::2001
Pool Name        : ""
Dhcp6 Server Addr : N/A
Dhcp6 ServerId (DUID): N/A
Dhcp6 InterfaceId : open-dhcp-4
Dhcp6 RemoteId    : N/A
Radius sub-if prefix : N/A
Router adv. policy : N/A
```

```
Lease Info origin : UserDb

ServerLeaseStart   : 05/09/2017 16:22:26
ServerLastRenew    : 05/09/2017 16:29:25
ServerLeaseEnd     : 05/09/2017 16:59:25
Session-Timeout    : N/A
IPoE|PPP session   : No
Radius User-Name    : "open-dhcp-4@domain1"
```

```
-----
Service ID      : 1
IP Address      : 2001:db8:0:1::100/128
Client HW Address : 00:0c:29:00:00:24
Subscriber-interface : sub-int-1
Group-interface  : grp-int-1
SAP              : [1/1/1:24]
```

```
---snip---
```

```
-----
Service ID      : 1
```

```

IP Address      : 2001:db8:40:100::/56
Client HW Address : 00:0c:29:00:00:14
Subscriber-interface : sub-int-1
Group-interface  : grp-int-1
SAP              : [1/1/1:14]

---snip---

-----
Service ID      : 1
IP Address      : 2001:db8:40:200::/56
Client HW Address : 00:0c:29:00:00:24
Subscriber-interface : sub-int-1
Group-interface  : grp-int-1
SAP              : [1/1/1:24]
Termination Type : local
Up Time         : 0d 00:05:44
Remaining Lease Time : 0d 00:24:16
Remaining SessionTime: N/A
Persistence Key  : N/A

Sub-Ident       : "open-pppoe-4|1"
Sub-Profile-String : "rad-sub"
SLA-Profile-String : "rad-sla"
App-Profile-String : ""
Lease ANCP-String : ""
Lease Int Dest Id : "open-auth-vport"
Category-Map-Name : ""
Dhcp6 ClientId (DUID): 0001000120a45933000c29000024
Dhcp6 IAID      : 1
Dhcp6 IAID Type  : prefix
Dhcp6 Client Ip   : fe80::20c:29ff:fe00:24
Primary-Dns       : 2001:db8::2000
Secondary-Dns     : 2001:db8::2001
Pool Name        : ""
Dhcp6 Server Addr : N/A
Dhcp6 ServerId (DUID): N/A
Dhcp6 InterfaceId : open-pppoe-4
Dhcp6 RemoteId    : N/A
Radius sub-if prefix : N/A
Router adv. policy : N/A

Lease Info origin : UserDb

ServerLeaseStart   : 05/09/2017 16:30:32
ServerLastRenew    : 05/09/2017 16:30:32
ServerLeaseEnd     : 05/09/2017 17:00:32
Session-Timeout    : N/A
IPoE|PPP session   : PPP
Radius User-Name    : "open-pppoe-4@domain1"

-----
Number of lease states : 4
=====
*A:BNG-1#

```

Troubleshooting Commands

The following output shows the debugging commands which can be used to troubleshoot problems with the different authentication models.

```
debug
```

```
router "Base"
  radius
    packet-type authentication accounting coa
    detail-level medium
  exit
exit
router "1"
  ip
    dhcp
      detail-level high
      mode egr-ingr-and-dropped
    exit
    dhcp6
      mode egr-ingr-and-dropped
      detail-level high
    exit
  exit
  local-dhcp-server "svc-1-dhcp4"
    detail-level high
    mode egr-ingr-and-dropped
  exit
  local-dhcp-server "svc-1-dhcp6"
    detail-level high
    mode egr-ingr-and-dropped
  exit
exit
service
  id 1
    ppp
      packet
        mode egr-ingr-and-dropped
        detail-level high
        discovery
        ppp
        dhcp-client
      exit
    exit
  exit
  id 2
    dhcp
      mode egr-ingr-and-dropped
      detail-level high
    exit
    dhcp6
      mode all
      detail-level high
    exit
    ppp
      packet
        mode dropped-only
        detail-level high
        discovery
        ppp
        dhcp-client
      exit
    exit
  exit
exit
subscriber-mgmt
  local-user-db "ludb-1"
    detail all
  exit
  local-user-db "ludb-2"
    detail all
```

```
        exit
        local-user-db "ludb-3"
            detail all
        exit
        local-user-db "ludb-4"
            detail all
        exit
    exit
exit
```

```
configure
log
    log-id 1
        from debug-trace
        to session
        no shutdown
    exit
exit
exit
```

Conclusion

The flexible authentication model allows access to various sources (LUDB, RADIUS, and Python) of subscriber parameters during the subscriber establishment phase. This model can be utilized for IPoE, PPPoE or L2TP subscribers in IES or VPRN services (including a wholesale/retail VRF model). A typical use case would be in a wholesale/retail environment where the wholesaler enforces its own rules via the LUDB before it passes the authentication request to the retailer's RADIUS server.

GTP Access

This chapter provides information about GTP access.

Topics in this chapter include:

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

Applicability


The information and configuration in this chapter are based on SR OS Release 16.0.R3.

Overview

The GPRS Tunneling Protocol (GTP) is defined by 3GPP for carrying data through mobile backhaul networks. GTP-U (User plane) is used to forward User Equipment (UE) traffic between the Radio Access Network (RAN) and the core network. GTPv2-C (Control plane) is used within the 4G Evolved Packet Core (EPC) to establish and maintain these GTP-U tunnels. Basic Evolved Packet System (EPS) knowledge is assumed throughout this chapter.

Nokia provides Enhanced Subscriber Management (ESM) features to stationary wireless subscribers over GTP. This offers service providers the means to provide broadband services to areas that cannot be easily or sufficiently covered using traditional fixed access technologies. GTP is another type of access to the BNG, and the set of concepts applying to this context are generally referred to as fixed wireless access. In the Nokia GTP solution, the Serving Gateway (S-GW) and the Packet Data Network (PDN) Gateway (P-GW) functions are integrated with the BNG, as shown in [Figure 76: GTP Access to the BNG](#).

Figure 76: GTP Access to the BNG



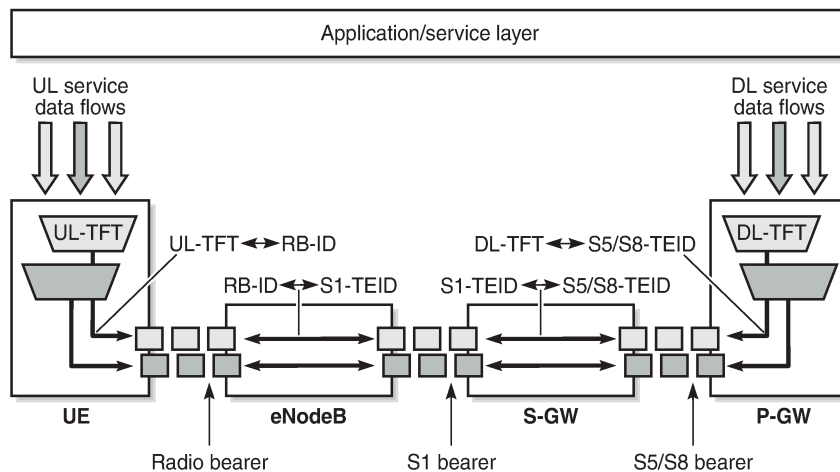
27982

In 3GPP architectures, specific communication channels between two entities are called interfaces. These 3GPP-defined interfaces are logical interfaces and unrelated to physical router interfaces. For example, the S11 interface covers the communication between the Mobility Management Entity (MME) and the S-GW, the S1-U interface the data path communication between the eNodeB and the S-GW, and the S5 interface the communication between the S-GW and the P-GW. The S1-MME interface covers the communication between the eNB and the MME. S11 and S1-U use GTP as their encapsulation protocol.

When a UE is activated, the UE attachment process is started, where the initial signaling passes via the eNB, the MME, and the S-GW to the P-GW to establish an EPS bearer. The EPS bearer runs from the UE via the eNB and the S-GW to the P-GW; the MME is not on the data path. Some subsequent signaling can run on the straight path from the UE to the P-GW.


The EPS bearer is a transmission path running end-to-end between a UE and a P-GW, which carries the user data; see [Figure 77: EPS Bearer Across the Different Interfaces](#). The EPS bearer is the concatenation of a radio bearer on the air interface, a GTP tunnel on the S1-U interface between the eNB and the S-GW, and a GTP tunnel on the S5 interface between the S-GW and the P-GW.

Figure 77: EPS Bearer Across the Different Interfaces



GTP-C signaling procedures are used to establish an S1 and an S5 bearer, during which Tunneling End ID (TEID) values are assigned and exchanged. These TEIDs are locally significant, and the upstream and downstream channels can have different TEIDs. With an S1 and S5 bearer established, user data can pass through using GTP-U encapsulation; see [Figure 78: GTP-C and GTP-U Encapsulation](#). UDP ports 2123 and 2152 are used for GTP-C and GTP-U, respectively.

Figure 78: GTP-C and GTP-U Encapsulation




27984

As part of the UE attachment procedure, the UE is assigned an IPv4/IPv6 address by the P-GW and a default bearer is established. The default bearer remains established throughout the lifetime of the PDN connection to provide the UE with always-on IP connectivity to that PDN. The default bearer QoS parameters can be provided by either the MME, Policy Control and Charging Rules Function (PCRF), RADIUS server, or locally.

The GTP-U tunnel encapsulation used in the network is shown in [Figure 79: GTP-U in Up and Downstream](#). The IPv4/IPv6 address assigned to the UE during the attachment process is used end-to-end. For the outer header of the S1 GTP tunnel, the eNB and the S-GW IPv4 addresses are used. Similarly, for the S5 GTP tunnel, the S-GW and P-GW IPv4 addresses are used.

Figure 79: GTP-U in Up and Downstream



27985

Because the S-GW and P-GW functionalities are integrated with the BNG, the S5 logical interface is internal to the system. SR OS supports TPSDA over GTP access tunnels initiated over the S11 (GTPv2-C) interface, where the UE data traffic is on the S1-U (GTP-U) interface. This is sometimes referred to as ESM over GTP (ESMoGTP).

Both IPv4 and IPv6 connectivity over GTP is supported. GTP session authentication can be performed using LUDB, RADIUS, or NASREQ.

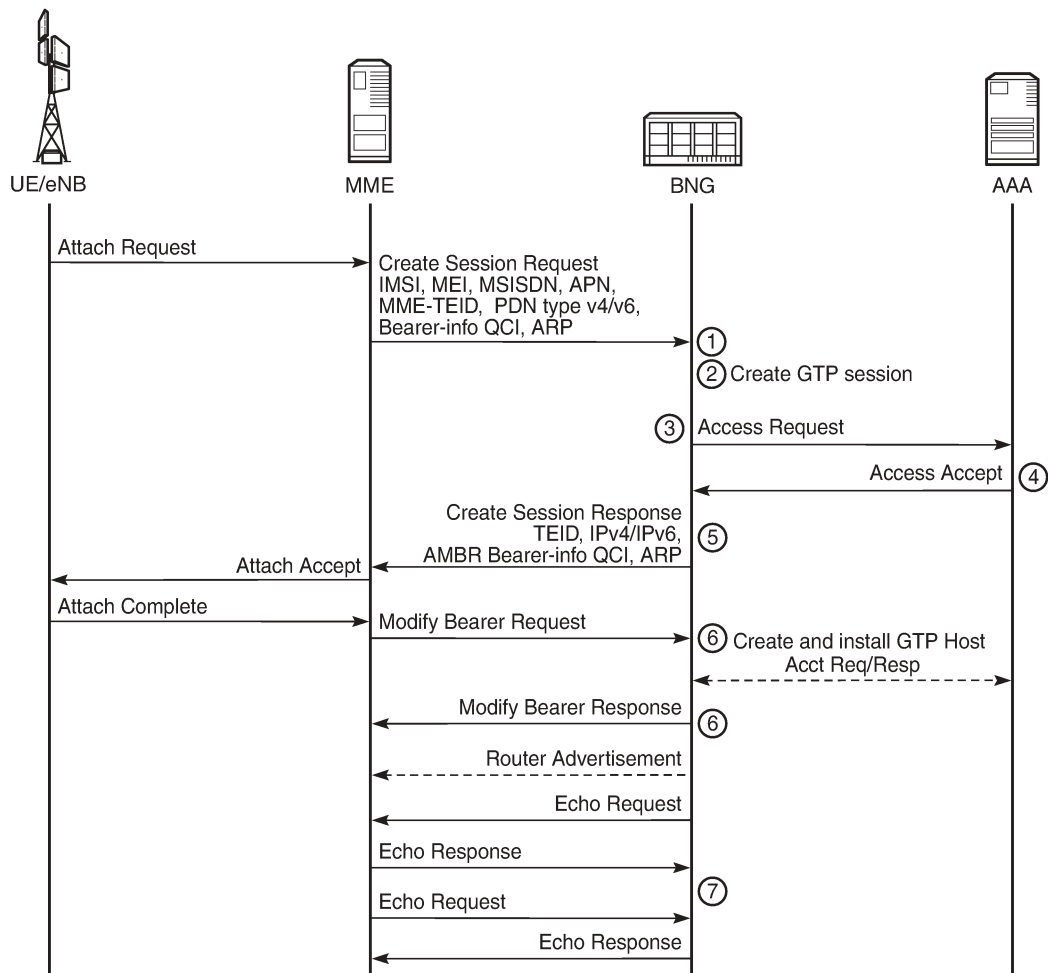
ESMoGTP requires PXC FPE in order to support L3 re-routing in the access network and in-line encapsulation/decapsulation of the GTP protocol.

GTP-C Control - GTP Host Creation and Host Deletion

Figure 80: GTP-C Control - GTP Host Creation shows the steps used for creating a GTP host.

1. A GTP S11 interface must be created and enabled for the BNG to accept and process the Create Session Request message emitted by the MME. This message includes IMSI, IMEI, MS-ISDN, APN, QCI, and so on.
2. The APN received is mapped via the APN policy to an authentication method (RADIUS, LUDB, or NASREQ), and to an optional default service and group interface. Also, a peer profile can be applied to that interface, defining keepalive timer (KA), default GTP Information Elements (IEs), Python-derived policies, and so on.
3. By default, the IMSI is used for authentication toward AAA.
4. The Access Accept message sent by AAA returns the standard set of ESM parameters, such as the IP address, QoS overrides, service ID, and group interface where the GTP host must be created.
5. The Create Session response includes the locally assigned TEID, IPv4/IPv6 information, AMBR, QCI, and the S1-U TEID and the S1-U endpoint. The GTP host is created.
6. The MME sends a Modify Bearer Request Message to update the default bearer information with the eNB final information (IP, TEID, and so on).
7. The BNG starts a keep-alive timer per S11 peer, using the timer values defined by the peer profile. Also, the MME can start a keep-alive timer, independent from the BNG.

Figure 80: GTP-C Control - GTP Host Creation




27986

Figure 81: GTP-C Control - GTP Host Deletion shows the actions taken for deleting a GTP host when triggered by a Delete Session Request message sent by the MME: on an incoming Delete Session Request message, the ESM host termination process is executed, and a Delete Session Response message is sent to the MME. This is marked as event 1 in **Figure 81: GTP-C Control - GTP Host Deletion**.

However, the BNG can also delete GTP hosts because of the echo timers expiring.

The BNG can autonomously start the GTP host deletion process, including the deletion of the GTP session; for example, when idle timers expire. The MME is notified through the Delete Bearer Request message so that it can delete the default bearer for the UE specified.


Figure 81: GTP-C Control - GTP Host Deletion



Configuration

Figure 82: GTP Access Topology shows the topology used for demonstrating GTP Access.

Figure 82: GTP Access Topology



On the BNG, the traffic entering on port 1/2/1 is internally routed to port 1/2/4, which is configured as a PXC port; see the *Port Cross-Connect (PXC)* chapter in the 7450 ESS, 7750 SR, 7950 XRS, and VSR Interface Configuration Advanced Configuration Guide for Classic CLI for more information. The PXC port 1/2/4 additionally requires subscriber management extensions to be enabled in the forwarding plane (FPE), see the *VXLAN Forwarding Path Extension* chapter in the 7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Advanced Configuration Guide for Classic CLI, and the **encap-type** is set to **qinq**:

```
# on BNG-1
configure
port 1/2/1
ethernet
```

```

        mode access
        encap-type dot1q
    exit
    no shutdown
exit
port 1/2/4
    ethernet
        mode hybrid
        encap-type dot1q
    exit
    no shutdown
exit
port-xc
    pxc 1 create
        port 1/2/4
        no shutdown
    exit
exit
port-xc
    pxc 1 create
        port 1/2/4
        no shutdown
    exit
exit
port pxc-1.a
    ethernet
        encap-type qinq
    exit
    no shutdown
exit
port pxc-1.b
    ethernet
        encap-type qinq
    exit
    no shutdown
exit
exit

```

```

configure
    fwd-path-ext
        fpe 1 create
            path pxc 1
            sub-mgmt-extensions
        exit
    exit
exit

```

The status of the PXC port can be verified as follows:

```
*A:BNG-1# show port-xc
```

```
=====
Port Cross-Connect Information
=====
```

PXC Id	Admin State	Oper State	Port Id	Description
1	Up	Up	1/2/4	(Not Specified)

```
-----
No. of PXCs: 1
=====
```

```
*A:BNG-1#
```

With subscriber management extensions enabled, SR OS can dynamically create SAPs on which GTP subscribers can be terminated at a later stage. The status for FPE 1 shows that the subscriber management extensions are enabled, as follows. The operational state indicated as N/A at the end of the same line is not relevant.

```
*A:BNG-1# show fwd-path-ext fpe 1

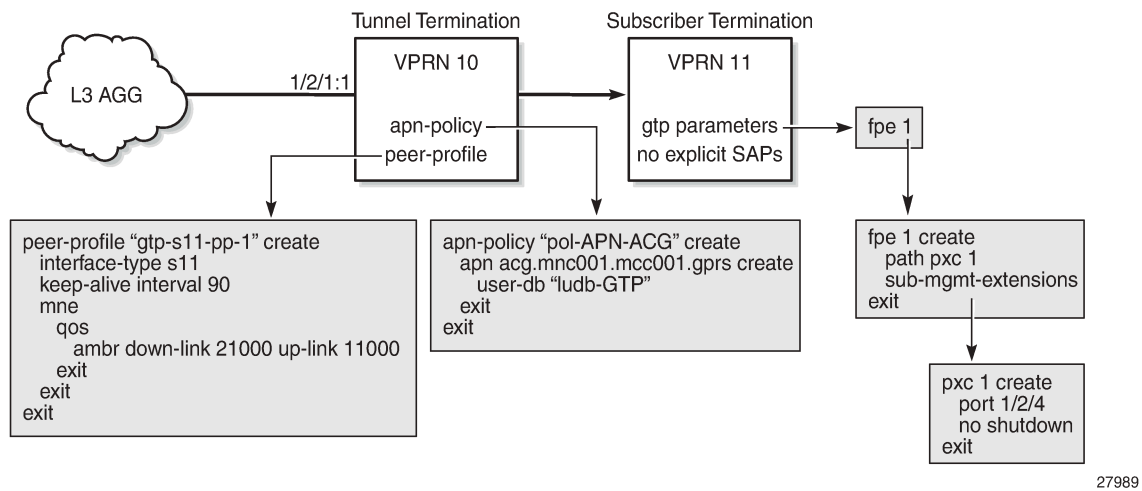
=====
FPE Id: 1
=====
Description      : (Not Specified)
Path             : pxc 1
Pw Port         : Disabled
Sub Mgmt Extension : Enabled          Oper      : down
Vxlan Termination : Disabled          Oper      : down
=====
*A:BNG-1#
```

The APN policy *pol-APN-ACG* is defined in the subscriber management **gtp** context, as follows. The other profiles are required for ESM host and subscriber setup.

```
configure
  subscriber-mgmt
    gtp
      apn-policy "pol-APN-ACG" create
      apn acg.mnc001.mcc001.gprs create
      user-db "ludb-GTP"
      exit
    exit
  exit
  sla-profile "sla-prof-1" create
  exit
  sub-profile "sub-prof-1" create
  exit
  sub-ident-policy "sub-id-pol-direct" create
  sub-profile-map
    use-direct-map-as-default
  exit
  sla-profile-map
    use-direct-map-as-default
  exit
  exit
exit
exit
exit
```

Two services are used to provide GTP access to the BNG, as shown in [Figure 83: GTP Tunnel and Subscriber Termination Configuration Logic](#). The first routed service provides the GTP tunnel termination, and the second provides the infrastructure to terminate the GTP subscribers.

Figure 83: GTP Tunnel and Subscriber Termination Configuration Logic



GTP Tunnel Terminating VPRN

In [Figure 83: GTP Tunnel and Subscriber Termination Configuration Logic](#), VPRN 10 implements the GTP tunnel termination. The GTP traffic enters VPRN 10 on interface *int-BNG-L3AGG* with MAC address 00:00:5e:00:53:01 on SAP 1/2/1:1. The interface *int-GTP-endpoint* is defined as a loopback interface, potentially providing redundancy to the L3 aggregation network. Static routes and a static ARP entry are defined in VPRN 10 for reaching the MME and the eNB through the L3 aggregation network, providing control over the L3 aggregation test environment. The APN policy provides the link between the tunnel termination VPRN (10) and the subscriber terminating VPRN (11). The peer-profile map defines which peer profile is used for which MME. In this example, the APN policy is *pol-APN-ACG*, and the peer profile is *gtp-s11-pp-1*. VPRN 10 is configured as follows:

```
configure
service
  vprn 10 name "GTP-tunnel-termination" customer 1 create
  description "GTP Tunnel termination"
  route-distinguisher 64496:10
  interface "int-GTP-endpoint" create
  description "Tunnel endpoint IP"
  address 192.0.2.10/32
  loopback
  exit
  interface "int-BNG-L3AGG" create
  address 192.168.0.1/24
  mac 00:00:5e:00:53:01
  static-arp 192.168.0.2 00:00:5e:00:53:00
  sap 1/2/1:1 create
  exit
  static-route-entry 192.168.100.0/24
  next-hop 192.168.0.2
  no shutdown
  exit
  static-route-entry 192.168.101.0/24
  next-hop 192.168.0.2
```

```

        no shutdown
    exit
exit
gtp
    s11
        interface "int-GTP-endpoint" create
            apn-policy "pol-APN-ACG"
        exit
        peer-profile-map
            address 192.168.100.1/32 peer-profile "gtp-s11-pp-1"
        exit
    exit
exit
no shutdown
exit
exit
exit
exit

```

Peer Profile

A peer-profile policy controls the keep-alive interval, and the QoS configuration to be used for MMEs (AMBR, ARP, QCI, ...). The peer profile *gtp-s11-pp-1* used for the example shown in Figure 8 is defined as follows. AMBR downlink and uplink are set to 21000 kb/s and 11000 kb/s, respectively, and QCI is set to 8.

```

configure
  subscriber-mgmt
    gtp
      peer-profile "gtp-s11-pp-1" create
        interface-type s11
        keep-alive interval 90
        mme
          qos
            ambr down-link 21000 up-link 11000
            arp 1
            down-link gbr 0 mbr 0
            qci 8
            up-link gbr 0 mbr 0
          exit
        exit
      exit
    exit
  exit
exit
exit
exit

```

APN Policy

An APN policy controls the authentication method per APN (LUDB, RADIUS, or Diameter). The APN policy *pol-APN-ACG* used for the example shown in [Figure 83: GTP Tunnel and Subscriber Termination Configuration Logic](#) is defined as follows. Only one APN is defined (*acg.mnc001.mcc001.gprs*), and the UEs using this APN are authenticated through the *ludb-GTP* LUDB.

```

configure
  subscriber-mgmt
    gtp
      apn-policy "pol-APN-ACG" create
        apn acg.mnc001.mcc001.gprs create
        user-db "ludb-GTP"

```

```

        exit
    exit
exit
exit

```

For authenticating GTP users using an LUDB, the LUDB must define hosts in the IPoE section of the LUDB. In this example, host matching is based on the IMSI, which is automatically populated in the IPoE string match criterion. Therefore, the **match-list string** command is used, but matching can also be based on a Derived ID provided by a Python script. Individual hosts can be matched using the **string** command for host-identification. To avoid the definition of many hosts, a *default* host can be used. In both cases, the output provides the service and the interface where the GTP subscriber must be implemented, as well as the IPv4/IPv6 address information required for the Local Address Assignment (LAA), as follows:

```

configure
  subscriber-mgmt
    local-user-db "ludb-GTP" create
    ipoe
      match-list string
      host "host-1" create
        host-identification
          string "00102000000111"
        exit
        address pool "pool4-1"
        identification-strings 254 create
          subscriber-id "GTP-subscriber-111"
          sla-profile-string "sla-prof-1"
          sub-profile-string "sub-prof-1"
        exit
        interface "int-GRP" service-id 11
        ipv6-slaac-prefix-pool "pool6-1"
        no shutdown
      exit
      host "default" create
        address pool "pool4-1"
        interface "int-GRP" service-id 11
        ipv6-slaac-prefix-pool "pool6-1"
        no shutdown
      exit
    exit
  no shutdown
exit
exit

```

GTP Subscriber Terminating VPRN

In [Figure 82: GTP Access Topology](#), VPRN 11 implements the GTP subscriber termination. Because GTP subscribers are ESM subscribers, the regular ESM concepts apply, meaning that subscriber management profiles are required; however, no explicit SAP definitions are required. LAA is used, where server *lcl-DHCPs-ip4* is used for IPv4, and *lcl-DHCPs-ip6* is used for IPv6. To support SLAAC, the IPv6 router-advertisements context has **prefix-options autonomous** enabled. The **gtp-parameters** context in the group interface defines the FPE to be used and must be enabled (no shutdown), as follows:

```

configure
  service
    vprn 11 name "GTP-subsc-termination" customer 1 create
    description "GTP subscriber termination"

```

```

---snip---

route-distinguisher 64496:11
subscriber-interface "int-SUSBC" create
  address 192.168.50.1/24
  ipv6
    delegated-prefix-len 56
    subscriber-prefixes
      prefix 2001:db8:101::/48 wan-host
    exit
  exit
group-interface "int-GRP" gtp create
  gtp-parameters
    fpe 1
    no shutdown
  exit
  ipv6
    router-advertisements
      prefix-options
        autonomous
      exit
      no shutdown
    exit
  exit
  local-address-assignment
    server "lcl-DHCPs-ip4"
    client-application ipoe-v4
    ipv6
      client-application ipoe-slaac
      server "lcl-DHCPs-ip6"
    exit
    no shutdown
  exit
  sap-parameters
    sub-sla-mgmt
      def-sla-profile "sla-prof-1"
      def-sub-profile "sub-prof-1"
      sub-ident-policy "sub-id-pol-direct"
    exit
  exit
exit
exit
no shutdown
exit
exit
exit

```

Enabling the **gtp-parameters** context (**no shutdown**) results in SR OS creating SAP pxc-1.b:1.8, as follows. The square brackets indicate that the SAP is created automatically, without any explicit configuration.

```
*A:BNG-1# show service id 11 sap
```

```
=====
SAP(Summary), Service 11
=====
```

PortId	SvcId	Ing. QoS	Ing. Fltr	Egr. QoS	Egr. Fltr	Adm	Opr
[pxc-1.b:1.8]	11	1	none	1	none	Up	Up

```
-----
Number of SAPs : 1

```

```
=====
*A:BNB-1#
```

When the UE with IMSI 001020000000111 connects, an S11 session is created, as follows. The 00:03:ff:f0:01:20 MAC address and pxc-1.b:1.8 SAP are internally generated to link the GTP session with an ESM IPoE session. The remote S1-U address is 192.168.101.1 for Bearer 6. The locally and remotely assigned TEID values are also shown.

```
*A:BNB-1# show subscriber-mgmt gtp s11 session imsi 001020000000111
```

```
=====
GTP S11 sessions
=====
```

```
IMSI                : 001020000000111
APN                  : acg.mnc001.mcc001.gprs
-----
Peer router          : 10
Peer address         : 192.168.100.1
Remote control TEID  : 22
Local control TEID   : 4293919008
PDN TEID             : 4293919008
Charging characteristics : (None)
Uplink AMBR (kbps)   : 10000
Downlink AMBR (kbps) : 20000
IpoE-session SAP     : [pxc-1.b:1.8]
IpoE-session Mac Address : 00:03:ff:f0:01:20
Bearer 6
  Rem S1-U address   : 192.168.101.1
  rem TEID           : 22
  loc TEID           : 4293919014
  uplink GBR (kbps)  : 0
  uplink MBR (kbps)  : 0
  downlink GBR (kbps) : 0
  downlink MBR (kbps) : 0
  QoS Class ID       : 8
  alloc/ret priority  : 1
-----
```

```
No. of GTP S11 sessions: 1
=====
```

```
*A:BNB-1#
```

Additionally, an ESM subscriber and an IPoE session is created on SAP pxc-1.b:1.8, as follows. The subscriber identifier is taken from the LUDB (GTP-subscriber-111). The IPoE session is created on service 11 and uses the internally generated MAC address. The IPv4 and the IPv6 addresses are combined in the same IPoE session, with MAC address 00:03:ff:e0:01:20.

```
*A:BNB-1# show service active-subscribers hierarchy
```

```
=====
Active Subscribers Hierarchy
=====
```

```
-- 001020000000111 (sub-prof-1)
|
+-- sap:[pxc-1.b:1.8] - sla:sla-prof-1
|
|-- IPOE-session - mac:00:03:ff:e0:01:20 - svc:11
|
|-- 192.168.50.12 - GTP
```

```

+-- 2001:db8:101:1::/64 - GTP
-----
Number of active subscribers : 1
Flags: (N) = the host or the managed route is in non-forwarding state
=====
*A:BNB-1#

```

Because an IPv4 address and an IPv6 address is used, two subscriber hosts are created on VPRN 11, as follows:

```

*A:BNB-1# show service id 11 subscriber-hosts
=====
Subscriber Host table
=====
Sap      Subscriber
IP Address
MAC Address      PPPoE-SID Origin      Fwding State
-----
[pxc-1.b:1.8]    001020000000111
192.168.50.12    00:03:ff:e0:01:20    N/A      GTP      Fwding
[pxc-1.b:1.8]    001020000000111
2001:db8:101:1::/64  00:03:ff:e0:01:20    N/A      GTP      Fwding
-----
Number of subscriber hosts : 2
=====
*A:BNB-1#

```

With two UEs connected to the BNG, the IPv4 and IPv6 addresses present in the routing tables are as follows. The addresses assigned to the UE that attached first are shown in bold.

```

*A:BNB-1# show router 11 route-table ipv4
=====
Route Table (Service: 11)
=====
Dest Prefix[Flags]      Type      Proto      Age      Pref
Next Hop[Interface Name]      Metric
-----
192.168.50.0/24          Local     Local      21h22m33s  0
int-SUSBC                0
192.168.50.12/32        Remote    Sub Mgmt    00h28m52s  0
[int-GRP]                 0
192.168.50.13/32        Remote    Sub Mgmt    00h01m13s  0
[int-GRP]                 0
-----
No. of Routes: 3
Flags: n = Number of times nexthop is repeated
      B = BGP backup route available
      L = LFA nexthop available
      S = Sticky ECMP requested
=====
*A:BNB-1#

```

```

*A:BNB-1# show router 11 route-table ipv6
=====
IPv6 Route Table (Service: 11)
=====

```

```
=====
Dest Prefix[Flags]                                Type  Proto  Age      Pref
  Next Hop[Interface Name]                        Metric
-----
2001:db8:101::/48                                Local  Local   21h23m02s  0
      int-SUSBC                                  0
2001:db8:101:1::/64                             Remote Sub Mgmt 00h29m21s  0
      [int-GRP]                                  0
2001:db8:101:2::/64                             Remote  Sub Mgmt 00h01m42s  0
      [int-GRP]                                  0
-----
No. of Routes: 3
Flags: n = Number of times nexthop is repeated
      B = BGP backup route available
      L = LFA nexthop available
      S = Sticky ECMP requested
=====
*A:BNG-1#
```

Debug

Debugging is useful when troubleshooting GTP scenarios. The debug configuration used is as follows:

```
debug
  router service-name "GTP-subsc-termination"
    ip
      dhcp
        detail-level high
        mode egr-ingr-and-dropped
      exit
    exit
  local-dhcp-server "lcl-DHCPs-ip4"
    detail-level high
    mode egr-ingr-and-dropped
  exit
  subscriber-mgmt
    local-user-db "ludb-GTP"
    detail all
  exit
  gtp
    event
    packet
      detail-level high
      mode all
    exit
  exit
exit
```

The trace for the UE with IMSI 001020000000111 initiating a connection to APN *acg.mnc001.mcc001.gprs* is as follows; also see [Figure 80: GTP-C Control - GTP Host Creation](#). Step 1 is the Create Session request, with the parameters supplied by the test-tool (IMSI, APN, RAT TYPE, PDN TYPE, and so on). Step 9 indicates that a GTP downlink is created. Step 12 indicates that the LUDB is accessed successfully. Steps 13, 14, and 15 indicate that LAA is used for address assignment. Step 28 indicates that an S1-U session is successfully established.

```
*A:BNG-1# show log log-id 1 ascending
```

```
=====
Event Log 1
=====
Description : (Not Specified)
Memory Log contents [size=100  next event=31  (not wrapped)]

1 2018/09/18 10:39:53.567 CEST MINOR: DEBUG #2001 vprn10 GTP
"GTP: GTPv2_INGRESS
IP Hdr: Src: 192.168.100.1, Dst: 192.0.2.10, Len: 205
UDP Hdr: Src: 2123, Dst: 2123, Len: 185
GTPv2 Hdr: Len: 173, Seq: 105, TEID: 0x0
GTPv2_INGRESS| S11-C: 192.168.100.1 | Rx: Create Session Req
[IMSI] : 001020000000111
[APN] : acg.mnc001.mcc001.gprs
[RAT TYPE] : EUTRAN
[CSID] : inst: 0 Len: 7 Val: 01010203040102
[PDN TYPE] : IPv4v6
[INDICATION] : GTP DAF
[S11 MME-C F-TEID] : 0x00000028 IPv4: 192.168.100.1
[BEARER CXT] : Add :0x6
Bearer Qos: PVI: 0x00 PL: 0x0f PCI: 0x00 QCI: 0x09
MBR: UL: 1000 kbps, DL: 1000 kbps GMBR: UL: 0 kbps, DL: 0 kbps
[RECOVERY] : 1
[PROTO CFG OPTS] : 0x80 (PPP_USE_IPPDP)
    IPCP_PROTO_ID: REQ
    PRIDNS : 0.0.0.0
    PRINBNS: 0.0.0.0
    SECDNS : 0.0.0.0
    SECNBNS: 0.0.0.0

    DNSv6_CONT_ID:REQ
    PCSCF_CONT_ID:REQ
    IPV4_LINK_MTU:REQ "
```

```
2 2018/09/18 10:39:53.567 CEST MINOR: DEBUG #2001 Base GTP
"GTP: GTP_USER_CREATED
    imsi = 001020000000111 imei = 0000000000000000
    msIsdn = 0000000000000000
    bssid = 0:0:0:0:0:0
    ssid =
    uli = 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
"
```

```
3 2018/09/18 10:39:53.567 CEST MINOR: DEBUG #2001 Base GTP
"GTP: GTP_ADD_DOWNLINK_TO_USER
    imsi = 001020000000111 imei = 0000000000000000
    msIsdn = 0000000000000000
    bssid = 0:0:0:0:0:0
    ssid =
    uli = 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
"
```

```
4 2018/09/18 10:39:53.567 CEST MINOR: DEBUG #2001 Base GTP
"GTP: GTP_PEER_CREATED
    peer = {2, 192.168.100.1, 2123} ver 2 laddr 192.0.2.10
"
```

```
5 2018/09/18 10:39:53.567 CEST MINOR: DEBUG #2001 Base GTP
"GTP: GTP_ADD_DOWNLINK_TO_PEER
    peer = {2, 192.168.100.1, 2123} ver 2 laddr 192.0.2.10
"
```



```
6 2018/09/18 10:39:53.567 CEST MINOR: DEBUG #2001 vprn10 GTP
"GTP: PEER ADDED
  peer = 192.168.100.1 peer port = 2123 src = 192.0.2.10
  Gtpv2 restartcnt = 1 path mgmt state = disabled
  ka retry cnt= 0
"

7 2018/09/18 10:39:53.567 CEST MINOR: DEBUG #2001 Base GTP
"GTP: RETRANSMIT_CTXT_ALLOCATED
  type Create Session Resp peer {2, 192.168.100.1, 2123} seqnr 105
"

8 2018/09/18 10:39:53.567 CEST MINOR: DEBUG #2001 vprn10 GTP
"GTP: PEER PATH MGMT ENABLED
  peer = 192.168.100.1 peer port = 2123 src = 192.0.2.10
  Gtpv2 restartcnt = 1 path mgmt state = up
  ka retry cnt= 0
"

9 2018/09/18 10:39:53.567 CEST MINOR: DEBUG #2001 vprn10 GTP
"GTP: GTP_DOWNLINK_CREATED
  gtp proto=GTP
  gtp iftype=S11
  loc s11 teid=0
  rem s11 teid=28
  s5 teid=0
  pdn type=0
  pdn v4 address::
  pdn v6 address::
  vrf=2
  ingIfIndex=3
  gtp-c src ip=192.0.2.10
  gtp-c dst ip=192.168.100.1
  gtp-u src ip=192.0.2.10
  gtp-u dst ip::
  chargChar=0
  dnLkAMBR=0
  upLkAMBR=0
  debug flag=1
  imei none
  msIsdn none
  def bearerId=0
"

10 2018/09/18 10:39:53.567 CEST MINOR: DEBUG #2001 Base GTP
"GTP: BEARER_CTXT_ALLOCATED
  imsi = 001020000000111 apn = acg.mnc001.mcc001.gprs
  id = 6
"

11 2018/09/18 10:39:53.567 CEST MINOR: DEBUG #2001 Base GTP
"GTP: GTP_ADD_BEARER_TO_RX_CTXT
  type Create Session Resp peer {2, 192.168.100.1, 2123} seqnr 105
"

12 2018/09/18 10:39:53.567 CEST MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
  string:
    original: 001020000000111
    masked: 001020000000111
```

```
Host host-1 found in user data base ludb-GTP"

13 2018/09/18 10:39:53.568 CEST MINOR: DEBUG #2001 vprn11 DHCP server
"DHCP server: lcl-DHCPs-ip4
Rx internal <NULL>
  ciaddr      : 0.0.0.0
"

14 2018/09/18 10:39:53.568 CEST MINOR: DEBUG #2001 vprn11 DHCP server
"DHCP server: lcl-DHCPs-ip4
lease added for 192.168.50.15 state=internalOffered
"

15 2018/09/18 10:39:53.568 CEST MINOR: DEBUG #2001 vprn11 DHCP server
"DHCP server: lcl-DHCPs-ip4
lease update for 192.168.50.15 state=internal
"

16 2018/09/18 10:39:53.570 CEST MINOR: DEBUG #2001 Base GTP
"GTP: GTP_REMOVE_BEARER_FROM_RX_CTXT
  type Create Session Resp peer {2, 192.168.100.1, 2123} seqnr 105
"

17 2018/09/18 10:39:53.570 CEST MINOR: DEBUG #2001 Base GTP
"GTP: RETRANSMIT_CTXT_UPDATED
  type Create Session Resp peer {2, 192.168.100.1, 2123} seqnr 105
"

18 2018/09/18 10:39:53.570 CEST MINOR: DEBUG #2001 vprn10 GTP
"GTP: GTPv2_EGRESS
IP Hdr: Src: 192.0.2.10, Dst: 192.168.100.1, Len: 228
UDP Hdr: Src: 2123, Dst: 2123, Len: 208
GTPv2 Hdr: Len: 196, Seq: 105, TEID: 0x28
I:*** | A:*** | S11-C: 192.0.2.10 | Tx: Create Session Resp
  [CAUSE]                : SUCCESS
  [RECOVERY]              : 6
  [PROTO CFG OPTS]       : 0x80 (PPP_USE_IPPDP)
    IPCP_PROTO_ID: REJ
    PRIDNS : 0.0.0.0
    PRINBNS: 0.0.0.0
    SECDNS : 0.0.0.0
    SECNBNS: 0.0.0.0

    IPV4_LINK_MTU: 1400
  [S11/S4-C SGW F-TEID] : 0xffb00120 IPv4: 192.0.2.10
  [S5/S8-C PGW F-TEID] : 0xffb00120 IPv4: 192.0.2.10
  [APN RESTRICTION]    : No Context/Restriction
  [PDN ADDR ALLOC]     : IPv4v6 192.168.50.15 2001:db8:101:2:17:ffff:fe00:1/64
  [AMBR]               : UL: 11000 kbps, DL: 21000 kbps
  [BEARER CXT]         : Add :0x6 Cause: SUCCESS
    Bearer Qos: PVI: 0x00 PL: 0x01 PCI: 0x00 QCI: 0x08
    MBR: UL: 0 kbps, DL: 0 kbps GMBR: UL: 0 kbps, DL: 0 kbps
  [S1-U SGW F-TEID]    : 0xffb00126 IPv4: 192.0.2.10
  [S5/S8-U PGW F-TEID] : 0xffb00126 IPv4: 192.0.2.10"

19 2018/09/18 10:39:53.570 CEST MINOR: DEBUG #2001 vprn10 GTP
"GTP: GTP_ACCESS_ACCEPT
  gtp proto=GTP
  gtp iftype=S11
  loc s11 teid=ffb00120
  rem s11 teid=28
  s5 teid=ffb00120
  pdn type=3
```

```
pdn v4 address=192.168.50.15
pdn v6 address=2001:db8:101:2:17:ffff:fe00:1
vrf=2
ingIfIndex=3
gtp-c src ip=192.0.2.10
gtp-c dst ip=192.168.100.1
gtp-u src ip=192.0.2.10
gtp-u dst ip=0.0.0.0
chargChar=0
dnLkAMBR=21000
upLkAMBR=11000
debug flag=1
imei none
msIsdn none
def bearerId=6

"

20 2018/09/18 10:39:53.592 CEST MINOR: DEBUG #2001 vprn10 GTP
"GTP: GTPv2_INGRESS
IP Hdr: Src: 192.168.100.1, Dst: 192.0.2.10, Len: 62
UDP Hdr: Src: 2123, Dst: 2123, Len: 42
GTPv2 Hdr: Len: 30, Seq: 106, TEID: 0xffb00120
GTPv2_INGRESS| S11-C: 192.168.100.1 | Rx: Modify Bearer Req
[BEARER CXT]          : Add :0x6
[S1-U eNB F-TEID]    : 0x00000028 IPv4: 192.168.101.1"

21 2018/09/18 10:39:53.592 CEST MINOR: DEBUG #2001 vprn10 GTP
"GTP: GTP_DOWNLINK_IN_SESSION_MSG_REVD
    gtp proto=GTP
    gtp iftype=S11
    loc s11 teid=ffb00120
    rem s11 teid=28
    s5 teid=ffb00120
    pdn type=3
    pdn v4 address=192.168.50.15
    pdn v6 address=2001:db8:101:2:17:ffff:fe00:1
    vrf=2
    ingIfIndex=3
    gtp-c src ip=192.0.2.10
    gtp-c dst ip=192.168.100.1
    gtp-u src ip=192.0.2.10
    gtp-u dst ip=0.0.0.0
    chargChar=0
    dnLkAMBR=21000
    upLkAMBR=11000
    debug flag=1
    imei none
    msIsdn none
    def bearerId=6
Gtpv2 Modify Bearer Req
"

22 2018/09/18 10:39:53.592 CEST MINOR: DEBUG #2001 Base GTP
"GTP: RETRANSMIT_CTXT_ALLOCATED
    type Modify Bearer Resp peer {2, 192.168.100.1, 2123} seqnr 106
"

23 2018/09/18 10:39:53.592 CEST MINOR: DEBUG #2001 Base GTP
"GTP: GTP_ADD_BEARER_TO_RX_CTXT
    type Modify Bearer Resp peer {2, 192.168.100.1, 2123} seqnr 106
"

24 2018/09/18 10:39:53.593 CEST MINOR: DEBUG #2001 vprn10 GTP
```

```
"GTP:  GTP_REMOVE_BEARER_FROM_DOWNLINK
      gtp proto=GTP
      gtp iftype=S11
      loc s11 teid=fffb00120
      rem s11 teid=28
      s5 teid=fffb00120
      pdn type=3
      pdn v4 address=192.168.50.15
      pdn v6 address=2001:db8:101:2:17:ffff:fe00:1
      vrf=2
      ingIfIndex=3
      gtp-c src ip=192.0.2.10
      gtp-c dst ip=192.168.100.1
      gtp-u src ip=192.0.2.10
      gtp-u dst ip=0.0.0.0
      chargChar=0
      dnLkAMBR=21000
      upLkAMBR=11000
      debug flag=1
      imei none
      msIsdn none
      def bearerId=6

"

25 2018/09/18 10:39:53.593 CEST MINOR: DEBUG #2001 Base GTP
"GTP:  BEARER_CTXT_FREED
      imsi = 001020000000111 apn = acg.mnc001.mcc001.gprs
      id = 6

"

26 2018/09/18 10:39:53.593 CEST MINOR: DEBUG #2001 Base GTP
"GTP:  GTP_REMOVE_BEARER_FROM_RX_CTXT
      type Modify Bearer Resp peer {2, 192.168.100.1, 2123} seqnr 106
"

27 2018/09/18 10:39:53.593 CEST MINOR: DEBUG #2001 vprn10 GTP
"GTP:  GTPv2_EGRESS
IP Hdr: Src: 192.0.2.10, Dst: 192.168.100.1, Len: 92
UDP Hdr: Src: 2123, Dst: 2123, Len: 72
GTPv2 Hdr: Len: 60, Seq: 106, TEID: 0x28
I:*** | A:*** | S11-C: 192.0.2.10 | Tx: Modify Bearer Resp
[CAUSE]                : SUCCESS
[RECOVERY]              : 6
[S11/S4-C SGW F-TEID]   : 0xffb00120 IPv4: 192.0.2.10
[BEARER CXT]            : Add :0x6 Cause: SUCCESS
[S1-U SGW F-TEID]       : 0xffb00126 IPv4: 192.0.2.10"

28 2018/09/18 10:39:53.593 CEST MINOR: DEBUG #2001 vprn10 GTP
"GTP:  GTP_S1U_ESTABLISHED
      gtp proto=GTP
      gtp iftype=S11
      loc s11 teid=fffb00120
      rem s11 teid=28
      s5 teid=fffb00120
      pdn type=3
      pdn v4 address=192.168.50.15
      pdn v6 address=2001:db8:101:2:17:ffff:fe00:1
      vrf=2
      ingIfIndex=3
      gtp-c src ip=192.0.2.10
      gtp-c dst ip=192.168.100.1
      gtp-u src ip=192.0.2.10
```

```

gtp-u dst ip=192.168.101.1
chargChar=0
dnLkAMBR=21000
upLkAMBR=11000
debug flag=1
imei none
msIsdn none
def bearerId=6
"

29 2018/09/18 10:40:23.347 CEST MINOR: DEBUG #2001 Base GTP
"GTP: RETRANSMIT_CTXT_FREED
  type Modify Bearer Resp peer {2, 192.168.100.1, 2123} seqnr 106
"

30 2018/09/18 10:40:23.347 CEST MINOR: DEBUG #2001 Base GTP
"GTP: RETRANSMIT_CTXT_FREED
  type Create Session Resp peer {2, 192.168.100.1, 2123} seqnr 105
"
*A:BNG-1#

```

The trace for the UE with IMSI 001020000000111 disconnecting from APN *acg.mnc001.mcc001.gprs* is as follows; also see [Figure 81: GTP-C Control - GTP Host Deletion](#). Step 1 is the Delete Session Request, initiated by the eNB. Step 4 indicates that the IP address is released. The GTP downlink is deleted and the bearer is removed from the GTP downlink. Step 13 indicates that the GTP user is deleted, step 15 indicates that the GTP peer is deleted.

```

*A:BNG-1# show log log-id 1 ascending

=====
Event Log 1
=====
Description : (Not Specified)
Memory Log contents [size=100  next event=16  (not wrapped)]

1 2018/09/18 10:00:37.063 CEST MINOR: DEBUG #2001 vprn10 GTP
"GTP: GTPv2_INGRESS
IP Hdr: Src: 192.168.100.1, Dst: 192.0.2.10, Len: 45
UDP Hdr: Src: 2123, Dst: 2123, Len: 25
GTPv2 Hdr: Len: 13, Seq: 104, TEID: 0xffc00120
GTPv2_INGRESS| S11-C: 192.168.100.1 | Rx: Delete Session Req
[EBI] : 0x6 "

2 2018/09/18 10:00:37.063 CEST MINOR: DEBUG #2001 vprn10 GTP
"GTP: GTP_DOWNLINK_IN_SESSION_MSG_REVD
  gtp proto=GTP
  gtp iftype=S11
  loc s11 teid=ffc00120
  rem s11 teid=27
  s5 teid=ffc00120
  pdn type=3
  pdn v4 address=192.168.50.14
  pdn v6 address=2001:db8:101:2:17:ffff:fe00:1
  vrf=2
  ingIfIndex=3
  gtp-c src ip=192.0.2.10
  gtp-c dst ip=192.168.100.1
  gtp-u src ip=192.0.2.10
  gtp-u dst ip=192.168.101.1
  chargChar=0
  dnLkAMBR=21000
  upLkAMBR=11000

```

```
debug flag=1
imei none
msIsdn none
def bearerId=6
Gtpv2 Delete Session Req
"

3 2018/09/18 10:00:37.063 CEST MINOR: DEBUG #2001 Base GTP
"GTP: RETRANSMIT_CTXT_ALLOCATED
  type Delete Session Resp peer {2, 192.168.100.1, 2123} seqnr 104
"

4 2018/09/18 10:00:37.064 CEST MINOR: DEBUG #2001 vprn11 DHCP server
"DHCP server: lcl-DHCPs-ip4
lease deleted for 192.168.50.14 (rxd internal release)
"

5 2018/09/18 10:00:37.064 CEST MINOR: DEBUG #2001 vprn10 GTP
"GTP: GTPv2_EGRESS
IP Hdr: Src: 192.0.2.10, Dst: 192.168.100.1, Len: 51
UDP Hdr: Src: 2123, Dst: 2123, Len: 31
GTPv2 Hdr: Len: 19, Seq: 104, TEID: 0x27
I:*** | A:*** | S11-C: 192.0.2.10 | Tx: Delete Session Resp
[CAUSE] : SUCCESS
[RECOVERY] : 6"

6 2018/09/18 10:00:37.064 CEST MINOR: DEBUG #2001 vprn10 GTP
"GTP: GTP_DOWNLINK_DELETED
  gtp proto=GTP
  gtp iftype=S11
  loc s11 teid=ffc00120
  rem s11 teid=27
  s5 teid=ffc00120
  pdn type=3
  pdn v4 address=192.168.50.14
  pdn v6 address=2001:db8:101:2:17:ffff:fe00:1
  vrf=2
  ingIfIndex=3
  gtp-c src ip=192.0.2.10
  gtp-c dst ip=192.168.100.1
  gtp-u src ip=192.0.2.10
  gtp-u dst ip=192.168.101.1
  chargChar=0
  dnLkAMBR=21000
  upLkAMBR=11000
  debug flag=1
  imei none
  msIsdn none
  def bearerId=6
"

7 2018/09/18 10:00:37.064 CEST MINOR: DEBUG #2001 vprn10 GTP
"GTP: GTP_REMOVE_BEARER_FROM_DOWNLINK
  gtp proto=GTP
  gtp iftype=S11
  loc s11 teid=ffc00120
  rem s11 teid=27
  s5 teid=ffc00120
  pdn type=3
  pdn v4 address=192.168.50.14
  pdn v6 address=2001:db8:101:2:17:ffff:fe00:1
  vrf=2
  ingIfIndex=3
```

```
gtp-c src ip=192.0.2.10
gtp-c dst ip=192.168.100.1
gtp-u src ip=192.0.2.10
gtp-u dst ip=192.168.101.1
chargChar=0
dnLkAMBR=21000
upLkAMBR=11000
debug flag=1
imei none
msIsdn none
def bearerId=6

"

8 2018/09/18 10:00:37.064 CEST MINOR: DEBUG #2001 Base GTP
"GTP: BEARER_CTXT_FREED
    imsi = 001020000000111 apn = acg.mnc001.mcc001.gprs
id = 6

"

9 2018/09/18 10:00:37.064 CEST MINOR: DEBUG #2001 Base GTP
"GTP: RETRANSMIT_CTXT_FREED
    type Delete Session Resp peer {2, 192.168.100.1, 2123} seqnr 104

"

10 2018/09/18 10:00:37.064 CEST MINOR: DEBUG #2001 Base GTP
"GTP: GTP_REMOVE_DOWNLINK_FROM_PEER
    peer = {2, 192.168.100.1, 2123} ver 2 laddr 192.0.2.10

"

11 2018/09/18 10:00:37.064 CEST MINOR: DEBUG #2001 Base GTP
"GTP: GTP_PEER_DELETED
    peer = {2, 192.168.100.1, 2123} ver 2 laddr 192.0.2.10

"

12 2018/09/18 10:00:37.064 CEST MINOR: DEBUG #2001 Base GTP
"GTP: GTP_REMOVE_DOWNLINK_FROM_USER
    imsi = 001020000000111 imei = 0000000000000000
    msIsdn = 0000000000000000
    bssid = 0:0:0:0:0:0
    ssid =
    uli = 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0

"

13 2018/09/18 10:00:37.064 CEST MINOR: DEBUG #2001 Base GTP
"GTP: GTP_USER_DELETED
    imsi = 001020000000111 imei = 0000000000000000
    msIsdn = 0000000000000000
    bssid = 0:0:0:0:0:0
    ssid =
    uli = 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0

"

14 2018/09/18 10:00:37.064 CEST MINOR: DEBUG #2001 vprn10 GTP
"GTP: PEERDB Peer not used
    peer = 192.168.100.1 peer port = 2123 src = 192.0.2.10
    Gtpv2 state = Active

"

15 2018/09/18 10:00:37.064 CEST MINOR: DEBUG #2001 vprn10 GTP
```

```
"GTP: PEER DELETED
  peer = 192.168.100.1 peer port = 2123 src = 192.0.2.10
  Gtpv2 restartcnt = 1 path mgmt state = up
  ka retry cnt= 0
```

```
"
```

Conclusion

Using GTP access technology, service providers can offer broadband services to areas that cannot be easily and sufficiently covered using traditional fixed access technologies. By offering LTE users GTP access to the BNG, the service providers can leverage the potential of wireless access with traditional BNG service offerings, including HQoS and multicast.

High Scale QoS IOM in ESM Context: Expanded SLA Mode

This chapter describes the High Scale QoS (HSQ) IOM in the Enhanced Subscriber Management (ESM) context

Topics in this chapter include:

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

Applicability

This chapter is applicable to the 7750 SR-7/12/12e platforms and describes the High Scale QoS (HSQ) IOM in the Enhanced Subscriber Management (ESM) context. The information and configuration in this chapter are based on SR OS Release 15.0.R4.

Overview

This chapter describes the QoS operation and configuration of the High Scale QoS IOM (HSQ), with a focus on the expanded SLA mode in the Enhanced Subscriber Management (ESM) context.

This chapter can be considered as an extension of the [High Scale QoS IOM in ESM Context: Single SLA Mode](#) chapter and it is recommended that the [High Scale QoS IOM in ESM Context: Single SLA Mode](#) chapter be read first. Generic concepts and configuration that is common to both chapters will not be repeated in this chapter. For example, high-level HSQ description or generic ESM configuration (as in Appendix A in the [High Scale QoS IOM in ESM Context: Single SLA Mode](#) chapter) will not be repeated here. However, since QoS is a centerpiece of this chapter, the QoS-specific configuration will be provided, even for the parts that may overlap between the two chapters (for example, HS attachment policy configuration is the same in both chapters and it will be repeated here).

Expanded SLA mode on HSQ is the default mode of operation in the ESM context and in contrast to single SLA mode, it allows multiple SLA profile instances (SPIs) per subscriber. That is, expanded SLA mode allows multiple HSQ queue groups per subscriber. Multiple subscriber **hostsessions** are supported per each SPI and HSQ queue group.

While expanded SLA mode provides better QoS flexibility at the subscriber level than single SLA mode, its scale becomes restricted by the number of HS primary shapers on the HSQ IOM. In Expanded SLA Mode, one HS primary shaper is allocated per subscriber to limit the subscriber aggregate rate (which can now consist of multiple HSQ queue groups). Since the HSQ IOM supports 16k HS primary shapers, the scale of the subscribers on HSQ IOM becomes limited by the same number. These two factors (greater QoS flexibility at the subscriber level and reduced subscriber scale) are the two main differences between the two SLA modes of operation.



Note:

Some small number of HS primary shapers are used internally by the system, so are not available to subscribers.

The SLA mode of operation is enabled per subscriber and it cannot be changed online (while the subscriber is established). The reason why expanded SLA mode is the default mode, even though it reduces the subscriber scale, is that this mode adheres to the existing ESM principles in SR OS where the number of SPIs per subscriber is not restricted.

Like the Single SLA Mode chapter, this chapter is also conceptually divided into two parts:

- The focus of the first part is on the (egress) QoS configuration for two subscribers, named "sub-1" and "sub-2". Each subscriber has two hosts and consequently two SPIs and HSQ queue groups. Subscriber hosts will be referred to as:
 - Host 1-1: first host of sub-1
 - Host 1-2: second host of sub-1
 - Host 2-1: first host of sub-2
 - Host 2-2: second host of sub-2



Note:

Subscriber sessions are also supported. This example is, however, based on subscriber hosts while the subscriber session concept is left disabled.

- The second part focuses on examining HSQ traffic management capabilities in the ESM context by observing output traffic patterns in a congested system.

Topics related to the generic operation of HSQ IOM and ESM that are not directly described in this chapter are included in the following:

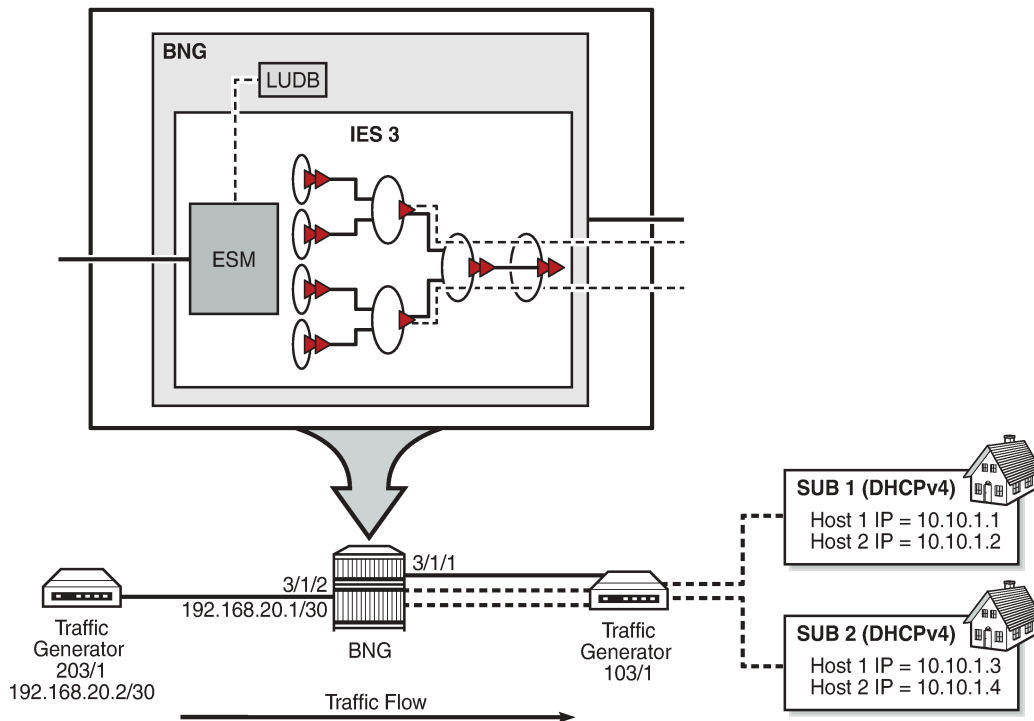
- Configuring single SLA mode in ESM is in the chapter [High Scale QoS IOM in ESM Context: Single SLA Mode](#).
- Configuring HSQ in the service context is in the chapter *High Scale QoS IOM: QoS, Service, and Network Configuration*.
- Generic ESM concepts are described in the *7450 ESS, 7750 SR, and VSR Triple Play Service Delivery Architecture Guide* and in numerous chapters in ACG part III – TPSDA.

Test Environment

Figure 84: Test Environment Example shows the test environment example with two subscribers in expanded SLA mode, each with two DHCPv4 hosts set up in the BNG. DHCPv4 control traffic (simulated subscriber hosts) is initiated from port 103/1 on a traffic generator. Subscribers are authenticated via local user database (LUDB) and instantiated on managed service access points (MSAPs) in the IES 3 service. Subscriber IP addresses are assigned statically via LUDB.

To explore traffic management capabilities under congestion, a number of traffic streams are generated in the downstream direction, from port 203/1 on the traffic generator toward the subscriber hosts on port 103/1.

Figure 84: Test Environment Example



26876

The QoS hierarchy associated with the subscribers is shown in [Figure 85: QoS Hierarchy in Expanded SLA Mode](#). At a high level, this hierarchy can be described as follows:

- Each of the two subscribers (sub-1 and "sub-2") has two hosts, each of which is associated with an SPI and HSQ queue group. HSQ queue group is always comprised of eight queues. However, not all the queues are required to be used by the subscriber. In this example, only six queues are used by each subscriber while the two remaining queues, although allocated to the subscriber, remain unused.



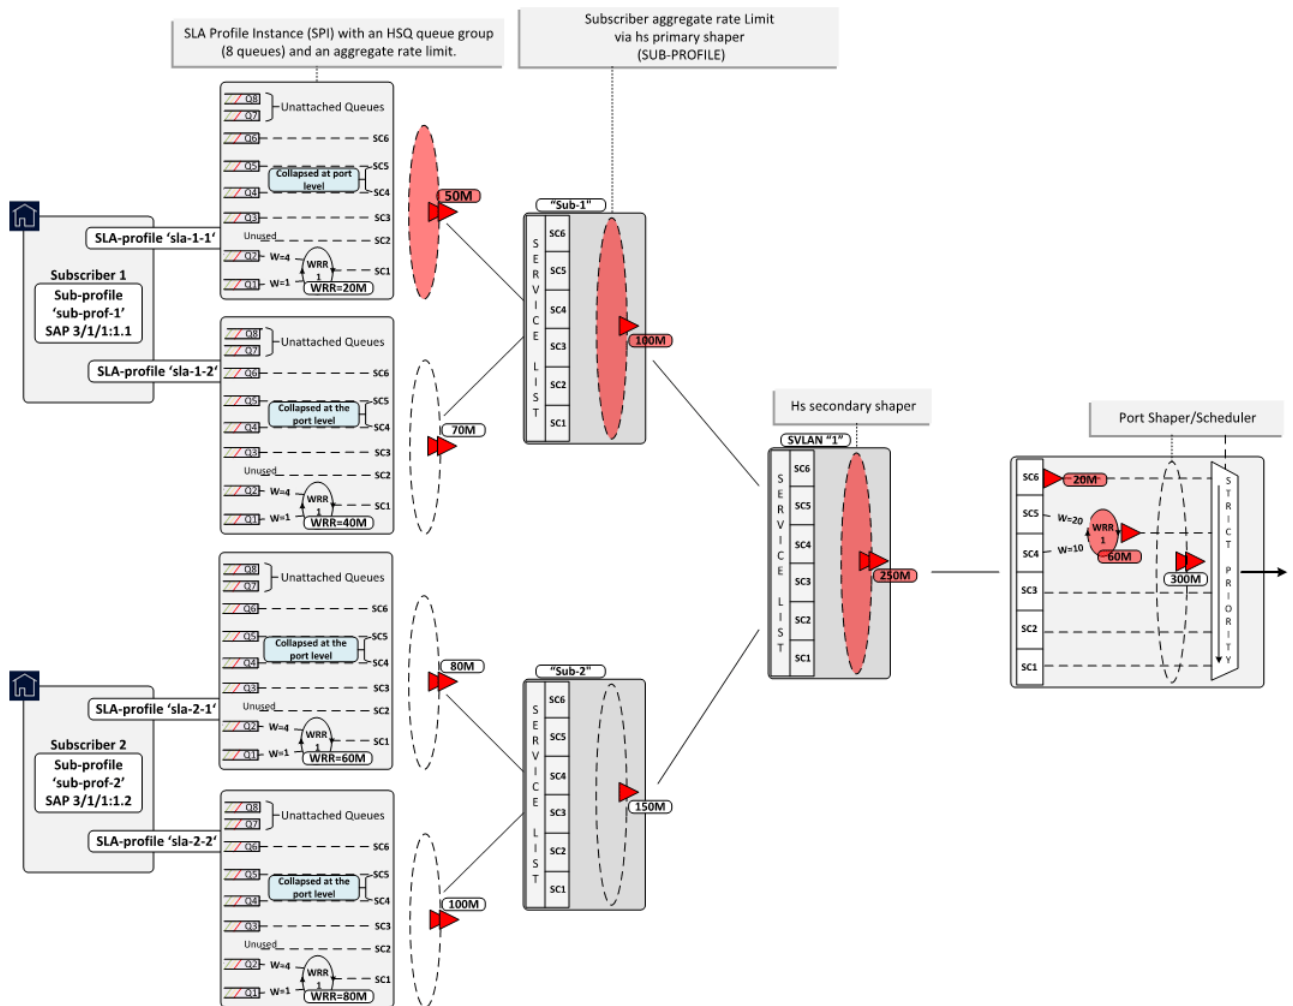
Note:

Subscriber sessions are also supported. This example is, however, based on subscriber hosts while the subscriber session concept is left disabled.

- All four hosts use the same HS attachment policy template. This means that the mapping between the queues, scheduling classes (SC), and WRR groups is the same for all four hosts.
- Queues 1 and 2 are attached to WRR group 1 at the local HSQ queue group level and served in 4:1 ratio. WRR group 1 is attached to the lowest scheduling class 1, whereas queues 3, 4, 5, and 6 are directly attached to scheduling classes 3, 4, 5, and 6, respectively.
- Each HSQ queue group is rate limited at the aggregate level. In this example, there is only one host associated with each HSQ queue group. However, multiple hosts per SPI and HSQ queue group can be configured.
- Each subscriber is rate limited at the aggregate level.
- Subscribers are mapped to the same HS secondary shapers.

- Mapping between the subscribers and their HS secondary shapers is achieved via outer VLANs. Sub-1 and sub-2 on SAPs with the outer VLAN 1 are mapped to the HS secondary shaper 1.
- The HS secondary shaper is associated with the HS port scheduler that has its own aggregate rate limit set. Because there is only one HS secondary shaper configured per port in this example, the smaller configured aggregate rate limit of the two (HS secondary shaper and the port shaper) will affect traffic going out this port.
- At the port level, the scheduling class 6 is rate limited while scheduling classes 4 and 5 are collapsed into a WRR group, which is rate limited.

Figure 85: QoS Hierarchy in Expanded SLA Mode



Configuration

The configuration section is split into two parts:

- ESM-specific configuration as it relates to HSQ

- QoS-specific configuration as it relates to HSQ

ESM-specific Configuration

For repeating parts of common ESM configuration, see the [High Scale QoS IOM in ESM Context: Single SLA Mode](#) chapter.

Each of the two subscribers (sub-1 and sub-2) have their own **sub-profile**, which will determine the HS SLA mode in which the subscriber operates, as well as the aggregate rate limit in kb/s for each subscriber. The aggregate rate of the subscriber is configured via its HS primary shaper. The **hs-agg-rate-limit** command in sub-profile is the only command that is applicable to HS primary shaper. Since the Expanded SLA Mode is the default mode, its configuration is not shown in the CLI unless requested by the **detail** keyword:

```
configure
  subscriber-mgmt
    sub-profile "sub-prof-1" create
      egress
        hs-agg-rate-limit 100000
      exit
    exit
    sub-profile "sub-prof-2" create
      egress
        hs-agg-rate-limit 150000
      exit
    exit
```

With **info detail** or **admin display-config detail** commands, the Expanded SLA Mode configuration is shown (irrelevant parts are removed from the output):

```
subscriber-mgmt
  sub-profile "sub-prof-1" create
    hs-sla-mode expanded
    egress
      hs-agg-rate-limit 100000
    exit
```

Each subscriber host has its own SPI that references the QoS SAP policy on egress (in this case, QoS SAP egress policies 10, 20, 30, and 40). Each SLA profile defines the aggregate rate in kb/s that will be applied to the corresponding HSQ queue group, and in this case to each subscriber host.

The egress QoS SAP policy defines the QoS characteristics at the subscriber host level (or HSQ queue group level), such as:

- traffic classification
- queue/WRR rates and weights
- mapping between the queues, WRR groups, and scheduling classes by referencing the HS attachment policy.

```
configure
  subscriber-mgmt
    sla-profile "sla-1-1" create
      egress
        qos 10
      exit
      hs-agg-rate-limit 50000
```

```

    exit
exit
sla-profile "sla-1-2" create
    egress
        qos 20
    exit
    hs-aggr-rate-limit 70000
    exit
exit
sla-profile "sla-2-1" create
    egress
        qos 30
    exit
    hs-aggr-rate-limit 80000
    exit
exit
sla-profile "sla-2-2" create
    egress
        qos 40
    exit
    hs-aggr-rate-limit 100000
    exit
exit
exit
exit

```

For detailed descriptions of the configuration options in MSAP policy, see the [High Scale QoS IOM in ESM Context: Single SLA Mode](#) chapter, because this configuration is shared:

```

configure
  subscriber-mgmt
    msap-policy "msaps" create
      sub-sla-mgmt
        def-inter-dest-id use-top-q
        sub-ident-policy "sub_ident_pol"
        no multi-sub-sap
        single-sub-parameters
        profiled-traffic-only
      exit
    exit
    ies-vprn-only-sap-parameters
      ingress
        qos 1 service-queuing
      exit
    exit
  exit
exit

```

QoS-specific configuration

At the subscriber level (or HSQ queue group level), the HS attachment policy defines mapping between the queues, WRR groups, and scheduling classes. In this example, queues 1 and 2 are mapped to a WRR group 1, which is mapped into scheduling class 1. Queues 3 to 6 are mapped to respective scheduling classes (SC 3 to 6). Queues 7 and 8, although allocated, are unattached, so will drop any traffic that they receive. A maximum of two WRR groups are supported at the HSQ queue group level and in this example, only WRR 1 is used.

The low burst max class parameter is set to 3, which ensures that the three lower scheduling classes (1 to 3) are stopped being served first if HSQ queue group aggregate congestion occurs. The HSQ queue group aggregate shaper has two burst tolerance thresholds: when the first threshold is reached, scheduling

classes 1 to 3 will be removed from the service list, followed by the higher-level scheduling classes (4 to 6) being removed when the second burst tolerance threshold is reached. This designation of scheduling classes in two tiers ensures lower latency for traffic associated with higher scheduling classes during short-lived congestion periods.

The following defined HS attachment policy is applied to the subscriber hosts through the SAP egress policy referenced in the SLA profile. In this example, all four subscriber hosts use the same HS attachment policy:

```
configure
qos
  hs-attachment-policy "hs-attach-1-1" create
    low-burst-max-class 3
    queue 1 wrr-group 1
    queue 2 wrr-group 1
    queue 3 sched-class 3
    queue 4 sched-class 4
    queue 5 sched-class 5
    queue 6 sched-class 6
    queue 7 unattached
    queue 8 unattached
    wrr-group 1 sched-class 1
    wrr-group 2 unattached
  exit
exit
```

Besides referencing the HS attachment policy, the SAP egress policy defines traffic classification, as well as characteristics of queues and WRR groups at the subscriber level. In this example, the SAP egress policy 10 is associated with the host 1 of subscriber sub-1. The queues 1 and 4 are in the context of WRR group 1 serviced in the ratio 1:4.



Note:

Although the SAP egress policy syntax implies that such policy is applied per SAP, in the ESM context, this policy is instantiated via the SLA profile; therefore, the queue/policer instantiations are performed per SPI and not per SAP.

The aggregate rate of WRR 1 is set to 20Mb/s. Mapping of forwarding classes to queues is self-explanatory. Separate SAP egress policies (20, 30, and 40) are applied to the remaining hosts (host 1-2, host 2-1, and host 2-2). The only difference between the QoS SAP egress policies for the subscriber hosts is the rate of the WRR group 1, which is 40Mb/s for host 1-2, 60Mb/s for host 2-1, and 80Mb/s for host 2-2.

The SAP egress policy is applied in the SLA profile for the subscriber. The WRR group rate (along with other QoS parameters) can be dynamically overridden via RADIUS/Diameter during authentication or while the host/session is online. This functionality would allow having one SAP egress policy configured where parameters can be dynamically overridden, when the policy is applied to the subscriber host or session.

```
configure
qos
  sap-egress 10 create
    hs-attachment-policy "hs-attach-1-1"
    queue 1 create
      hs-wrr-weight 1
    exit
    queue 2 create
      hs-wrr-weight 4
    exit
    queue 3 create
    exit
    queue 4 create
```

```

exit
queue 5 create
exit
queue 6 create
exit
hs-wrr-group 1
  rate 20000
exit
fc af create
  queue 3
exit
fc be create
  queue 1
exit
fc h1 create
  queue 5
exit
fc h2 create
  queue 6
exit
fc l1 create
  queue 4
exit
fc l2 create
  queue 2
exit
dscp af12 fc "af"
dscp be fc "be"
dscp af22 fc "h1"
dscp ef fc "h2"
dscp af21 fc "l1"
dscp af11 fc "l2"
exit

```

It was already shown in the [ESM-specific Configuration](#) section that the HSQ queue group aggregate rate limit is configured in the SLA profile while the aggregate rate of the subscriber is managed by the HS primary shaper and is configured in the sub-profile.

The next hierarchy level in the chain (up from the subscriber level) is performed by the HS secondary shapers that are directly configured in the egress context of the port. In this example, there is only one configured HS secondary shaper that corresponds to the outer VLAN on the subscriber SAPs. The network representation of the HS secondary shaper is an access node downstream from the BNG.

The HS secondary shaper "1" is configured with the aggregate rate of 250Mb/s. The rates are configured in kb/s. Similarly, as at the subscriber level, the **low-burst-max-class 3** command maps all scheduling classes at or below level 3 to the low burst tolerance threshold, while all scheduling classes above level 3 are mapped to the high burst tolerance threshold at the HS secondary shaper level. Therefore, in case of short congestion periods, objects associated with scheduling classes 3 and below will be removed from the service list before the objects associated with scheduling classes 4 and above.

```

configure
  port 3/1/1
    ethernet
      mode access
      encaps-type qinq
      egress
        hs-scheduler-policy "hs1"
        hs-secondary-shaper "1" create
          aggregate
            rate 250000
            low-burst-max-class 3

```



```

        exit
    exit
exit
no shutdown
exit

```

The last configuration block in the scheduling hierarchy is an HS port scheduler, which is associated with the port via the command **hs-scheduler-policy** "hs1" in the preceding CLI code.

The HS port scheduler characteristics in this example are defined as follows:

- The maximum rate is set to 300Mb/s. However, in this example, this rate limit has no effect on traffic because it is higher than the 250Mb/s rate configured on HS secondary shaper 1. Because HS secondary shaper 1 is the only one configured under the port, the lower rate of this HS secondary shaper will limit traffic before it reaches the 300Mb/s limit configured on the port. More realistic deployment scenarios would contain multiple secondary shapers per port.
- Scheduling classes 4 and 5 are collapsed into a single scheduling priority (5) and they are served in a 1:2 ratio. This is performed via WRR group 1.

This collapsing of scheduling classes 4 and 5 occurs at the port level, whereas scheduling classes 1 and 2 are collapsed at the subscriber level (HSQ queue group level).

- The WRR 1 at the port level is rate limited to 60Mb/s.
- The highest priority (6) scheduling class is rate limited to 20Mb/s at the port level.



Note:

All the rates under the HS port scheduler are configured in Mb/s.

```

configure
qos
  hs-scheduler-policy "hs1" create
    max-rate 200
    group 1 rate 60
    scheduling-class 4 group 1 weight 10
    scheduling-class 5 group 1 weight 20
    scheduling-class 6 rate 20
  exit
exit

```

The delta between the low and high burst tolerance thresholds at the subscriber and HS secondary shaper levels can be adjusted with the **hs-fixed-high-thresh-delta** command that is configured at the card level. In this example, the difference between the thresholds is set to 12kbytes.

```

configure
card 3
  card-type iom4-e-hs
  mda 1
    mda-type me10-10gb-sfp+
    no shutdown
  exit
  fp 1
    egress
      hs-fixed-high-thresh-delta 12000
    exit
  exit
no shutdown

```

Operational Commands

Operational commands are used to troubleshoot the system and monitor its operational state. The focus of this section will be on **show**, **clear**, and **tools dump** commands. The **debug** commands are omitted because there are no debug commands related to QoS on HSQ IOM and the **debug** commands related to ESM are described in other chapters. Also, there are no log events related to QoS on HSQ IOM.

Show Commands

show commands in this section are divided into two groups:

- **show** commands that display association between ESM and QoS objects where all displayed information is static and it does not change autonomously over time.
- **show** commands that display QoS hierarchy with the running rates (where the state is changing autonomously).

show commands have filters (or CLI parameters) that can be used to control the amount of the output information. This section will provide a few **show** commands; it is left to the user to explore all the options available for a particular **show** command.

Subscriber Management Related Show Commands

Examining subscriber associations with QoS objects should begin with the **show service active-subscribers** command. The output of this command provides information about the subscriber context and the output can vary in detail depending on the options with which this command is run. Besides the subscriber name, underlying subscriber SAPs, and SLA/sub-profile names, the following information is provided:

- SLA mode in which the subscriber operates on HSQ (single versus expanded SLA mode)
- Aggregate rate of the subscriber
- Ingress and egress QoS policies
- Subscriber association with an HS secondary shaper and the **inter dest id** string that is used to make this association
- Ingress queue/policer statistics
- Egress queues statistics

For brevity, only the information for subscriber sub-1 (both hosts) is shown in the following output.

```
*A:PE-1# show service active-subscribers subscriber "sub-1" detail

=====
Active Subscribers
=====
-----
Subscriber sub-1 (sub-prof-1)
-----
I. Policer Ctrl. : N/A
E. Policer Ctrl. : N/A
I. vport-hashing : Disabled
I. sec-sh-hashing: Disabled
Q Frame-Based Ac*: Disabled
```

```

Acct. Policy      : N/A                               Collect Stats   : Disabled
ANCP Pol.        : N/A
Accu-stats-pol   : (Not Specified)
HostTrk Pol.     : N/A
IGMP Policy      : N/A
MLD Policy       : N/A
PIM Policy       : N/A
Sub. MCAC Policy : N/A
NAT Policy       : N/A
Firewall Policy  : N/A
UPnP Policy      : N/A
NAT Prefix List  : N/A
Def. Encap Offset: none                               Encap Offset Mode: none
Avg Frame Size   : N/A
Vol stats type   : full
Preference       : 5
LAG hash class   : 1
LAG hash weight  : 1
Sub. ANCP-String : "sub-1"
Sub. Int Dest Id : "1"
Igmp Rate Adj    : N/A
RADIUS Rate-Limit: N/A
Oper-Rate-Limit  : 100000
-----
Radius Accounting
-----
Policy           : N/A
Session Opti.Stop: False
-----
HS
-----
SLA-mode         : expanded                           E Agg Rate Limit : 100000
Hs Second Shaper : "1"
* indicates that the corresponding row element may have been truncated.
-----
(1) SLA Profile Instance
  - sap:[3/1/1:1.1] (IES 3 - group-int-1)
  - sla:sla-1-1
-----
Description      : (Not Specified)
Host Limits      : No Limit
Egr Sched-Policy : N/A
Ingress Qos-Policy : 1                               Egress Qos-Policy : 10
Ingress Queuing Type : Service-queuing (Not Applicable to Policer)
Ingr IP Fltr-Id  : N/A                               Egr IP Fltr-Id    : N/A
Ingr IPv6 Fltr-Id : N/A                               Egr IPv6 Fltr-Id  : N/A
Ingress Report-Rate : Maximum
Egress Report-Rate : Maximum
Egress Remarking  : from Sap Qos
Credit Control Pol. : N/A
Category Map      : (Not Specified)
Use ing L2TP DSCP : false
Hs-Agg-Rate-Limit   : 50000
Egress HS Q stat mode: no-override
Hs-Oper-Rate-Limit  : 50000
Egr hqos mgmt status : disabled
-----
IP Address
-----
MAC Address      Session      Origin      Svc      Fwd
-----
10.10.1.1        00:00:64:01:01:01  N/A        DHCP      3        Y
-----

```

SLA Profile Instance statistics

	Packets	Octets
Off. HiPrio	: 0	0
Off. LowPrio	: 0	0
Off. Uncolor	: 0	0
Off. Managed	: 0	0
Queueing Stats (Ingress QoS Policy 1)		
Dro. HiPrio	: 0	0
Dro. LowPrio	: 0	0
For. InProf	: 0	0
For. OutProf	: 0	0
Queueing Stats (Egress QoS Policy 10)		
Dro. In/InplusProf	: 0	0
Dro. Out/ExcProf	: 89016218	89016218000
For. In/InplusProf	: 0	0
For. Out/ExcProf	: 74180626	74180626000

SLA Profile Instance per Queue statistics

	Packets	Octets
Ingress Queue 1 (Unicast) (Priority)		
Off. HiPrio	: 0	0
Off. LowPrio	: 0	0
Dro. HiPrio	: 0	0
Dro. LowPrio	: 0	0
For. InProf	: 0	0
For. OutProf	: 0	0
Egress Queue 1		
Dro. In/InplusProf	: 0	0
Dro. Out/ExcProf	: 34006842	34006842000
For. In/InplusProf	: 0	0
For. Out/ExcProf	: 3083545	3083545000
Egress Queue 2		
Dro. In/InplusProf	: 0	0
Dro. Out/ExcProf	: 24756213	24756213000
For. In/InplusProf	: 0	0
For. Out/ExcProf	: 12334174	12334174000
Egress Queue 3		
Dro. In/InplusProf	: 0	0
Dro. Out/ExcProf	: 0	0
For. In/InplusProf	: 0	0
For. Out/ExcProf	: 29672356	29672356000
Egress Queue 4		
Dro. In/InplusProf	: 0	0
Dro. Out/ExcProf	: 14981296	14981296000
For. In/InplusProf	: 0	0
For. Out/ExcProf	: 7272638	7272638000
Egress Queue 5		
Dro. In/InplusProf	: 0	0
Dro. Out/ExcProf	: 7708660	7708660000
For. In/InplusProf	: 0	0
For. Out/ExcProf	: 14545275	14545275000

```

Egress Queue 6
Dro. In/InplusProf : 0                0
Dro. Out/ExcProf   : 7563207          7563207000
For. In/InplusProf : 0                0
For. Out/ExcProf   : 7272638          7272638000

-----
(2) SLA Profile Instance
- sap:[3/1/1:1.1] (IES 3 - group-int-1)
- sla:sla-1-2
-----
Description          : (Not Specified)
Host Limits          : No Limit
Egr Sched-Policy     : N/A
Ingress Qos-Policy   : 1                Egress Qos-Policy : 20
Ingress Queuing Type : Service-queuing (Not Applicable to Policer)
Ingr IP Fltr-Id      : N/A                Egr IP Fltr-Id   : N/A
Ingr IPv6 Fltr-Id    : N/A                Egr IPv6 Fltr-Id : N/A
Ingress Report-Rate  : Maximum
Egress Report-Rate   : Maximum
Egress Remarking     : from Sap Qos
Credit Control Pol. : N/A
Category Map         : (Not Specified)
Use ing L2TP DSCP    : false
Hs-Agg-Rate-Limit    : 70000
Egress HS Q stat mode: no-override
Hs-Oper-Rate-Limit   : 70000
Egr hqos mgmt status : disabled
-----
IP Address
MAC Address          Session          Origin          Svc          Fwd
-----
10.10.1.2            00:00:64:01:01:02    N/A             DHCP          3            Y
-----
SLA Profile Instance statistics
-----
Packets              Octets
Off. HiPrio          : 0                0
Off. LowPrio         : 0                0
Off. Uncolor         : 0                0
Off. Managed         : 0                0

Queueing Stats (Ingress QoS Policy 1)
Dro. HiPrio          : 0                0
Dro. LowPrio         : 0                0
For. InProf          : 0                0
For. OutProf         : 0                0

Queueing Stats (Egress QoS Policy 20)
Dro. In/InplusProf   : 0                0
Dro. Out/ExcProf     : 118692005        118692005000
For. In/InplusProf   : 0                0
For. Out/ExcProf     : 74182782         74182782000

-----
SLA Profile Instance per Queue statistics
-----
Packets              Octets

```

```
Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio      : 0
Off. LowPrio     : 0
Dro. HiPrio      : 0
Dro. LowPrio     : 0
For. InProf      : 0
For. OutProf     : 0

Egress Queue 1
Dro. In/InplusProf : 0
Dro. Out/ExcProf   : 48844432
For. In/InplusProf : 0
For. Out/ExcProf   : 3083637

Egress Queue 2
Dro. In/InplusProf : 0
Dro. Out/ExcProf   : 39593533
For. In/InplusProf : 0
For. Out/ExcProf   : 12334536

Egress Queue 3
Dro. In/InplusProf : 0
Dro. Out/ExcProf   : 0
For. In/InplusProf : 0
For. Out/ExcProf   : 29673215

Egress Queue 4
Dro. In/InplusProf : 0
Dro. Out/ExcProf   : 14981730
For. In/InplusProf : 0
For. Out/ExcProf   : 7272849

Egress Queue 5
Dro. In/InplusProf : 0
Dro. Out/ExcProf   : 7708883
For. In/InplusProf : 0
For. Out/ExcProf   : 14545696

Egress Queue 6
Dro. In/InplusProf : 0
Dro. Out/ExcProf   : 7563427
For. In/InplusProf : 0
For. Out/ExcProf   : 7272849
```

To reveal the subscriber hierarchy in a terse form with respect to the sub/SLA-profiles and the SAP, the following command can be run:

```
*A:PE-1# show service active-subscribers hierarchy

=====
Active Subscribers Hierarchy
=====
-- sub-1 (sub-prof-1)
|
|-- sap:[3/1/1:1.1] - sla:sla-1-1
|
|   |-- 10.10.1.1 - mac:00:00:64:01:01:01 - DHCP - svc:3
|
|-- sap:[3/1/1:1.1] - sla:sla-1-2
|
|   |-- 10.10.1.2 - mac:00:00:64:01:01:02 - DHCP - svc:3
```

```
-- sub-2 (sub-prof-2)
|
|-- sap:[3/1/1:1.2] - sla:sla-2-1
|
|   |-- 10.10.1.3 - mac:00:00:64:01:01:03 - DHCP - svc:3
|
|-- sap:[3/1/1:1.2] - sla:sla-2-2
|
|   |-- 10.10.1.4 - mac:00:00:64:01:01:04 - DHCP - svc:3
|
-----
Number of active subscribers : 2
Flags: (N) = the host or the managed route is in non-forwarding state
=====
```

The following SAP related **show** command confirms that the SAP queues are removed from the underlying subscriber SAP (under the **stats** section at the end of the output). This was ensured by configuring the **profiled-traffic-only** command in the MSAP policy, with the purpose of reducing the queue consumption on ingress and egress.

```
*A:PE-1# show service id 3 sap 3/1/1:1.1 detail

=====
Service Access Points(SAP)
=====
Service Id      : 3
SAP             : 3/1/1:1.1          Encap           : qinq
QinQ Dot1p     : Default
Description     : Managed SAP - Capture Svc 10 3/1/1:*.
Admin State    : Up                  Oper State      : Up
Flags          : None
Multi Svc Site : None
Last Status Change : 09/22/2017 17:17:22
Last Mgmt Change  : 09/22/2017 17:21:50
Sub Type       : managed
Capture Service Id : 10              Capture SAP    : 3/1/1:*.
MSAP Policy    : msaps
Idle           : no                  Sticky          : no
Dot1Q Ethertype : 0x8100             QinQ Ethertype  : 0x8100
Split Horizon Group: (Not Specified)

Admin MTU       : 1522                Oper MTU        : 1522
Ingr IP Fltr-Id : n/a                 Egr IP Fltr-Id  : n/a
Ingr Mac Fltr-Id : n/a                 Egr Mac Fltr-Id  : n/a
Ingr IPv6 Fltr-Id : n/a                 Egr IPv6 Fltr-Id : n/a
qinq-pbit-marking : both

Q Frame-Based Acct : Disabled          Egr Agg Rate Limit: max
Limit Unused BW    : Disabled

Acct. Pol         : None               Collect Stats     : Disabled

Anti Spoofing : Ip-Mac             Dynamic Hosts     : Enabled
Avl Static Hosts  : 0                 Tot Static Hosts  : 0
Calling-Station-Id : n/a

Application Profile: None
Transit Policy    : None
AARP Id          : None

Oper Group        : (none)             Monitor Oper Grp  : (none)
Host Lockout Plcy : n/a
Lag Link Map Prof : (none)
Bandwidth         : Not-Applicable
```

```

Oper DCpu Prot Pol*: _default-access-policy

-----
ETH-CFM SAP specifics
-----
Tunnel Faults      : accept                AIS                : Disabled
MC Prop-Hold-Timer : n/a
Squelch Levels     : None
Collect Lmm Stats  : Disabled
LMM FC Stats       : None
LMM FC In Prof     : None

-----
QOS
-----
Ingress qos-policy : 1                      Egress qos-policy : 1
Ingress FP QGrp    : (none)                  Egress Port QGrp  : (none)
Ing FP QGrp Inst   : (none)                  Egr Port QGrp Inst: (none)
Shared Q plcy      : n/a                     Multipoint shared  : Disabled
I. Sched Pol       : (Not Specified)
E. Sched Pol       : (Not Specified)
I. Policer Ctl Pol : (Not Specified)
E. Policer Ctl Pol : (Not Specified)
E. HS Sec. Shaper  : (Not Specified)
I. QGrp Redir. List: (Not Specified)
E. QGrp Redir. List: (Not Specified)

-----
Subscriber Management
-----
Admin State        : Up                      MAC DA Hashing     : False
Def Sub-Id         : Use auto-sub-id
Def Sub-Profile    : None
Def SLA-Profile    : None
Def Inter-Dest-Id : (Use top-q-tag)
Def App-Profile    : None
Sub-Ident-Policy   : sub_ident_pol

Subscriber Limit : 1
Single-Sub-Parameters
Prof Traffic Only : True
Non-Sub-Traffic    : N/A

Static host management
MAC learn options  : N/A

-----
Sap Statistics
-----
Last Cleared Time  : N/A

CPM Ingress        : 5                      Packets            : 1077
Octets

Forwarding Engine Stats
Dropped            : 0                      0
Received Valid     : 0                      0
Off. HiPrio        : 0                      0
Off. LowPrio       : 0                      0
Off. Uncolor       : 0                      0
Off. Managed       : 0                      0

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio        : 0                      0
Dro. LowPrio       : 0                      0

```



```

For. InProf      : 0          0
For. OutProf     : 0          0

Queueing Stats(Egress QoS Policy 1)
Dro. In/InplusProf : 0          0
Dro. Out/ExcProf   : 0          0
For. In/InplusProf : 0          0
For. Out/ExcProf   : 0          0
-----
Sap per Queue stats
-----
                                Packets          Octets
No entries found
=====
* indicates that the corresponding row element may have been truncated.

```

QoS-related Show Commands in ESM Context

Examination of the subscriber QoS hierarchy on HSQ can start at the subscriber (HS queue group and SAP) level and gradually move through the HS secondary shaper, and finally the port level. The output of the **show** commands should confirm that the subscriber is associated with the QoS object as intended by the configuration.

For example, the following output shows that HS attachment policy "hs-attach-1-1" is associated with the QoS SAP egress policy 10. It was determined previously that QoS SAP egress policy 10 is associated with host 1 of the subscriber 1 ("sub-1-1"). In this case, a two-step process was necessary to track the association between the subscriber and the HS attachment policy.

```

*A:PE-1# show qos sap-egress 10 association

=====
QoS Sap Egress
=====

-----
Sap Egress Policy (10)
-----
Policy-id          : 10          Scope          : Template
Ethernet-ctag     : False       Parent-loc   : default
Name              : (Not Specified)
Description       : (Not Specified)
Policy Active     : True        Plcrs HQoS Managed : False
Post Plcr Mapping Policy: (Not Specified)
HS Attachment Policy : hs-attach-1-1

-----
Dynamic Configuration Information
-----
PccRule Insert Point : n/a          DynPlcr Insert Point : n/a
CBS                  : Def          MBS                  : Def
Parent              : (Not Specified)
Level               : 1             Weight                : 1
Packet Byte Offset  : 0
Stat Mode           : minimal

-----

Associations
-----
SLA Profiles :

```

```
- sla-1-1

-----
HSM DA Associations
-----
No Associations Found.

=====
```

The output from the HS attachment policy reflects the QoS configuration state at the subscriber level that is shown in Figure 2:

- Queues 1 and 2 are attached to WRR 1.
- Queues 3 to 6 are directly attached to the corresponding scheduling classes (3 to 6).
- Queues 7 and 8 are unattached.
- WRR 2 is unattached.
- HS attachment policy "hs-attach-1" is associated with QoS SAP egress policies 10, 20, 30, and 40 that correspond to subscriber hosts "sub-1-1", "sub-1-2", "sub-2-1", and "sub-2-2".

```
*A:PE-1# show qos hs-attachment-policy "hs-attach-1-1" detail

=====
HS Attachment Policy Information
=====
Policy Name      : hs-attach-1-1
Description      : (Not Specified)
Low Burst Max Class : 3

-----
Queue            Scheduling Class      WRR Group
-----
1                (Not-Applicable)        1
2                (Not-Applicable)        1
3                3                (Not-Applicable)
4                4                (Not-Applicable)
5                5                (Not-Applicable)
6                6                (Not-Applicable)
7                unattached        unattached
8                unattached        unattached

-----
WRR Group        Scheduling Class
-----
1                1
2                unattached

-----
Associations
-----
Network-Queue Policy
-----
No Matching Entries

Sap-Egress Policy
-----
10
20
30
40
```

```
Egress Queue-Group Templates
-----
No Matching Entries
```

Association between the HS secondary shaper and the subscribers can be verified with the following command. HS secondary shaper is allocated per port (or LAG). The two subscribers sub-1 and sub-2 are instantiated on SAPs with the outer VLAN tag "1" and consequently they are both associated with the HS secondary shaper 1. The HS secondary shaper 1 is rate limited to 120Mb/s while its scheduling classes are left open (max rate).

```
*A:PE-1# show port 3/1/1 hs-secondary-shaper "1" associations

=====
Ethernet Port 3/1/1 Egress HS Secondary Shaper Information
=====
Policy Name      : 1
Description      : (Not Specified)
Rate             : 250000 Kbps
Low Burst Max Class: 3

-----
Class              Rate
-----
1                  max
2                  max
3                  max
4                  max
5                  max
6                  max
-----

-----
Service Associations
-----
Service ID        Service Type        SAP
-----
No Service Associations Found.
-----

-----
Subscriber Associations
-----
Subscriber ID
-----
sub-1
sub-2
-----
Number of subscriber associations : 2
```

The port scheduler information can be obtained with the following command. This command is run in the QoS context (as opposed to being run in the port context, which was the case for HS secondary schedulers). The reason for this is that the HS secondary scheduler is configured directly under the port, while the HS port scheduler is configured in an HS scheduler policy (in QoS context), which is then applied to a port.

```
*A:PE-1# show qos hs-scheduler-policy "hs1" detail

=====
HS Scheduler Policy Information
=====
```

Policy Name	:	hs1		
Description	:	(Not Specified)		
Max Rate	:	300 Mbps		

Scheduling Class	Rate	Group	Weight in Group	

1	max	0	1	
2	max	0	1	
3	max	0	1	
4	max	1	10	
5	max	1	20	
6	20 Mbps	0	1	

Group	Rate			

1	60 Mbps			

Port Ethernet Egress Associations				

3/1/1				

Show Commands with Dynamically Changing Information

The following commands show the QoS hierarchy with the running rates of the objects (queues, WRR groups, scheduling classes, secondary shapers, and port shapers) and the queue buffer depths in the QoS hierarchy. The command output in this section is based on the scenario described in the [Traffic Management on HSQ](#) section.



Note:

Running rates are the dynamically calculated rates while traffic is running.

The following subscriber hierarchy (for sub-1) shows the rates per subscriber host, per scheduling priority (and consequently, scheduling class), starting at the queue level and moving up toward the port level. For example, scheduling priority 1 starts with the rates of two subscriber host queues (1 and 2) that are mapped to the WRR group and then it moves up to the rate at the HS secondary shaper level (the summed rate of all the entities at scheduling priority 1 at the HS secondary shaper level), ending with the rate of the scheduling class 1 at the port level. The aggregate rate of the HS secondary shaper and the port in the subscriber hierarchy are also provided.

In this case, the rates between the two HSQ queue groups (or subscriber hosts) are evenly spread. This distribution will be explained in the [Traffic Management on HSQ](#) section. The subscriber hosts can be differentiated by the two different SLA profiles with which they are associated.

```
*A:PE-1# show qos hs-scheduler-hierarchy subscriber "sub-1" egress

=====
Hs Scheduler Hierarchy Information
=====
PortId                : 3/1/1
SAP                   : [3/1/1:1.1]
SLA Profile           : sla-1-1
Hs Sched Policy Name  : hs1

Port Max-Rate : 250 Mbps
Hs-Sec-Shaper:1 Agg-Rate : 250479 Kbps
```

```

Scheduler Priority 6
  Scheduler Class 6 Rate : 20 Mbps
    Hs-Sec-Shaper:1 Class 6 Rate : 20040 Kbps
      Queue : 6 Rate : 4912 Kbps

Scheduler Priority 5 Group 1
  Scheduler Class 5 Rate : 40 Mbps Weight : 20
    Hs-Sec-Shaper:1 Class 5 Rate : 40098 Kbps
      Queue : 5 Rate : 9824 Kbps
  Scheduler Class 4 Rate : 20 Mbps Weight : 10
    Hs-Sec-Shaper:1 Class 4 Rate : 20024 Kbps
      Queue : 4 Rate : 4912 Kbps

Scheduler Priority 3
  Scheduler Class 3 Rate : 81 Mbps
    Hs-Sec-Shaper:1 Class 3 Rate : 81746 Kbps
      Queue : 3 Rate : 20032 Kbps

Scheduler Priority 2
  Scheduler Class 2 Rate : 0 Mbps
    Hs-Sec-Shaper:1 Class 2 Rate : 0 Kbps

Scheduler Priority 1
  Scheduler Class 1 Rate : 88 Mbps
    Hs-Sec-Shaper:1 Class 1 Rate : 88568 Kbps
      Queue : 1 Group : 1 Rate : 2080 Kbps
      Queue : 2 Group : 1 Rate : 8336 Kbps
=====
PortId : 3/1/1
SAP : [3/1/1:1.1]
SLA Profile : sla-1-2
Hs Sched Policy Name : hs1

Port Max-Rate : 250 Mbps
Hs-Sec-Shaper:1 Agg-Rate : 250414 Kbps

Scheduler Priority 6
  Scheduler Class 6 Rate : 20 Mbps
    Hs-Sec-Shaper:1 Class 6 Rate : 20040 Kbps
      Queue : 6 Rate : 4912 Kbps

Scheduler Priority 5 Group 1
  Scheduler Class 5 Rate : 40 Mbps Weight : 20
    Hs-Sec-Shaper:1 Class 5 Rate : 40065 Kbps
      Queue : 5 Rate : 9808 Kbps
  Scheduler Class 4 Rate : 20 Mbps Weight : 10
    Hs-Sec-Shaper:1 Class 4 Rate : 20040 Kbps
      Queue : 4 Rate : 4912 Kbps

Scheduler Priority 3
  Scheduler Class 3 Rate : 81 Mbps
    Hs-Sec-Shaper:1 Class 3 Rate : 81714 Kbps
      Queue : 3 Rate : 20040 Kbps

Scheduler Priority 2
  Scheduler Class 2 Rate : 0 Mbps
    Hs-Sec-Shaper:1 Class 2 Rate : 0 Kbps

Scheduler Priority 1
  Scheduler Class 1 Rate : 88 Mbps
    Hs-Sec-Shaper:1 Class 1 Rate : 88552 Kbps
      Queue : 1 Group : 1 Rate : 2080 Kbps
      Queue : 2 Group : 1 Rate : 8336 Kbps

```

The following command provides the running rates at the HS secondary shaper and port levels:

```
*A:PE-1# show qos hs-scheduler-hierarchy port 3/1/1 hs-secondary-shapers
```

```
=====
Hs Scheduler Hierarchy Information
=====
```

```
Hs Sched Policy Name      : hs1
```

```
Port Max-Rate : 250 Mbps
```

```
Scheduler Priority 6
  Scheduler Class 6  Rate : 20 Mbps
```

```
Scheduler Priority 5  Group 1
  Scheduler Class 5  Rate : 40 Mbps      Weight : 20
  Scheduler Class 4  Rate : 20 Mbps      Weight : 10
```

```
Scheduler Priority 3
  Scheduler Class 3  Rate : 81 Mbps
```

```
Scheduler Priority 2
  Scheduler Class 2  Rate : 0 Mbps
```

```
Scheduler Priority 1
  Scheduler Class 1  Rate : 88 Mbps
```

```
-----
HS Secondary Shaper Rates
-----
```

```
Hs-Sec-Shaper:1 Agg-Rate : 249043 Kbps
  Class 6      Rate : 19926 Kbps
  Class 5      Rate : 39853 Kbps
  Class 4      Rate : 19926 Kbps
  Class 3      Rate : 81273 Kbps
  Class 2      Rate : 0 Kbps
  Class 1      Rate : 88062 Kbps
```

```
Hs-Sec-Shaper:default Agg-Rate : 0 Kbps
  Class 6      Rate : 0 Kbps
  Class 5      Rate : 0 Kbps
  Class 4      Rate : 0 Kbps
  Class 3      Rate : 0 Kbps
  Class 2      Rate : 0 Kbps
  Class 1      Rate : 0 Kbps
-----
```

Another important parameter to monitor is the depth of the queues. This provides information about the congestion at the queue level. The buffer space per queue is allocated automatically by the system and, in the following case, the buffers are rather large. Deep buffering causes longer delays. To avoid this, the queue buffers can be adjusted by the **mbs** command under the queue definition in the QoS SAP egress policy. For example, 15kbytes would accommodate roughly fifteen 1000byte packets in a buffer queue.

```
*A:PE-1>config>qos>sap-egress# info
```

```
-----
queue 1 create
  mbs 15 kilobytes
```

The following output shows that, except for queues 3, all queues have their buffers fully used. Because there is no congestion on queues 3, their buffer depth is low.

```
*A:PE-1# show hs-pools port 3/1/1 egress subscriber "sub-1"
| match "Queue Information" pre-lines 1 post-lines 200
-----
Queue Information
-----
Queue Name      : Sub=sub-1:sla-1-1 3->3/1/1:1.1->1
FC Map          : be ef nc
Admin PIR       : 20000                      Oper PIR           : 0
Admin MBS       : 64 KB                      Oper MBS           : 64 KB
HS Wrr Group    : 1
HS Wrr Class Weight: 1                      HS Wrr Weight      : 1
Depth          : 58 KB
HS Class        : 1                      HS Alt Port Class Pool : No
HS Slope Policy : _tmnx_hs_default

Queue Name      : Sub=sub-1:sla-1-1 3->3/1/1:1.1->2
FC Map          : l2
Admin PIR       : 20000                      Oper PIR           : 0
Admin MBS       : 64 KB                      Oper MBS           : 64 KB
HS Wrr Group    : 1
HS Wrr Class Weight: 1                      HS Wrr Weight      : 4
Depth          : 58 KB
HS Class        : 1                      HS Alt Port Class Pool : No
HS Slope Policy : _tmnx_hs_default

Queue Name      : Sub=sub-1:sla-1-1 3->3/1/1:1.1->3
FC Map          : af
Admin PIR       : Max                      Oper PIR           : Max
Admin MBS       : 375000 B                  Oper MBS           : 375296 B
HS Wrr Group    : (not-applicable)
HS Wrr Class Weight: 1                      HS Wrr Weight      : 0
Depth          : 1 KB
HS Class        : 3                      HS Alt Port Class Pool : No
HS Slope Policy : _tmnx_hs_default

Queue Name      : Sub=sub-1:sla-1-1 3->3/1/1:1.1->4
FC Map          : l1
Admin PIR       : Max                      Oper PIR           : Max
Admin MBS       : 375000 B                  Oper MBS           : 375296 B
HS Wrr Group    : (not-applicable)
HS Wrr Class Weight: 1                      HS Wrr Weight      : 0
Depth          : 333 KB
HS Class        : 4                      HS Alt Port Class Pool : No
HS Slope Policy : _tmnx_hs_default

Queue Name      : Sub=sub-1:sla-1-1 3->3/1/1:1.1->5
FC Map          : h1
Admin PIR       : Max                      Oper PIR           : Max
Admin MBS       : 375000 B                  Oper MBS           : 375296 B
HS Wrr Group    : (not-applicable)
HS Wrr Class Weight: 1                      HS Wrr Weight      : 0
Depth          : 333 KB
HS Class        : 5                      HS Alt Port Class Pool : No
HS Slope Policy : _tmnx_hs_default

Queue Name      : Sub=sub-1:sla-1-1 3->3/1/1:1.1->6
FC Map          : h2
Admin PIR       : Max                      Oper PIR           : Max
Admin MBS       : 375000 B                  Oper MBS           : 375296 B
HS Wrr Group    : (not-applicable)
```

```

HS Wrr Class Weight: 1          HS Wrr Weight          : 0
Depth                : 332 KB
HS Class             : 6          HS Alt Port Class Pool : No
HS Slope Policy       : _tmnx_hs_default

Queue Name           : Sub=sub-1:sla-1-2 3->3/1/1:1.1->1
FC Map               : be ef nc
Admin PIR             : 40000      Oper PIR              : 0
Admin MBS             : 64 KB      Oper MBS             : 64 KB
HS Wrr Group          : 1
HS Wrr Class Weight: 1          HS Wrr Weight          : 1
Depth                : 56 KB
HS Class             : 1          HS Alt Port Class Pool : No
HS Slope Policy       : _tmnx_hs_default

Queue Name           : Sub=sub-1:sla-1-2 3->3/1/1:1.1->2
FC Map               : l2
Admin PIR             : 40000      Oper PIR              : 0
Admin MBS             : 64 KB      Oper MBS             : 64 KB
HS Wrr Group          : 1
HS Wrr Class Weight: 1          HS Wrr Weight          : 4
Depth                : 58 KB
HS Class             : 1          HS Alt Port Class Pool : No
HS Slope Policy       : _tmnx_hs_default

Queue Name           : Sub=sub-1:sla-1-2 3->3/1/1:1.1->3
FC Map               : af
Admin PIR             : Max        Oper PIR              : Max
Admin MBS             : 375000 B   Oper MBS             : 375296 B
HS Wrr Group          : (not-applicable)
HS Wrr Class Weight: 1          HS Wrr Weight          : 0
Depth                : 1 KB
HS Class             : 3          HS Alt Port Class Pool : No
HS Slope Policy       : _tmnx_hs_default

Queue Name           : Sub=sub-1:sla-1-2 3->3/1/1:1.1->4
FC Map               : l1
Admin PIR             : Max        Oper PIR              : Max
Admin MBS             : 375000 B   Oper MBS             : 375296 B
HS Wrr Group          : (not-applicable)
HS Wrr Class Weight: 1          HS Wrr Weight          : 0
Depth                : 333 KB
HS Class             : 4          HS Alt Port Class Pool : No
HS Slope Policy       : _tmnx_hs_default

Queue Name           : Sub=sub-1:sla-1-2 3->3/1/1:1.1->5
FC Map               : h1
Admin PIR             : Max        Oper PIR              : Max
Admin MBS             : 375000 B   Oper MBS             : 375296 B
HS Wrr Group          : (not-applicable)
HS Wrr Class Weight: 1          HS Wrr Weight          : 0
Depth                : 332 KB
HS Class             : 5          HS Alt Port Class Pool : No
HS Slope Policy       : _tmnx_hs_default

Queue Name           : Sub=sub-1:sla-1-2 3->3/1/1:1.1->6
FC Map               : h2
Admin PIR             : Max        Oper PIR              : Max
Admin MBS             : 375000 B   Oper MBS             : 375296 B
HS Wrr Group          : (not-applicable)
HS Wrr Class Weight: 1          HS Wrr Weight          : 0
Depth                : 333 KB
HS Class             : 6          HS Alt Port Class Pool : No
HS Slope Policy       : _tmnx_hs_default

```


Clear Commands

Clear commands in the HSQ context are used to clear statistics associated with the HS secondary shaper:

```
clear port <port-id> hs-secondary-shaper <name> statistics
```

Resources Monitoring

The **tools dump system-resources** command is used to monitor resource on an HSQ IOM.

Some of the resources (HSQ queue groups, HS primary shapers, HS secondary shapers, HS turbo queue groups, and so on) are allocated for internal use, thereby reducing the number of resources in the **Free** column. Such internally consumed resources are not available to be part of the user configuration.



Note:

100G ports support port-based (access or network) HS queue groups on egress that can be configured to support higher throughput rates. Such port-based high-rate HS queue groups on egress are referred to as egress HS turbo queue groups.

The number of internally consumed resources depends on the configuration and MDA types. In the following example, the number of allocated HS queue groups is 40. Each of the four subscriber hosts consume one HS queue group, which means that 36 HS queue groups are internally allocated.

Out of 23 allocated HS primary shapers, two are allocated for the two subscribers and one is allocated to fulfill the HSQ hierarchy chain in conjunction with the default HS secondary shaper. This leaves 20 HS primary shapers consumed internally by the HSQ IOM.

Similarly, out of 22 allocated HS secondary shapers in total, only two are user related (and visible via **show** commands): a default HS secondary shaper and the HS secondary shaper 1. This means that 20 HS secondary shapers are internally consumed. The same logic can be followed for turbo HS queue groups, which are out of the scope in this chapter.



Note:

On HSQ IOM, a complete scheduling hierarchy must be maintained between each attached queue at the SPI level and its port priority. That is, the HSQ has no provision for bypassing the primary and secondary shaping/scheduling levels.

```
*A:PE-1# tools dump system-resources 3
Resource Manager info at 004 h 09/22/17 22:09:11.322:

Hardware Resource Usage for Slot #3, CardType iom4-e-hs, Cmplx #0:
-----|-----|-----|-----
          | Total | Allocated | Free
-----|-----|-----|-----
SAP Ingress QoS Policies | 1791 | 1 | 1790
Dynamic Egress Classification + | 2047 | 5 | 2042
SAP Egress QoS Policies - | 5 | 0 |
Network Egress Classification - | 0 | 0 |
Ingress Queues | 131072 | 502 | 130570
Egress Queues | 786432 | 143 | 786289
Egress HS Turbo Queue Groups | 64 | 10 | 54
Egress HS Queue Groups | 98240 | 40 | 98200
```

Primary Shapers +	16384	23	16361
Explicit Primary Shapers -		2	
Managed Primary Shapers -		21	
Secondary Shapers	4096	22	4074
Ingress Policers	511999	1	511998
Egress Policers	262143	1	262142
Ingress Policer Stats	511967	0	511967
Egress Policer Stats	262111	0	262111
Qos Ingress Root Arbiters	65535	1	65534
Qos Egress Root Arbiters	65535	1	65534
Qos Intermediate Arbiters	262143	0	262143
Egress QoS Bypass	131071	0	131071
Ingress ACL Entries	65536	2	65534
Ingress QoS Entries	16384	2	16382
Ingress IPv6 ACL Entries	28672	2	28670
Ingress IPv6 QoS Entries	4096	2	4094
Egress ACL Entries	32768	2	32766
Egress QoS Entries	14336	2	14334
Egress IPv6 ACL Entries	16384	2	16382
Egress IPv6 QoS Entries	2048	2	2046
Ingress ACL Filters	2047	0	2047
Ingress IPv6 ACL Filters	2047	0	2047
Egress ACL Filters	2047	0	2047
Egress IPv6 ACL Filters	2047	0	2047
QoS User Schedulers	98303	0	98303
QoS User Scheduler Overrides	196607	0	196607
Sap IngQGrp RedirLst Entries	31999	0	31999
Dynamic Service Entries +	131071	4	131067
Subscriber Hosts -	131071	4	131067
Encap Group Members -	65535	0	65535
Egr Network Queue Group Mappings -	131071	0	131071
SapInst EgrQGrp RedirLst Entries -	31999	0	31999
Dynamic Nexthop Entries +	511999	4	511995
Subscriber Nexthops -	511999	4	511995
Ipssec tunnels -	511999	0	511999
Subscriber SPI QoS Overrides	131072	0	131072
Mac Fdb Entries	511999	0	511999
Egress TLS Mcast Entries	368639	1	368638

Traffic Management on HSQ

This section examines traffic output on HSQ IOM during congestion. [Figure 86: Managing Congestion on HSQ in Expanded SLA Mode](#) is a graphic representation of the configuration described previously, but with traffic streams running through the HSQ IOM.

Six traffic streams are sent in the downstream direction toward each of the four subscriber hosts (host 1-1, host 1-2, host 2-1, and host 2-2): in total, there are 24 traffic streams. The traffic streams are shown on the left side of [Figure 86: Managing Congestion on HSQ in Expanded SLA Mode](#), with their names, offered rates (IN column) and measured output rates (OUT column). The traffic streams are sent and analyzed by the traffic generator.

The four shaded squares in the center of [Figure 86: Managing Congestion on HSQ in Expanded SLA Mode](#) represent the four subscriber hosts and their scheduling classes. The QoS hierarchy is shown on the right side. Red shaded areas (shapers) represent points of congestion on HSQ IOM caused by the 24 traffic streams.

Each of the six traffic streams per subscriber host is fed into the six queues of each subscriber host (or HSQ queue groups). The first digit in the traffic stream name represents the subscriber, the second digit

represents the subscriber host, and the third digit represents the queue to which this stream is sent. For example, STRM 1-2-3 represents a traffic stream sent to queue 3 of the second host for subscriber 1.

To summarize the scenario shown in [Figure 86: Managing Congestion on HSQ in Expanded SLA Mode](#):

- Traffic streams 1 and 2 of each subscriber host are mapped to subscriber queues 1 and 2, which are in turn associated with WRR group 1 at the subscriber host (HSQ queue group) level.
- WRR group 1 is, depending on the subscriber host, rate limited to 20Mb/s, 40Mb/s, 60Mb/s, and 80 Mb/s. Weight ratio between queues 2 and 1 is 4:1. WRR group 1 is then attached to scheduling class 1.
- Traffic stream 3 is via queue 3 directly mapped to scheduling class 3.
- Scheduling class 2 is unused in this example.
- Traffic streams 4 and 5 are via queues 4 and 5 mapped to scheduling classes 4 and 5, which are at the port level collapsed into WRR group 1 with an aggregate rate limit of 60Mb/s.
- Traffic stream 6 is the highest priority stream that is via queue 6 mapped to scheduling class 6. At the port level, scheduling class 6 is rate limited to 20Mb/s.
- Subscriber host (HSQ queue group) aggregate rates are set to 50Mb/s, 70Mbps, 80 Mb/s, and 100Mb/s, respectively.
- Subscribers (sub-1 and sub-2) aggregate rates (HS primary shapers) are set to 100Mb/s and 150 Mb/s, respectively.
- Subscribers sub-1 and sub-2 are mapped to HS secondary shaper "1" (via outer VLAN on their SAPs) with the aggregate rate of 250 Mb/s.
- HS port scheduler is rate limited to 300Mb/s.

The configured rate limits at the subscriber level are L2 rates, while configured rate limits at the HS secondary shaper level and the HS port level include L1 overhead and are, therefore, on-the-wire rates. On-the-wire rates account for 20 additional bytes in each Ethernet frame (8 bytes preamble and 12 bytes IFG).

All traffic streams are sent with constant rates (no added burstiness) and fixed packet size (1000 bytes). Therefore, the difference between the L2 rates and the on-the-wire rates for the 1000byte packets is $1000/1020 = 0.98$ or 2%. That is, on-the-wire rates are 2% higher than the L2 rates. The name for this 2% delta factor in this chapter will be the Rate Conversion Factor (RCF).

Figure 86: Managing Congestion on HSQ in Expanded SLA Mode

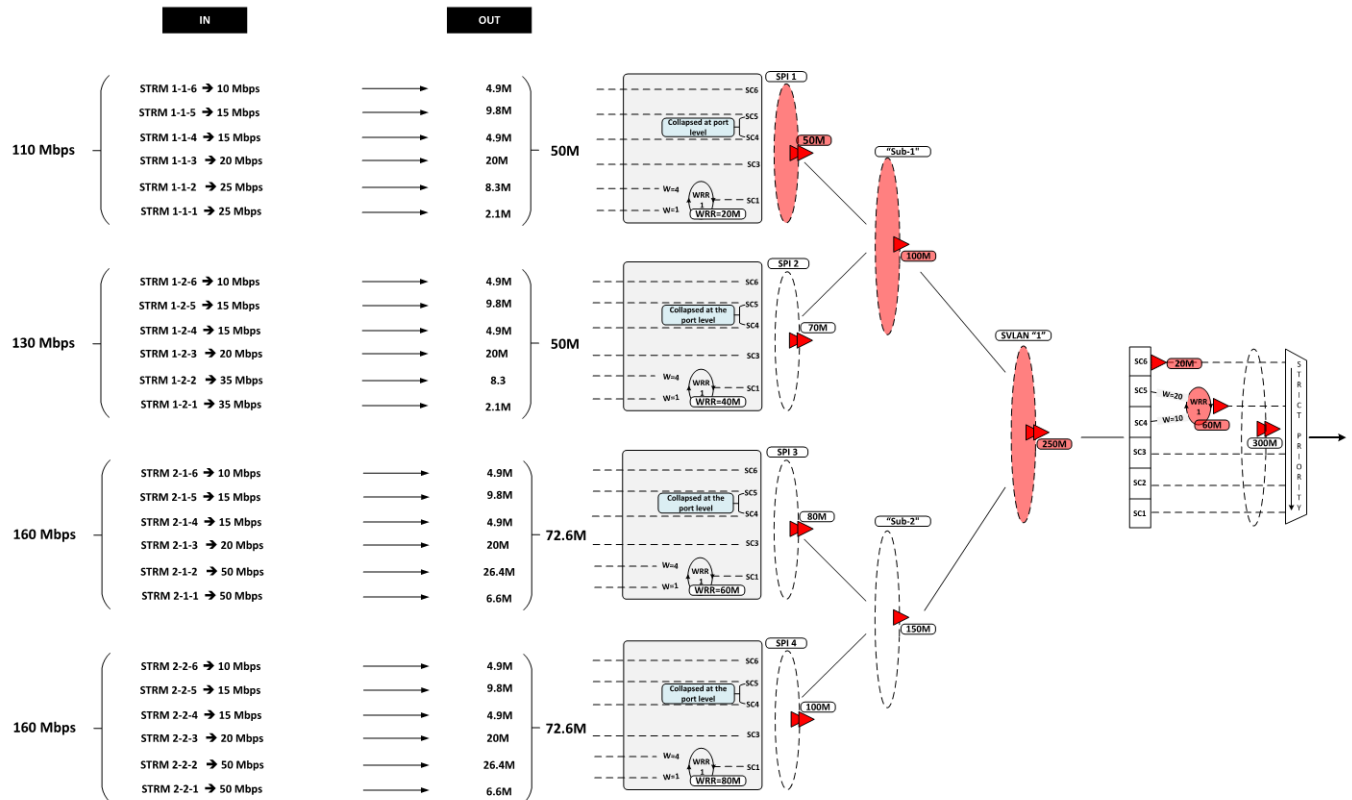


Table 22: Input/Output Rates lists the input/output rates throughout the subscriber QoS hierarchy. Red shaded table cells represent congested objects. The numbers in blue represent the L2 output rates measured on the traffic generator. The numbers in red (in parentheses) are configured aggregate rates for the object in the hierarchy. The numbers above them are either the operational rates measured by SR OS and observed via the `show qos hs-scheduler-hierarchy` command, or the manually summed rates under the hierarchy.



Note:

Running aggregate rates per HSQ queue group or subscriber (HS primary shaper) are now available via **show** commands.

Some of the output columns in this table require additional explanation:

- Per stream output rates (first column under the Output Rate section) are L2 rates as measured by the traffic generator. The analysis of the results will be based on these rates.

However, those rates can also be displayed within the system via the **show qos hs-scheduler-hierarchy subscriber <sub-name> egress** command. RCF must be used to correctly interpret the results in respective rate domains (L1 versus L2).

- WRR rates at the subscriber level (second column under the Output Rate section) are populated by manually adding rates measured by the traffic generator for traffic streams 1 and 2. The 4:1 notation in parentheses represents the weight ratio between the two streams. There are no means to observe the WRR group rates at the subscriber level directly within the system (via **show** commands).

- Aggregate rates per host (third column under the Output Rate section) are populated by manually adding the rates measured by the traffic generator from all six traffic streams for each subscriber. There are no means to observe the aggregate subscriber rate directly within the system (via show commands). Those rates also represent the HSQ queue group aggregate rate because each host in this example is associated with an HSQ queue group.
- Aggregate rates per subscriber (fourth column under the Output Rate section) are populated by manually adding the rates measured by the traffic generator from all six traffic streams for each subscriber. There are no means to observe the aggregate subscriber rate directly within the system (via show commands).
- Aggregate rates per HS secondary shaper (fifth column under the Output Rate section) are on-the-wire rates displayed via the **show qos hs-scheduler-hierarchy port 3/1/1 hs-secondary-shapers** command. The displayed rates are rounded to the nearest Mb/s.
- Per scheduling class rates at the port level (sixth column under the Output Rate section) are displayed via the **show qos hs-scheduler-hierarchy port 3/1/1** command. The displayed rates are rounded by the system to the nearest Mb/s. The discrepancy between the actual rates for scheduling class 3 (80Mb/s) and the displayed rate (81Mb/s) is due to the measuring and rounding inaccuracy at display time. The actual rate at scheduling class 3 at the port level can be calculated by summing the rates measured (by the traffic generator) of each traffic stream 3 for each subscriber and adjusting the sum for the on-the-wire rate (RCF).
The same logic applies to scheduling class 1 (actual rate 86.8Mb/s vs 88Mb/s measured rate).
- The WRR-1 rate at the port level (seventh column under the Output Rate section) is manually calculated by adding the rates measured by the traffic generator of all streams mapped to scheduling classes 5 and 4 (8 traffic streams in total, two per each subscriber host), then converting this rate into the on-the-wire rate. At the port level, traffic on scheduling classes 5 and 4 is weighted in a 2:1 ratio (this is noted in parentheses).
- The port rate (last column under the Output Rate section) is collected from two places:
 - The top number is obtained via the **show qos hs-scheduler-hierarchy** command and represents the on-the-wire rate.
 - The bottom number in blue is the number measured by the traffic generator, which only measures L2 rates. On-the-wire rates can be converted to L2 rates by multiplying the on-the-wire rates by the RCF.

Table 22: Input/Output Rates

Strm	Input rate in [mbps]		Output rate in [mbps]							
	Per strm	Agg per host	Per strm (L2)	WRR on subscr level	Agg per host (L2)	Agg per subscr (L2)	Agg per secondary shaper (wire)	Per sch class on port level (wire)	WRR-1 on port level (wire)	Port
1-1-6	10	110	4.9	-	50 (50)	100 (100)	250 (250)	SC 6	-	250 (300) 245.1 (L2)
1-1-5	15		9.8	-				20	-	
1-1-4	15		4.9	-				(20)	-	

Strm	Input rate in [mbps]		Output rate in [mbps]								
	Per strm	Agg per host	Per strm (L2)	WRR on subscr level	Agg per host (L2)	Agg per subscr (L2)	Agg per secondary shaper (wire)	Per sch class on port level (wire)	WRR-1 on port level (wire)	Port	
1-1-3	20		20	-					-		
1-1-2	25		8.3	10.4				(4:1)	SC 5 40		60 (60) (2:1)
1-1-1	25		2.1								
1-2-6	10	130	4.9	-	50 (70)			SC 4 20			
1-2-5	15		9.8	-							
1-2-4	15		4.9	-							
1-2-3	20		20	-							
1-2-2	35		8.3	10.4 (4:1)							
1-2-1	35		2.1								
2-1-6	10	160	4.9	-	72.6 (80)	145.2 (150)		SC 3 81	-		
2-1-5	15		9.8	-					-		
2-1-4	15		4.9	-					-		
2-1-3	20		20	-					-		
2-1-2	50		26.4	33				SC 2 -	-		
2-1-1	50		6.6	(4:1)					-		
2-2-6	10	160	4.9	-	72.6 (100)			SC 1 88	-		
2-2-5	15		9.8	-					-		
2-2-4	15		4.9	-					-		
2-2-3	20		20	-					-		
2-2-2	50		26.4	33				-			
2-2-1	50		6.6	(4:1)				-			

Analysis of Results

The analysis of results begins with the rates per stream measured by the traffic generator. The expected behavior is that those rates agree with the theoretical rate calculations based on our understanding of QoS on HSQ IOM.

Considering the six strict priority classes in the HSQ scheduling mechanism, the expectation is that the traffic is serviced in the order of priority, from the highest scheduling class 6 to the lowest scheduling class 1. Consequently, the traffic analysis starts with the streams that are mapped to the highest priority scheduling class 6 (streams 1-1-6, 1-2-6, 2-1-6, 2-2-6, that is, stream 6 of each subscriber host).

Scheduling class 6 (streams 1-1-6, 1-2-6, 2-1-6, and 2-2-6) – the highest priority scheduling class

Due to the aggregate rate limit of 20Mb/s for scheduling class 6 at the port level, it is expected that each subscriber host receives an equal amount of traffic on scheduling class 6:

$$(20 \text{ Mb/s (aggregate rate of SC 6)}) \div (4 \text{ streams (one per subscriber host)}) = 5 \text{ Mb/s}$$

The measured results on the traffic generator for traffic streams 6 in Table 1 show that each subscriber host receives 4.9Mb/s out of an offered 10Mb/s. The slight difference between the expected (5Mb/s) and the measured (4.9Mb/s) rate is caused by the discrepancy between the L2 rates at the subscriber level (as measured by the traffic generator), and on-the-wire rates enforcement (20Mb/s) at the scheduling class 6 at the port level. Multiplying 5Mb/s by the RCF ($1000/1020$) will align the results.

This slight discrepancy between the L2 and on-the-wire rates is common throughout the remaining analysis. Because this is well understood and expected, it will not be mentioned again.

Scheduling classes 5 and 4 (streams 1-1-5, 1-1-4, 1-2-5, 1-2-4, 2-1-5, 2-1-4, 2-2-5, and 2-2-4)

Scheduling classes 5 and 4 are collapsed at the port level into WRR group 1, which is at the next scheduling priority to be served. These two scheduling classes are served by WRR group 1 in a 2:1 ratio. The aggregate rate limit for WRR 1 at the port level is set to 60Mb/s.

Expected rates for combined scheduling classes 5 and 4 of each host are

$$60\text{Mb/s} \div 4 = 15\text{Mb/s}$$

When 15Mb/s is distributed between the two scheduling classes in the 2:1 ratio, each host should receive 10Mb/s on scheduling class 5 and 5Mb/s on scheduling class 4, for the total of 15Mb/s (out of an offered 30Mb/s).

The measured results in [Table 22: Input/Output Rates](#) are 9.8Mb/s on scheduling class 5 and 4.9Mb/s on scheduling class 4 for traffic streams 4 and 5 of each individual host. This is in line with the expected results (the slight difference is due to L2 rates measured by the traffic generator and enforced on-the-wire rates at the port level).

Scheduling class 3 (streams 1-1-3, 1-2-3, 2-1-3, and 2-2-3)

Traffic streams mapped to scheduling class 3 do not have any rate restriction at the scheduling class level. Those streams can be only limited by the congestion at the HSQ queue group aggregate level, the subscriber aggregate level, the HS secondary shaper aggregate level, or the port aggregate level. Because the total amount of traffic so far is below congestion level at each point in the hierarchy, it is expected that traffic stream 3 flows unimpeded. Consequently, each subscriber host should receive the full input rate of 20Mb/s on scheduling class 3. The actual results in [Table 22: Input/Output Rates](#) are aligned with the expected results.

To confirm that there is no congestion at the aggregate level so far for the subscriber hosts, the subscribers, the HS secondary shapers, and the port, a calculation shows that each subscriber host has received an equal amount of bandwidth so far: 39.6Mb/s. This is below the configured aggregate rate limits at the host, the subscriber, the HS secondary scheduler and the port levels:

- 39.6Mb/s is below the configured limit of 50Mb/s for host 1-1.
- 39.6Mb/s is below the configured limit of 70Mb/s for host 1-2.
- 39.6Mb/s is below the configured limit of 80Mb/s for host 2-1.
- 39.6Mb/s is below the configured limit of 100Mb/s for host 2-2.
- 79.2Mb/s ($39.6\text{Mb/s} \times 2$) is below the configured limit of 100Mb/s for sub-1.
- 79.2Mb/s ($39.6\text{Mb/s} \times 2$) is below the configured limit of 150Mb/s for sub-2.
- Sub-1 and sub-2 compete for the bandwidth at HS secondary shaper "1", and their combined rate of 158.4 Mb/s ($39.6\text{Mb/s} \times 4$) is below the configured aggregate rate of the HS secondary shaper "1" (250Mb/s) or the port shaper (300Mb/s).

Scheduling classes 2 and 1 (streams 1-1-2, 1-1-1, 1-2-2, 1-2-1, 2-1-2, 2-1-1, 2-2-2, 2-2-1) – the lowest priority scheduling classes

The total amount of offered (input) traffic for lowest priority scheduling classes 1 and 2 across all four subscriber hosts is 560Mb/s (110Mb/s for host 1-1, 130Mb/s for host 1-2, 160Mb/s for host 2-1, and 160Mb/s for host 2-2). This additional amount of traffic will cause congestion on the aggregate level for host 1-1, subscriber sub-1, and HS secondary shaper 1.

Up to this point, the spare capacity on the HS secondary shaper 1 is:

250 Mb/s (HS secondary shaper 1 aggregate on-the-wire rate limit)

- 40.3 Mb/s (on-the-wire traffic from host 1-1 up to this point)
- 40.3 Mb/s (on-the-wire traffic from host 1-2 up to this point)
- 40.3 Mb/s (on-the-wire traffic from host 2-1 up to this point)
- 40.3 Mb/s (on-the-wire traffic from host 2-2 up to this point)

= 88.8 Mb/s (capacity left on the HS secondary shaper 1 up to this point)



Note:

1. 40.3Mb/s corresponds to 39.6Mb/s converted to L2 rate.

2. Spare capacity for HS secondary shaper "1" is calculated based on on-the-wire rates while the subscriber aggregate capacity is calculated in L2 rates because the aggregate rate limit in HS secondary shaper in SR OS is configured in on-the-wire rates while the subscriber aggregate rates are configured in L2 rates.

Considering that each host has so far received 40.3Mb/s (on-the-wire rate), the 88.8Mb/s left on the HS secondary shaper will be distributed between scheduling classes 2 and 1 for each of the hosts (1-1, 1-2, 2-1, and 2-2) in the following manner:

- Host 1-1, limited by its configured aggregate rate limit of 50Mb/s (L2 rate) will receive (50Mb/s – 39.6Mb/s =) 10.4Mb/s.
- The amount of bandwidth that host 1-2 receives is limited by the sub-1 configured aggregate rate limit (100Mb/s). Because host 1-1 and host 1-2 belong to the same subscriber (sub-1) and host 1-1 is limited by its own configured limit of 50Mb/s, this will leave 50Mb/s of aggregate bandwidth for host 2-1. That is, host 1-2 is not limited by its own configured aggregate rate (70Mb/s) but by that of the sub-1, which is now congested. Therefore, host 1-2 will receive (50Mb/s – 39.6Mb/s =) 10.4Mb/s of bandwidth on scheduling classes 2 and 1.
- The remaining 66.2Mb/s (L2 rate) is delegated to sub-2, and this will not be enough to congest any aggregate rate at the sub-2 level. Therefore, this bandwidth will be equally divided over host 2-1 and 2-2, each receiving 33.1Mb/s (L2 rate) on scheduling classes 2 and 1.

Considering the preceding bandwidth distribution for scheduling classes 2 and 1 for all four hosts and the 4:1 weight ratio between scheduling classes 2 and 1, the following conclusion can be reached:

- Host 1-1 → 10.4Mb/s in 4:1 ratio:
 - scheduling class 2 will receive 8.32Mb/s
 - scheduling class 1 will receive 2.08Mb/s
- Host 1-2 → 10.4Mb/s in 4:1 ratio:
 - scheduling class 2 will receive 8.32Mb/s
 - scheduling class 1 will receive 2.08Mb/s
- Host 2-1 → 33.1Mb/s in 4:1 ratio:
 - scheduling class 2 will receive 26.5Mb/s
 - scheduling class 1 will receive 6.62Mb/s
- Host 2-2 → 33.1Mb/s in 4:1 ratio:
 - scheduling class 2 will receive 26.5Mb/s
 - scheduling class 1 will receive 6.62Mb/s

These numbers match the rates measured by the traffic generator for streams <1-1-2>, <1-1-1>, <1-2-2>, <1-2-1>, <2-1-2>, <2-1-1>, <2-2-2>, <2-2-1> in [Table 22: Input/Output Rates](#).

Conclusion

This chapter described traffic management capabilities and configuration of HSQ IOM with ESM in expanded SLA mode. This chapter is an extension of the chapter [High Scale QoS IOM in ESM Context: Single SLA Mode](#). Both single and expanded SLA modes rely on the unique properties of HSQ IOM, while each mode has a unique set of characteristics:

- Single SLA mode provides higher subscriber scale per HSQ IOM, but allows only one SLA profile instance (SPI) per subscriber. This means that only a single SAP per subscriber (or residence) is supported (service per SAP model is not supported).
- In contrast, expanded SLA mode supports multiple SPIs per subscriber (and SAPs), but with the reduced subscriber scale. The reduction in subscriber scale is caused by the tie-in between the subscribers and HS primary shapers. That is, the aggregate subscriber rate is enforced through HS primary shapers, which means that the mapping between a subscriber and an HS primary shaper is 1:1. Therefore, 16k HS primary shapers on HSQ will determine the subscriber scale in expanded SLA mode.



Note:

On HSQ, a subscriber cannot be dissociated from an HS primary shaper, or from the QoS hierarchy.

Expanded SLA Mode can be used in service per SAP deployments where each subscriber is assigned to multiple SAPs, or it can be used on a single SAP where various services require that each of them is assigned its own HSQ queue group (or SPI).

HSQ IOM is a module of choice in an environment that demands high performance (200Gb/s per IOM) in combination with QoS functionality and a high number of egress queues.

High Scale QoS IOM in ESM Context: Single SLA Mode

This chapter describes the High Scale QoS (HSQ) IOM in the Enhanced Subscriber Management (ESM) context.

Topics in this chapter include:

- [Applicability](#)
- [Overview](#)
- [Configuration](#)
- [Conclusion](#)
- [Appendix A — Generic ESM configuration](#)

Applicability

This chapter describes the High Scale QoS (HSQ) IOM in the Enhanced Subscriber Management (ESM) context. The information and configuration in this chapter are based on SR OS Release 15.0.R4.

Overview

This chapter describes the QoS operation and configuration of the High Scale QoS (HSQ) IOM, with a focus on single SLA mode in the Enhanced Subscriber Management ESM context.

Single SLA mode on HSQ is characterized by a single SLA profile instance (SPI) per subscriber. The subscriber can have multiple hosts or sessions, which are all sharing this single SPI. Single SLA mode can be enabled per subscriber and, while enabled, any attempt to establish more than one SPI per subscriber will be rejected by the system.

The default SLA mode on HSQ IOM is expanded SLA mode, which supports multiple SPIs per subscriber. This topic is covered in chapter [High Scale QoS IOM in ESM Context: Expanded SLA Mode](#).

HSQ IOM in the context of ESM provides the following benefits:

- Traffic management functionality on egress with seven tiers of shaping hierarchy, six strict priority scheduling levels, and weighted round-robin (WRR) groups at the subscriber and port level.
- Increased scale with 786k egress queues.
- Aggregate throughput in the range of 200 Gb/s full-duplex per HSQ IOM.
- ~96k subscribers per HSQ IOM in 1:1 (sub/SAP) scenario, each with 8 egress queues in single SLA mode.

With HSQ IOM, the ESM continues to support existing SAP deployment models in the following way:

- Subscriber per SAP (1:1) – single SLA mode and expanded SLA mode
- Multiple subscribers per SAP (N:1) – single SLA mode and expanded SLA mode
- Multiple SAPs per subscriber (service per SAP) – only expanded SLA mode



Note:

1. The example in this chapter is based on subscriber hosts and the subscriber session concept is disabled.
2. A few tens of egress queues are allocated for internal use, so are unavailable for subscribers.
3. The number of subscribers per HSQ IOM in expanded SLA mode depends on the number of available HS primary shapers (16k in total). Because some of the HS primary shapers may be used internally, it is recommended that the exact number of free resources on an HSQ card is periodically checked with the `tools dump system-resources` command.
4. Because each SAP under the same subscriber requires its own SPI, this model is not supported in single SLA mode.

HSQ IOM supports an enhanced egress QoS architecture to provide scalable network, service, and subscriber QoS. At ingress, the HSQ supports regular FP3 QoS with a high ingress policer scaling.

The emphasis in this chapter is on egress shaping and scheduling in the ESM context. Buffer pool management on HSQ is not described in this chapter. Generic ESM concepts are described in other chapters, but for the sake of completeness, a basic ESM configuration as it applies to the example in this chapter is outlined in [Appendix A — Generic ESM configuration](#).

This chapter is conceptually divided into two parts:

- The focus of the first part is on the (egress) QoS configuration for three subscribers, named "sub-1", "sub-2", and "sub-3".
- The second part focuses on examining HSQ traffic management capabilities in the ESM context by observing output traffic patterns in a congested system.

Topics related to the generic operation of HSQ IOM and ESM that are not directly described in this chapter are included in the following:

- Configuring expanded SLA mode in ESM is in the chapter [High Scale QoS IOM in ESM Context: Expanded SLA Mode](#).
- Configuring HSQ in the service context is in the chapter [High Scale QoS IOM: QoS, Service and Network Configuration](#).
- Generic ESM concepts are described in the *7450 ESS, 7750 SR, and VSR Triple Play Service Delivery Architecture Guide* and in numerous TPSDA chapters in the ACG.

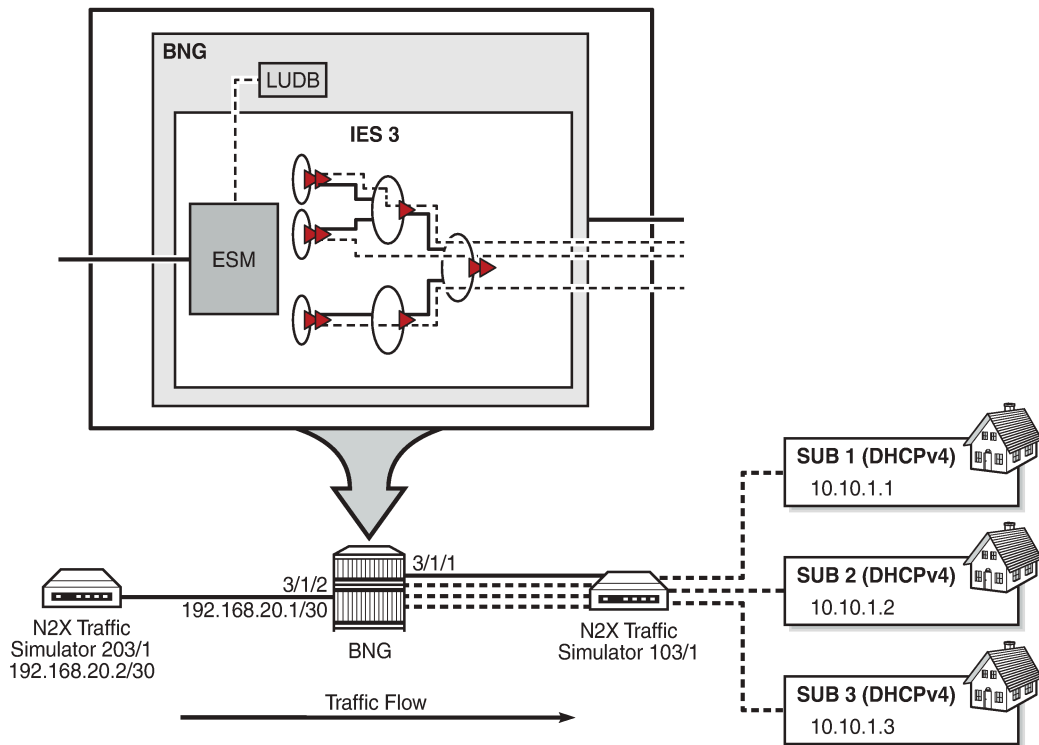
Test environment

[Figure 87: Test environment example](#) shows the test environment example with three subscribers in single SLA mode, each with a single DHCPv4 host setup in the BNG.

DHCPv4 traffic (simulated subscriber hosts) is initiated from port 103/1 on a traffic generator. Subscribers are authenticated via local user database (LUDB) and instantiated on managed SAPs (MSAPs) in the IES 3 service. Subscriber IP addresses are assigned statically via LUDB.

To explore traffic management capabilities under congestion, a number of traffic streams are generated in the downstream direction, from port 203/1 on the same traffic generator toward the subscribers on port 103/1.

Figure 87: Test environment example



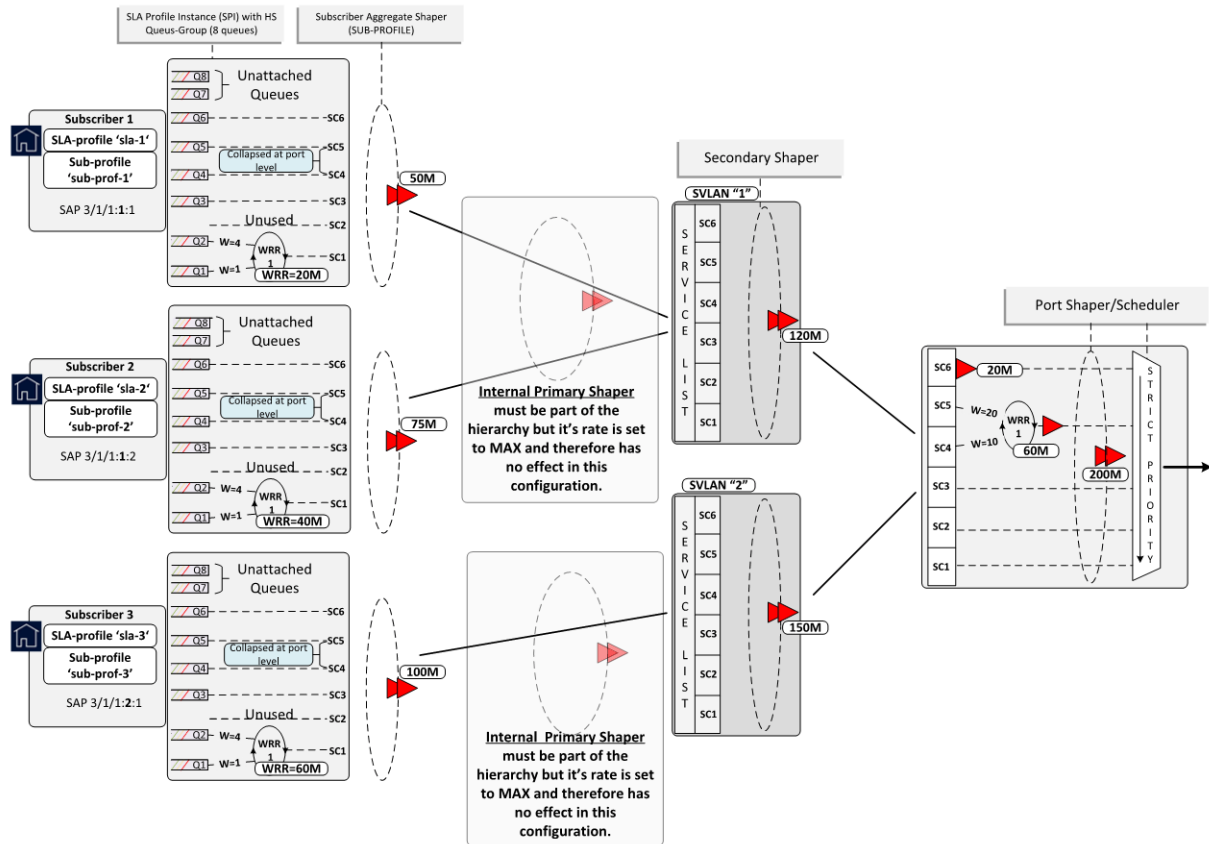
26891

The QoS hierarchy associated with the subscribers is shown in [Figure 88: QoS hierarchy in single SLA mode](#). At a high level, this hierarchy can be described as follows:

- Each of the three subscribers (sub-1, sub-2, and sub-3) is associated with an HSQ queue group. The HSQ queue group always comprises eight queues. However, not all the queues are required to be used by the subscriber. In this example, only six queues are used by each subscriber, while the two remaining queues, although allocated to the subscriber, remain unused.
- All three subscribers share the same HS attachment policy. This means that the mapping between the queues, scheduling classes (SC), and WRR groups is the same for all three subscribers.
- Queues 1 and 2 are attached to WRR group 1 at the local HSQ queue group level and served in 4:1 ratio. WRR group 1 is attached to the lowest scheduling class 1, whereas queues 3, 4, 5, and 6 are directly attached to scheduling classes 3, 4, 5, and 6, respectively.
- Each subscriber is rate limited at the aggregate level.
- Subscribers are mapped to two HS secondary shapers. Limiting the aggregate rates of these shapers will ensure that the bandwidth capacity of the corresponding access nodes in the network is not overrun.
- Mapping between the subscribers and their HS secondary shapers is achieved via outer VLANs. Sub-1 and sub-2 on SAPs with the outer VLAN 1 are mapped to HS secondary shaper 1 while sub-3 on outer VLAN 3 is mapped to HS secondary shaper 2.
- The two HS secondary shapers are associated with the HS port scheduler that has its own aggregate rate limit set.

- At the port level, scheduling class 6 is rate limited while scheduling classes 4 and 5 are collapsed into a WRR group, which is rate limited.

Figure 88: QoS hierarchy in single SLA mode



Configuration

The configuration section is split into two parts:

- [HSQ-specific ESM configuration](#)
- [HSQ-specific QoS configuration](#)

HSQ-specific ESM configuration

The generic ESM configuration used in this example is provided in [Appendix A — Generic ESM configuration](#), for the sake of completeness. This section focuses only on HSQ-specific ESM configuration.

Each of the three subscribers have their own **sub-profile**, which will determine the HS SLA mode in which the subscriber operates, as well as the aggregate rate limit (in kb/s) for each subscriber:

```
configure
```

```
subscriber-mgmt
  sub-profile "sub-prof-1" create
    hs-sla-mode single
    egress
      hs-agg-rate-limit 50000
    exit
  exit
  sub-profile "sub-prof-2" create
    hs-sla-mode single
    egress
      hs-agg-rate-limit 75000
    exit
  exit
  sub-profile "sub-prof-3" create
    hs-sla-mode single
    egress
      hs-agg-rate-limit 100000
    exit
  exit
```

Each subscriber has its own SPI that references the QoS SAP policy on egress (in this case, QoS SAP egress policies 10, 20, and 30). The egress QoS SAP policy defines the QoS characteristics at the local subscriber level (or HSQ queue group level), such as:

- traffic classification
- queue/WRR rates and weights
- mapping between the queues, WRR groups, and scheduling classes by referencing the HS attachment policy



Note: The subscriber, and all the hosts within the subscriber, will share this SPI.

```
configure
  subscriber-mgmt
    sla-profile "sla-1" create
      egress
        qos 10
      exit
    exit
  exit
  sla-profile "sla-2" create
    egress
      qos 20
    exit
  exit
  sla-profile "sla-3" create
    egress
      qos 30
    exit
  exit
exit
```

MSAP policy is a mandatory configuration for dynamically created SAPs (Managed SAP or MSAPs) where, in the context of HSQ, a few key parameters are provisioned:

- The **def-inter-dest-id use-top-q** command associates subscribers with the HS secondary shaper by matching the outer VLAN ID of the subscriber SAP to the HS secondary shaper name. For example, a subscriber on SAP 3/1/1:1.2 will be mapped to HS secondary shaper 1:

```
configure
  port 3/1/1
  ethernet
    mode access
    encap-type qinq
    egress
      hs-scheduler-policy "hs1"
      hs-secondary-shaper "1" create
      aggregate
        rate 150000
      exit
    exit
  exit
exit
exit
exit
```

- The **qos 1 service-queuing** command disables shared queueing on ingress because ingress shared queueing is not supported on HSQ. With shared queueing disabled and service queueing enabled, every configured subscriber or SAP queue on ingress is potentially allocated multiple times, which leads to inefficient use of resources. Every configured queue on ingress is multiplied by the number of possible destination forwarding complexes. For this reason, Nokia recommends that policers are used on ingress.
- The **profiled-traffic-only** command provides a resource optimization that removes queues from the subscriber SAP. By default, one ingress and one egress queue is instantiated per subscriber SAP due to a default QoS SAP policy "1" that is applied to a SAP. In most cases, these SAP queues remain unused because each subscriber uses its own sets of separate subscriber queues that are allocated per SPI. Therefore, the SAP queues can be safely removed.



Note:

1. The default SAP egress policy can be replaced with another, non-default SAP egress policy. However, the SAP egress policy cannot be removed from a SAP.
2. A SAP queue in the ESM context on HSQ is used when multicast per SAP replication mode for the subscriber is enabled. In multicast per SAP replication mode, a single copy of each multicast stream is sent for all hosts of the subscriber via a SAP queue. If the SAP queue is removed, multicast traffic will flow via failover queues or via a queue group that must be explicitly configured.

For example, if there are 50k subscribers on an HSQ IOM, with each subscriber on its own SAP, then removing the default SAP queues:

- preserves 50k queues on ingress. These ingress queues are allocated from the 128k ingress queue pool and are now available for ingress subscriber queueing.

Removing the SAP queues on ingress becomes even more relevant on HSQ, which does not support shared queueing on ingress. Non-shared queueing on ingress means that the number of ingress queues allocated per each SAP would increase proportionally with the number of destination forwarding complexes in the system.

- preserves 50k queues on egress from the 768k egress queue pool. Because each subscriber requires 8 queues on egress, these 50k egress queues can be used for an additional ~6k subscribers.

SAP queue removal works only for subscribers in 1:1 model (subscriber per SAP) and where the number of subscribers per SAP is limited to 1. This requires one additional command: **no multi-sub-sap** (or alternatively **multi-sub-sap limit 1**).

The following command can be used to verify that the SAP queues have been removed:

```
show service id <srvc-id> sap <sap-id> stats
```

The resulting **msap-policy** configuration is as follows:

```
configure
  subscriber-mgmt
    msap-policy "msaps" create
      sub-sla-mgmt
        def-inter-dest-id use-top-q
        sub-ident-policy "sub_ident_pol"
        no multi-sub-sap
        single-sub-parameters
        profiled-traffic-only
      exit
    exit
  ies-vprn-only-sap-parameters
    ingress
      qos 1 service-queuing
    exit
  exit
exit
```

HSQ-specific QoS configuration

At the subscriber level (or HSQ queue group level), the HS attachment policy defines mapping between the queues, WRR groups, and scheduling classes. In this example, queues 1 and 2 are mapped to a WRR group 1, which is mapped into scheduling class 1. Queues 3 to 6 are mapped to respective scheduling classes (SC 3 to 6). Queues 7 and 8, although allocated, are unattached, so will drop any traffic that they receive. A maximum of two WRR groups are supported at the HSQ queue group level and in this example, only WRR 1 is used.

The **low-burst-max-class** parameter is set to 3, which ensures that the 3 lower scheduling classes (1 to 3) are stopped being served first if HSQ queue group aggregate congestion occurs. The HSQ queue group aggregate shaper has two burst tolerance thresholds: when the first threshold is reached, scheduling classes 1 to 3 will be removed from the service list, followed by the higher level scheduling classes (4 to 6) being removed when the second threshold is reached. This designation of scheduling classes in two tiers ensures lower latency for traffic associated with higher scheduling classes during short-lived congestion periods.

The following defined HS attachment policy is applied to the subscriber through the SAP egress policy referenced in the SLA profile. In this example, all three subscribers use the same HS attachment policy:

```
configure
  qos
    hs-attachment-policy "hs-attach-1-1" create
      low-burst-max-class 3
      queue 1 wrr-group 1
      queue 2 wrr-group 1
      queue 3 sched-class 3
```

```

queue 4 sched-class 4
queue 5 sched-class 5
queue 6 sched-class 6
queue 7 unattached
queue 8 unattached
wrr-group 1 sched-class 1
wrr-group 2 unattached
exit
exit

```

Besides referencing the HS attachment policy, the SAP egress policy defines traffic classification, as well as characteristics of queues and WRR groups at the subscriber level. In this example, the SAP egress policy 10 is associated with sub-1. The queues 1 and 4 are in the context of WRR group 1 serviced in the ratio 1:4.



Note:

Although the SAP egress policy syntax implies that such policy is applied per SAP, in the ESM context this policy is instantiated via the SLA profile; therefore, the queue/policer instantiations are performed per SPI and not per SAP.

The aggregate rate of WRR 1 is set to 20 Mb/s. Mapping of forwarding classes to queues is self-explanatory. Similarly, SAP egress policies are applied to sub-2 (policy 20) and sub-3 (policy 30). The only difference between the QoS SAP egress policies for the subscribers is the rate of the WRR group 1, which is 40 Mb/s for sub-2 and 60 Mb/s for sub-3.

The SAP egress policy is applied in the SLA profile for the subscriber. The WRR group rate (along with other QoS parameters) can be dynamically overridden via RADIUS/Diameter during authentication or while the host/session is online. This functionality would allow having only one SAP egress policy configured where parameters can be dynamically overridden, when the policy is applied to the subscriber host or session.

```

configure
qos
  sap-egress 10 create
    hs-attachment-policy "hs-attach-1-1"
    queue 1 create
      hs-wrr-weight 1
    exit
    queue 2 create
      hs-wrr-weight 4
    exit
    queue 3 create
    exit
    queue 4 create
    exit
    queue 5 create
    exit
    queue 6 create
    exit
    hs-wrr-group 1
      rate 20000
    exit
    fc af create
      queue 3
    exit
    fc be create
      queue 1
    exit
    fc h1 create
      queue 5

```

```

exit
fc h2 create
  queue 6
exit
fc l1 create
  queue 4
exit
fc l2 create
  queue 2
exit
dscp af12 fc "af"
dscp be fc "be"
dscp af22 fc "h1"
dscp ef fc "h2"
dscp af21 fc "l1"
dscp af11 fc "l2"
exit

```

The next hierarchy level in the chain (up from the subscriber level) is performed by the two HS secondary shapers that are directly configured in the egress context of the port. The names of the two HS secondary shapers correspond to the outer VLANs on the subscriber SAPs. The network represents the HS secondary shapers as the access nodes downstream from BNG.

The HS secondary shapers "1" and "2" are configured with the aggregate rates of 120 Mb/s and 150 Mb/s, respectively. The rates are configured in kb/s. Similarly, as at the subscriber level, the **low-burst-max-class 3** command maps all scheduling classes at or below level 3 to the low burst tolerance threshold, while all scheduling classes above level 3 are mapped to the high burst tolerance threshold at the HS secondary shaper level. Therefore, in case of congestion, objects associated with scheduling classes 3 and below will be removed from the service list before the objects associated with scheduling classes 4 and above.

```

configure
  port 3/1/1
    ethernet
      mode access
      encap-type qinq
      egress
        hs-scheduler-policy "hs1"
        hs-secondary-shaper "1" create
          aggregate
            rate 120000
            low-burst-max-class 3
          exit
        exit
        hs-secondary-shaper "2" create
          aggregate
            rate 150000
            low-burst-max-class 3
          exit
        exit
      exit
    exit
  no shutdown
exit

```

The last configuration block in the scheduling hierarchy is an HS port scheduler, which is associated with the port via the command **hs-scheduler-policy "hs1"** in the preceding CLI code.

The HS port scheduler characteristics in this example are defined as follows:

- The maximum rate is set to 200 Mb/s.

- Scheduling classes 4 and 5 are collapsed into a single scheduling priority (5) and they are served in a 1:2 ratio. This is performed via WRR group 1.

This collapsing of scheduling classes 4 and 5 occurs at the port level, whereas scheduling classes 1 and 2 are collapsed at the subscriber level (HSQ queue group level).

- The WRR 1 at the port level is rate limited to 60 Mb/s.
- The highest priority (6) scheduling class is rate limited to 20 Mb/s at the port level.



Note: All the rates under the HS port scheduler are in Mb/s.

```
configure
qos
  hs-scheduler-policy "hsl" create
  max-rate 200
  group 1 rate 60
  scheduling-class 4 group 1 weight 10
  scheduling-class 5 group 1 weight 20
  scheduling-class 6 rate 20
exit
exit
```

The delta between the low and high burst tolerance thresholds at the subscriber and HS secondary shaper levels can be adjusted with the **hs-fixed-high-thresh-delta** command that is configured at the card level. In this example, the difference between the thresholds is set to 12 kbytes. The default value is 4000 bytes.

The lower threshold for the buckets is calculated automatically by the system, based on many internal inputs (clock frequency of rate timer, shaper type, heuristics, MDA type, and so on).

```
configure
card 3
  card-type iom4-e-hs
  mda 1
    mda-type me10-10gb-sfp+
    no shutdown
  exit
  fp 1
    egress
      hs-fixed-high-thresh-delta 12000
    exit
  exit
no shutdown
```

Operational commands

Operational commands are used to troubleshoot the system and monitor its operational state. The focus of this section will be on **show**, **clear**, and **tools dump** commands. The **debug** commands are omitted because there are no **debug** commands related to QoS on HSQ IOM and the **debug** commands related to ESM are described in other chapters. Also, there are no log events related to QoS on HSQ IOM.

Show commands

show commands in this section are divided into two groups:

- **show** commands that display association between ESM and QoS objects where all displayed information is static and does not change autonomously over time.
- **show** commands that display QoS hierarchy with the running rates (where the state is changing autonomously).

show commands have filters (or CLI parameters) that can be used to control the amount of the output information. This section will provide a few **show** command examples; it is left to the user to explore all the options available for a particular **show** command.

Subscriber management related show commands

Examining subscriber associations with QoS objects should begin with the **show service active-subscribers** command. The output of this command provides information about the subscriber context and the output can vary in detail depending on the options with which this command is run. Besides the subscriber name, underlying subscriber SAPs, and SLA/sub-profile names, the following information is provided:

- SLA mode in which the subscriber operates on HSQ (single versus expanded SLA mode)
- Aggregate rate of the subscriber
- Ingress and egress QoS policies
- Subscriber association with an HS secondary shaper and the **inter dest id** string that is used to make this association
- Ingress queue/policer statistics
- Egress queues statistics

For brevity, only the information for subscriber sub-1 is shown in the following output:

```
A:PE-1# show service active-subscribers subscriber "sub-1" detail

=====
Active Subscribers
=====
-----
Subscriber sub-1 (sub-prof-1)
-----
I. Policer Ctrl. : N/A
E. Policer Ctrl. : N/A
I. vport-hashing : Disabled
I. sec-sh-hashing: Disabled
Q Frame-Based Ac*: Disabled
Acct. Policy      : N/A
Collect Stats    : Disabled
ANCP Pol.        : N/A
Accu-stats-pol   : (Not Specified)
HostTrk Pol.     : N/A
IGMP Policy      : N/A
MLD Policy       : N/A
PIM Policy       : N/A
Sub. MCAC Policy : N/A
NAT Policy       : N/A
Firewall Policy  : N/A
UPnP Policy      : N/A
NAT Prefix List  : N/A
Def. Encap Offset: none
Encap Offset Mode: none
```

```

Avg Frame Size      : N/A
Vol stats type      : full
Preference          : 5
LAG hash class      : 1
LAG hash weight     : 1
Sub. ANCP-String    : "sub-1"
Sub. Int Dest Id   : "1"
Igmp Rate Adj       : N/A
RADIUS Rate-Limit   : N/A
Oper-Rate-Limit   : 50000
-----
Radius Accounting
-----
Policy              : N/A
Session Opti.Stop   : False
-----
HS
-----
SLA-mode          : single                      E Agg Rate Limit : 50000
Hs Second Shaper : "1"
* indicates that the corresponding row element may have been truncated.
-----
(1) SLA Profile Instance
    - sap:[3/1/1:1.1] (IES 3 - group-int-1)
    - sla:sla-1
-----
Description          : (Not Specified)
Host Limits          : No Limit
Egr Sched-Policy     : N/A
Ingress Qos-Policy : 1                      Egress Qos-Policy : 10
Ingress Queuing Type : Service-queuing (Not Applicable to Policer)
Ingr IP Fltr-Id      : N/A                      Egr IP Fltr-Id      : N/A
Ingr IPv6 Fltr-Id    : N/A                      Egr IPv6 Fltr-Id    : N/A
Ingress Report-Rate  : Maximum
Egress Report-Rate   : Maximum
Egress Remarking     : from Sap Qos
Credit Control Pol. : N/A
Category Map         : (Not Specified)
Use ing L2TP DSCP    : false
Hs-Agg-Rate-Limit    : Maximum
Egress HS Q stat mode: no-override
Hs-Oper-Rate-Limit   : Maximum
Egr hqos mgmt status : disabled
-----
IP Address          MAC Address          Session          Origin          Svc          Fwd
-----
10.10.1.1           00:00:64:01:01:01    N/A              DHCP             3            Y
-----
SLA Profile Instance statistics
-----
Packets          Octets
Off. HiPrio      : 0            0
Off. LowPrio     : 0            0
Off. Uncolor     : 0            0
Off. Managed     : 0            0
Queueing Stats (Ingress QoS Policy 1)
Dro. HiPrio      : 0            0
Dro. LowPrio     : 0            0

```

```

For. InProf      : 0          0
For. OutProf     : 0          0

Queueing Stats (Egress QoS Policy 10)
Dro. In/InplusProf : 88117477      88822416816
Dro. Out/ExcProf   : 68563861      69112371888
For. In/InplusProf : 41604295      41937129360
For. Out/ExcProf   : 87102263      87799081104

```

SLA Profile Instance per Queue statistics

	Packets	Octets
Ingress Queue 1 (Unicast) (Priority)		
Off. HiPrio	: 0	0
Off. LowPrio	: 0	0
Dro. HiPrio	: 0	0
Dro. LowPrio	: 0	0
For. InProf	: 0	0
For. OutProf	: 0	0
Egress Queue 1		
Dro. In/InplusProf	: 0	0
Dro. Out/ExcProf	: 62846249	63349018992
For. In/InplusProf	: 0	0
For. Out/ExcProf	: 2014638	2030755104
Egress Queue 2		
Dro. In/InplusProf	: 56805196	57259637568
Dro. Out/ExcProf	: 0	0
For. In/InplusProf	: 8055689	8120134512
For. Out/ExcProf	: 0	0
Egress Queue 3		
Dro. In/InplusProf	: 0	0
Dro. Out/ExcProf	: 256284	258334272
For. In/InplusProf	: 0	0
For. Out/ExcProf	: 51632424	52045483392
Egress Queue 4		
Dro. In/InplusProf	: 22176682	22354095456
Dro. Out/ExcProf	: 0	0
For. In/InplusProf	: 16739851	16873769808
For. Out/ExcProf	: 0	0
Egress Queue 5		
Dro. In/InplusProf	: 0	0
Dro. Out/ExcProf	: 5461328	5505018624
For. In/InplusProf	: 0	0
For. Out/ExcProf	: 33455201	33722842608
Egress Queue 6		
Dro. In/InplusProf	: 9135599	9208683792
Dro. Out/ExcProf	: 0	0
For. In/InplusProf	: 16808755	16943225040
For. Out/ExcProf	: 0	0

To reveal the subscriber hierarchy in a terse form with respect to the sub/SLA-profiles and the SAP, the following command can be run:

```
A:PE-1# show service active-subscribers hierarchy
```

```
=====
Active Subscribers Hierarchy
=====
-- sub-1 (sub-prof-1)
|
+-- sap:[3/1/1:1.1] - sla:sla-1
|
+-- 10.10.1.1 - mac:00:00:64:01:01:01 - DHCP - svc:3

-- sub-2 (sub-prof-2)
|
+-- sap:[3/1/1:1.2] - sla:sla-2
|
+-- 10.10.1.2 - mac:00:00:64:01:01:02 - DHCP - svc:3

-- sub-3 (sub-prof-3)
|
+-- sap:[3/1/1:2.1] - sla:sla-3
|
+-- 10.10.1.3 - mac:00:00:64:01:01:03 - DHCP - svc:3

-----
Number of active subscribers : 3
Flags: (N) = the host or the managed route is in non-forwarding state
=====
```

The following SAP related **show** command confirms that the SAP queues are removed from the underlying subscriber SAP (under the **stats** section at the end of the output). This was ensured by configuring the **profiled-traffic-only** command in the MSAP policy, with the purpose of reducing the queue consumption on ingress and egress.

```
A:PE-1# show service id 3 sap 3/1/1:1.1 detail

=====
Service Access Points(SAP)
=====
Service Id      : 3
SAP             : 3/1/1:1.1          Encap             : qinq
QinQ Dot1p     : Default
Description     : Managed SAP - Capture Svc 10 3/1/1:*.
Admin State    : Up                  Oper State       : Up
Flags          : None
Multi Svc Site : None
Last Status Change : 09/08/2017 15:16:20
Last Mgmt Change  : 09/11/2017 14:22:55
Sub Type       : managed
Capture Service Id : 10              Capture SAP      : 3/1/1:*.
MSAP Policy    : msaps
Idle           : no                  Sticky           : no
Dot1Q Ethertype : 0x8100             QinQ Ethertype   : 0x8100
Split Horizon Group: (Not Specified)

Admin MTU       : 1522                Oper MTU         : 1522
Ingr IP Fltr-Id : n/a                 Egr IP Fltr-Id  : n/a
Ingr Mac Fltr-Id : n/a                 Egr Mac Fltr-Id : n/a
Ingr IPv6 Fltr-Id : n/a                Egr IPv6 Fltr-Id : n/a
qinq-pbit-marking : both
Q Frame-Based Acct : Disabled          Egr Agg Rate Limit: max
Limit Unused BW   : Disabled

Acct. Pol        : None                Collect Stats    : Disabled

Anti Spoofing : Ip-Mac             Dynamic Hosts    : Enabled
```



```

Avl Static Hosts   : 0                               Tot Static Hosts   : 0
Calling-Station-Id : n/a

Application Profile: None
Transit Policy     : None
AARP Id           : None

Oper Group         : (none)                           Monitor Oper Grp  : (none)
Host Lockout Plcy  : n/a
Lag Link Map Prof  : (none)
Bandwidth          : Not-Applicable
Oper DCpu Prot Pol*: _default-access-policy

-----
ETH-CFM SAP specifics
-----
Tunnel Faults      : accept                           AIS                : Disabled
MC Prop-Hold-Timer : n/a
Squelch Levels     : None
Collect Lmm Stats  : Disabled
LMM FC Stats       : None
LMM FC In Prof     : None

-----
QOS
-----
Ingress qos-policy : 1                               Egress qos-policy : 1
Ingress FP QGrp    : (none)                           Egress Port QGrp  : (none)
Ing FP QGrp Inst   : (none)                           Egr Port QGrp Inst: (none)
Shared Q plcy      : n/a                               Multipoint shared  : Disabled
I. Sched Pol       : (Not Specified)
E. Sched Pol       : (Not Specified)
I. Policer Ctl Pol : (Not Specified)
E. Policer Ctl Pol : (Not Specified)
E. HS Sec. Shaper  : (Not Specified)
I. QGrp Redir. List: (Not Specified)
E. QGrp Redir. List: (Not Specified)

-----
Subscriber Management
-----
Admin State        : Up                               MAC DA Hashing    : False
Def Sub-Id         : Use auto-sub-id
Def Sub-Profile    : None
Def SLA-Profile    : None
Def Inter-Dest-Id  : (Use top-q-tag)
Def App-Profile    : None
Sub-Ident-Policy   : sub_ident_pol
Subscriber Limit   : 1
Single-Sub-Parameters
  Prof Traffic Only : True
  Non-Sub-Traffic   : N/A

Static host management
MAC learn options  : N/A

-----
Sap Statistics
-----
Last Cleared Time  : N/A

CPM Ingress        : 3                               Packets            : 3
                                                           Octets            : 443

Forwarding Engine Stats

```

```

Dropped          : 0          0
Received Valid   : 0          0
Off. HiPrio      : 0          0
Off. LowPrio     : 0          0
Off. Uncolor     : 0          0
Off. Managed     : 0          0

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio      : 0          0
Dro. LowPrio     : 0          0
For. InProf      : 0          0
For. OutProf     : 0          0

Queueing Stats(Egress QoS Policy 1)
Dro. In/InplusProf : 0          0
Dro. Out/ExcProf   : 0          0
For. In/InplusProf : 0          0
For. Out/ExcProf   : 0          0
-----
Sap per Queue stats
-----
Packets          Octets
No entries found
=====
* indicates that the corresponding row element may have been truncated.
A:PE-1#

```

QoS-related show commands in ESM context

Examination of the subscriber QoS hierarchy on HSQ can start at the subscriber (HSQ queue group and SAP) level and gradually move through the HS secondary shaper, and finally the port level. The output of the **show** commands should confirm that the subscriber is associated with the QoS object as intended by the configuration.

For example, the following output shows that HS attachment policy "hs-attach-1-1" is associated with QoS SAP egress policy 10. It was determined previously that QoS SAP egress policy 10 is associated with subscribersub-1. In this case, a two-step process was necessary to track the association between the subscriber and the HS attachment policy.

```

A:PE-1# show qos sap-egress 10 association

=====
QoS Sap Egress
=====

-----
Sap Egress Policy (10)
-----
Policy-id          : 10          Scope          : Template
Ethernet-ctag      : False       Parent-loc     : default
Name               : (Not Specified)
Description         : (Not Specified)
Policy Active       : True        Plcrs HQoS Managed : False
Post Plcr Mapping Policy: (Not Specified)
HS Attachment Policy : hs-attach-1-1

-----
Dynamic Configuration Information
-----
PccRule Insert Point : n/a          DynPlcr Insert Point : n/a

```

```

CBS           : Def      MBS           : Def
Parent        : (Not Specified)
Level         : 1        Weight        : 1
Packet Byte Offset : 0
Stat Mode     : minimal
-----
Associations
-----
SLA Profiles :
- sla-1
-----
HSMDA Associations
-----
No Associations Found.
=====

```

The output from the HS attachment policy reflects the QoS configuration state at the subscriber level that is shown in [Figure 88: QoS hierarchy in single SLA mode](#):

- Queues 1 and 2 are attached to WRR 1.
- Queues 3 to 6 are directly attached to the corresponding scheduling classes (3 to 6).
- Queues 7 and 8 are unattached.
- WRR 2 is unattached.
- HS attachment policy "hs-attach-1" is associated with QoS SAP egress policies 10, 20, and 30 that correspond to sub-1, sub-2, and sub-3.

```

A:PE-1# show qos hs-attachment-policy "hs-attach-1-1" detail
=====
HS Attachment Policy Information
=====
Policy Name      : hs-attach-1-1
Description      : (Not Specified)
Low Burst Max Class : 3
-----
Queue            Scheduling Class      WRR Group
-----
1                (Not-Applicable)          1
2                (Not-Applicable)          1
3                3                (Not-Applicable)
4                4                (Not-Applicable)
5                5                (Not-Applicable)
6                6                (Not-Applicable)
7                unattached        unattached
8                unattached        unattached
-----
WRR Group        Scheduling Class
-----
1                1
2                unattached
-----
Associations
-----

```

```

Network-Queue Policy
-----
No Matching Entries

Sap-Egress Policy
-----
10
20
30

Egress Queue-Group Templates
-----
No Matching Entries
-----

```

Association between the HS secondary shaper and the subscribers can be verified with the following command. The HS secondary shaper is allocated per port (or LAG). The two subscribers (sub-1 and sub-2) are instantiated on SAPs with the outer VLAN tag 1; consequently, they are both associated with HS secondary shaper 1. The HS secondary shaper 1 is rate limited to 120 Mb/s while its scheduling classes are left open (max rate).

```

A:PE-1# show port 3/1/1 hs-secondary-shaper "1" associations

=====
Ethernet Port 3/1/1 Egress HS Secondary Shaper Information
=====
Policy Name      : 1
Description      : (Not Specified)
Rate             : 120000 Kbps
Low Burst Max Class: 3

-----
Class            Rate
-----
1                max
2                max
3                max
4                max
5                max
6                max
-----

-----
Service Associations
-----
Service ID      Service Type      SAP
-----
No Service Associations Found.
-----

-----
Subscriber Associations
-----
Subscriber ID
-----
sub-1
sub-2
-----
Number of subscriber associations : 2
-----

```

The port scheduler information can be obtained with the following command. This command is run in the QoS context (as opposed to being run in the port context, which was the case for HS secondary schedulers). The reason for this is that the HS secondary scheduler is configured directly under the port, while the HS port scheduler is configured in an HS scheduler policy (in QoS context), which is then applied to a port.

```
A:PE-1# show qos hs-scheduler-policy "hs1" detail
```

```
=====
```

```
HS Scheduler Policy Information
```

```
=====
```

```
Policy Name           : hs1
Description            : (Not Specified)
Max Rate               : 200 Mbps
```

```
-----
```

Scheduling Class	Rate	Group	Weight in Group
1	max	0	1
2	max	0	1
3	max	0	1
4	max	1	10
5	max	1	20
6	20 Mbps	0	1

```
-----
```

Group	Rate
1	60 Mbps

```
-----
```

```
Port Ethernet Egress Associations
```

```
-----
```

```
3/1/1
```

```
-----
```

```
=====
```

Show commands with dynamically changing information

The following commands show the QoS hierarchy with the measured rates of the objects (queues, WRR groups, scheduling classes, secondary shapers, and port shapers) and the queue buffer depths in the QoS hierarchy. The command output in this section is based on the scenario described in the [Traffic management on HSQ](#) section.

The following subscriber hierarchy shows the rates per scheduling priority (and consequently, scheduling class), starting at the queue level and moving up toward the port level. For example, scheduling priority 1 starts with the rates of two subscriber queues (1 and 2) that are mapped to the WRR group. Scheduling priority 1 then moves to the rate at the HS secondary shaper level (the summed rate of all the entities at scheduling priority 1 at the HS secondary shaper level), ending with the rate of the scheduling class 1 at the port level. The aggregate rate of the HS secondary shaper and the port in the subscriber hierarchy are also provided.



Note: Running rates are the dynamically calculated rates while traffic is running.

```
A:PE-1# show qos hs-scheduler-hierarchy subscriber "sub-1" egress
```

```

=====
Hs Scheduler Hierarchy Information
=====
PortId                : 3/1/1
SAP                   : [3/1/1:1.1]
SLA Profile           : sla-1
Hs Sched Policy Name  : hs1

Port Max-Rate : 200 Mbps
Hs-Sec-Shaper:1 Agg-Rate : 120017 Kbps

Scheduler Priority 6
  Scheduler Class 6 Rate : 19 Mbps
    Hs-Sec-Shaper:1 Class 6 Rate : 13349 Kbps
      Queue : 6 Rate : 6528 Kbps

Scheduler Priority 5 Group 1
  Scheduler Class 5 Rate : 40 Mbps Weight : 20
    Hs-Sec-Shaper:1 Class 5 Rate : 26699 Kbps
      Queue : 5 Rate : 13072 Kbps
  Scheduler Class 4 Rate : 19 Mbps Weight : 10
    Hs-Sec-Shaper:1 Class 4 Rate : 13300 Kbps
      Queue : 4 Rate : 6528 Kbps

Scheduler Priority 3
  Scheduler Class 3 Rate : 61 Mbps
    Hs-Sec-Shaper:1 Class 3 Rate : 40800 Kbps
      Queue : 3 Rate : 20008 Kbps

Scheduler Priority 2
  Scheduler Class 2 Rate : 0 Mbps
    Hs-Sec-Shaper:1 Class 2 Rate : 0 Kbps

Scheduler Priority 1
  Scheduler Class 1 Rate : 58 Mbps
    Hs-Sec-Shaper:1 Class 1 Rate : 25867 Kbps
      Queue : 1 Group : 1 Rate : 752 Kbps
      Queue : 2 Group : 1 Rate : 3104 Kbps
=====

```

The following command provides the measured rates at the HS secondary shaper and port levels:

```

A:PE-1# show qos hs-scheduler-hierarchy port 3/1/1 hs-secondary-shapers

=====
Hs Scheduler Hierarchy Information
=====
Hs Sched Policy Name      : hs1

Port Max-Rate : 200 Mbps

Scheduler Priority 6
  Scheduler Class 6 Rate : 20 Mbps

Scheduler Priority 5 Group 1
  Scheduler Class 5 Rate : 40 Mbps Weight : 20
  Scheduler Class 4 Rate : 19 Mbps Weight : 10

Scheduler Priority 3
  Scheduler Class 3 Rate : 61 Mbps

Scheduler Priority 2
  Scheduler Class 2 Rate : 0 Mbps

```

```
Scheduler Priority 1
  Scheduler Class 1 Rate : 58 Mbps

-----
HS Secondary Shaper Rates
-----
Hs-Sec-Shaper:1 Agg-Rate : 119519 Kbps
  Class 6 Rate : 13268 Kbps
  Class 5 Rate : 26544 Kbps
  Class 4 Rate : 13284 Kbps
  Class 3 Rate : 40661 Kbps
  Class 2 Rate : 0 Kbps
  Class 1 Rate : 25761 Kbps

Hs-Sec-Shaper:2 Agg-Rate : 79764 Kbps
  Class 6 Rate : 6642 Kbps
  Class 5 Rate : 13300 Kbps
  Class 4 Rate : 6625 Kbps
  Class 3 Rate : 20342 Kbps
  Class 2 Rate : 0 Kbps
  Class 1 Rate : 32852 Kbps

Hs-Sec-Shaper:default Agg-Rate : 0 Kbps
  Class 6 Rate : 0 Kbps
  Class 5 Rate : 0 Kbps
  Class 4 Rate : 0 Kbps
  Class 3 Rate : 0 Kbps
  Class 2 Rate : 0 Kbps
  Class 1 Rate : 0 Kbps

-----
```

Another important parameter to monitor is the depth of the queues. This provides information about the degree of congestion at the queue level. The buffer space per queue is allocated automatically by the system and, in the following case, the buffers are rather large. Deep buffering causes longer delays. To avoid this, the queue buffers can be adjusted by the **mbs** command under the queue definition in the QoS SAP egress policy. For example, 15 kbytes would accommodate roughly fifteen 1000 byte packets in a buffer queue.

```
*A:PE-1>config>qos>sap-egress# info
-----
      queue 1 create
      mbs 15 kilobytes
```

The following output shows that, except for queue 3, all queues have their buffers fully used. Because there is no congestion on queue 3, its buffer is unused.

```
A:PE-1# show hs-pools port 3/1/1 egress subscriber "sub-1" | match "Queue Information" pre-
lines 1 post-lines 100
-----
Queue Information
-----
Queue Name      : Sub=sub-1:sla-1 3->3/1/1:1.1->1
FC Map          : be ef nc
Admin PIR       : 20000                               Oper PIR           : 0
Admin MBS       : 64 KB                               Oper MBS          : 64 KB
HS Wrr Group    : 1
HS Wrr Class Weight: 1                               HS Wrr Weight     : 1
Depth          : 58 KB
HS Class        : 1                               HS Alt Port Class Pool : No
```

```

HS Slope Policy      : _tmnx_hs_default

Queue Name          : Sub=sub-1:sla-1 3->3/1/1:1.1->2
FC Map              : l2
Admin PIR           : 20000                      Oper PIR           : 0
Admin MBS           : 64 KB                      Oper MBS           : 64 KB
HS Wrr Group        : 1
HS Wrr Class Weight : 1                      HS Wrr Weight      : 4
Depth               : 58 KB
HS Class            : 1                      HS Alt Port Class Pool : No
HS Slope Policy      : _tmnx_hs_default

Queue Name          : Sub=sub-1:sla-1 3->3/1/1:1.1->3
FC Map              : af
Admin PIR           : Max                      Oper PIR           : Max
Admin MBS           : 262500 B                  Oper MBS           : 262656 B
HS Wrr Group        : (not-applicable)
HS Wrr Class Weight : 1                      HS Wrr Weight      : 0
Depth               : 2 KB
HS Class            : 3                      HS Alt Port Class Pool : No
HS Slope Policy      : _tmnx_hs_default

Queue Name          : Sub=sub-1:sla-1 3->3/1/1:1.1->4
FC Map              : l1
Admin PIR           : Max                      Oper PIR           : Max
Admin MBS           : 262500 B                  Oper MBS           : 262656 B
HS Wrr Group        : (not-applicable)
HS Wrr Class Weight : 1                      HS Wrr Weight      : 0
Depth               : 233 KB
HS Class            : 4                      HS Alt Port Class Pool : No
HS Slope Policy      : _tmnx_hs_default

Queue Name          : Sub=sub-1:sla-1 3->3/1/1:1.1->5
FC Map              : h1
Admin PIR           : Max                      Oper PIR           : Max
Admin MBS           : 262500 B                  Oper MBS           : 262656 B
HS Wrr Group        : (not-applicable)
HS Wrr Class Weight : 1                      HS Wrr Weight      : 0
Depth               : 233 KB
HS Class            : 5                      HS Alt Port Class Pool : No
HS Slope Policy      : _tmnx_hs_default

Queue Name          : Sub=sub-1:sla-1 3->3/1/1:1.1->6
FC Map              : h2
Admin PIR           : Max                      Oper PIR           : Max
Admin MBS           : 262500 B                  Oper MBS           : 262656 B
HS Wrr Group        : (not-applicable)
HS Wrr Class Weight : 1                      HS Wrr Weight      : 0
Depth               : 233 KB
HS Class            : 6                      HS Alt Port Class Pool : No
HS Slope Policy      : _tmnx_hs_default

-----
=====
A:PE-1#

```

Clear commands

Clear commands in the HSQ context are used to clear statistics associated with the HS secondary shaper:

```
clear port <port-id> hs-secondary-shaper <name> statistics
```


Resources monitoring

The following command is used to monitor resource on an HSQ IOM. Some of the resources (HSQ queue groups, HS primary shapers, HS secondary shapers, HS turbo queue groups, and so on) are allocated for internal use, thereby reducing the number of resources in the **Free** column. Such internally consumed resources are not available to be part of the user configuration.

The number of internally consumed resources depends on the configuration and MDA types.

In the following example, the number of allocated HS queue groups is 39. Each of the three subscribers consumes one HS queue group, which means that 36 HS queue groups are internally allocated. Similarly, out of 23 allocated HS secondary shapers, only 3 are user related (and visible via **show** commands): a "default" HS secondary shaper, HS secondary shaper 1, and HS secondary shaper 2. This means that 20 HS secondary shapers are internally consumed.

The same logic can be used for turbo HS queue groups and HS primary shapers. In addition, one managed HS primary shaper is always allocated per each HS secondary shaper. Both turbo HS queue groups and HS primary shapers are out of the scope of this chapter.



Note:

100G ports support port-based (access or network) HSQ queue groups on egress that can be configured to support higher throughput rates. Such port-based high rate HSQ queue groups on egress are referred to as egress HS turbo queue groups.

Hardware Resource Usage for Slot #3, CardType iom4-e-hs, Cmplx #0:			
	Total	Allocated	Free
-----	-----	-----	-----
SAP Ingress QoS Policies	1791	1	1790
Dynamic Egress Classification +	2047	4	2043
SAP Egress QoS Policies -		4	
Network Egress Classification -		0	
Ingress Queues	131072	501	130571
Egress Queues	786432	137	786295
Egress HS Turbo Queue Groups	64	10	54
Egress HS Queue Groups	98240	39	98201
Primary Shapers +	16384	23	16361
Explicit Primary Shapers -		0	
Managed Primary Shapers -		23	
Secondary Shapers	4096	23	4073
Ingress Policers	511999	1	511998
Egress Policers	262143	1	262142
Ingress Policer Stats	511967	0	511967
Egress Policer Stats	262111	0	262111
Qos Ingress Root Arbiters	65535	1	65534
Qos Egress Root Arbiters	65535	1	65534
Qos Intermediate Arbiters	262143	0	262143
Egress QoS Bypass	131071	0	131071
Ingress ACL Entries	65536	2	65534
Ingress QoS Entries	16384	2	16382
Ingress IPv6 ACL Entries	28672	2	28670
Ingress IPv6 QoS Entries	4096	2	4094
Egress ACL Entries	32768	2	32766
Egress QoS Entries	14336	2	14334
Egress IPv6 ACL Entries	16384	2	16382
Egress IPv6 QoS Entries	2048	2	2046
Ingress ACL Filters	2047	0	2047
Ingress IPv6 ACL Filters	2047	0	2047
Egress ACL Filters	2047	0	2047
Egress IPv6 ACL Filters	2047	0	2047

QoS User Schedulers		98303		0		98303
QoS User Scheduler Overrides		196607		0		196607
Sap IngQGrp RedirLst Entries		31999		0		31999
Dynamic Service Entries	+	131071		3		131068
Subscriber Hosts	-	131071		3		131068
Encap Group Members	-	65535		0		65535
Egr Network Queue Group Mappings	-	131071		0		131071
SapInst EgrQGrp RedirLst Entries	-	31999		0		31999
Dynamic Nexthop Entries	+	511999		3		511996
Subscriber Nexthops	-	511999		3		511996
Isec tunnels	-	511999		0		511999
Subscriber SPI QoS Overrides		131072		0		131072
Mac Fdb Entries		511999		0		511999
Egress TLS Mcast Entries		368639		1		368638

Traffic management on HSQ

This section examines traffic output on HSQ IOM during congestion. [Figure 89: Managing Congestion on HSQ in Single SLA Mode](#) is a graphic representation of the configuration described previously, but with traffic streams running through the IOM.

Six traffic streams are sent in the downstream direction toward each of the three subscribers-in total there are 18 traffic streams. The traffic streams are shown on the left side of the figure, with their names, offered rates (IN column) and measured output rates (OUT column). The traffic streams are sent and analyzed by the traffic generator.

The three gray shaded squares in the center of the figure represent the three subscribers and their scheduling classes. The QoS hierarchy is shown on the right side. Red shaded areas (shapers) represent points of congestion on HSQ IOM caused by the 18 traffic streams.

Each of the six traffic streams per subscriber is fed into different subscriber queues. The first digit in the traffic stream name represents the subscriber, whereas the second digit represents the queue to which this stream is sent. For example, STRM 2-3 represents a traffic stream sent to sub-2, queue 3.

To summarize the scenario shown in [Figure 89: Managing Congestion on HSQ in Single SLA Mode](#):

- Traffic streams 1 and 2 of each subscriber are mapped to subscriber queues 1 and 2, which are in turn associated with WRR group 1 at the subscriber level.
- WRR group 1 is, depending on the subscriber, rate limited to 20 Mb/s, 40 Mb/s, or 60 Mb/s. Weight ratio between queues 2 and 1 is 4:1. WRR group 1 is then attached to scheduling class 1.
- Traffic stream 3 is via queue 3 directly mapped to scheduling class 3. Scheduling class 2 is unused in this example.
- Traffic streams 4 and 5 are via queues 4 and 5 mapped to scheduling classes 4 and 5, which are at the port level collapsed into WRR group 1 with an aggregate rate limit of 60 Mb/s.
- Traffic stream 6 is the highest priority stream that is via queue 6 mapped to scheduling class 6. At the port level, scheduling class 6 is rate limited to 20 Mb/s.
- Subscriber (sub-1, sub-2, and sub-3) aggregate rates are set to 50 Mb/s, 75 Mb/s, and 100 Mb/s, respectively.
- Subscribers sub-1 and sub-2 are mapped to HS secondary shaper 1 (via the outer VLANs on their SAPs) while sub-3 is mapped to HS secondary shaper 2.
- HS port scheduler is rate limited to 200 Mb/s.

The configured rate limits at the subscriber level are L2 rates, whereas the configured rate limits at the HS secondary shaper level and the HS port level include L1 overhead and are, therefore, on-the-wire rates. On-the-wire rates account for 20 additional bytes in each Ethernet frame (8 bytes preamble and 12 bytes IFG).

All traffic streams are sent with constant rates (no added burstiness) and fixed packet size (1000 bytes). Therefore, the difference between the L2 rates and the on-the-wire rates for the 1000 byte packets is $1000/1020 = 0.98$ or 2%. That is, on-the-wire rates are 2% higher than the L2 rates. The name for this 2% delta factor in this chapter will be the Rate Conversion Factor (RCF).

Figure 89: Managing Congestion on HSQ in Single SLA Mode

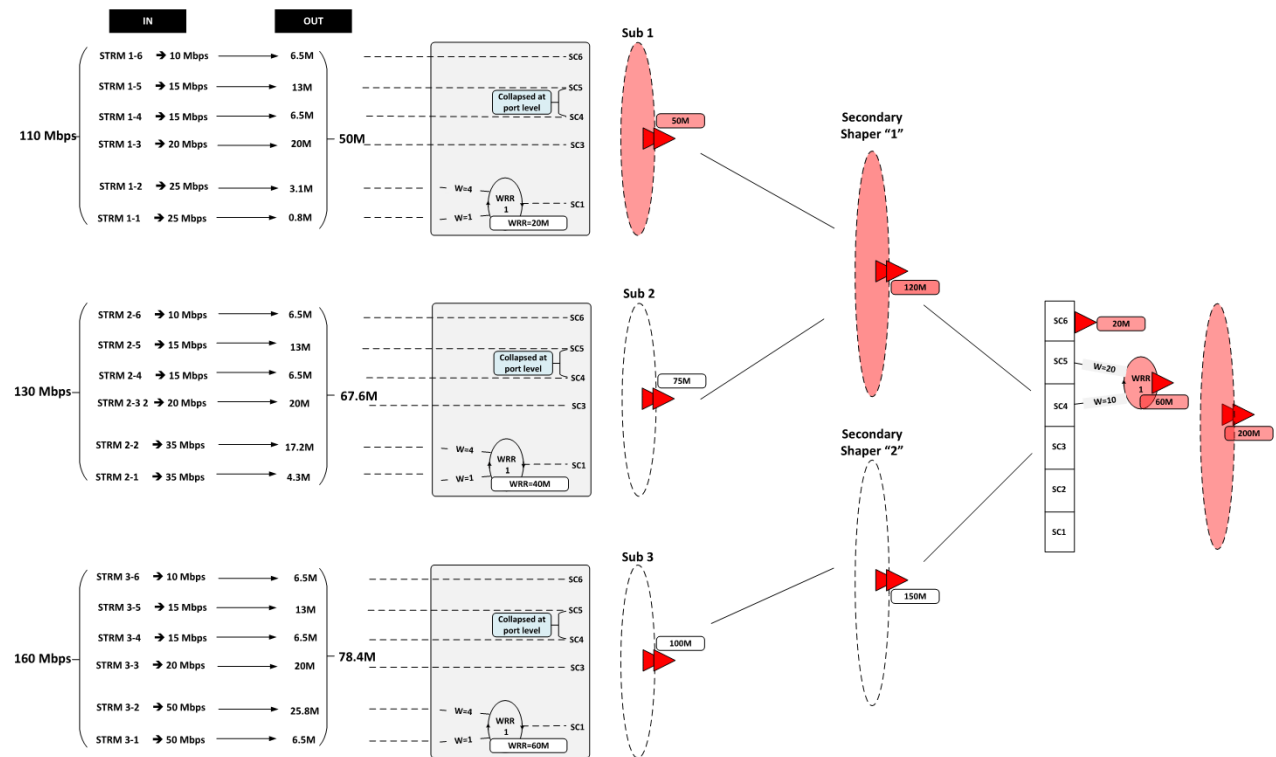


Table 23: Input and output rates throughout the subscriber QoS hierarchy lists the input/output rates throughout the subscriber QoS hierarchy. Red shaded table cells represent congested objects. The numbers on the blue background represent the output rates measured on the traffic generator. The numbers in red (in parentheses) are configured aggregate rates for the object in the hierarchy. The numbers above them are operational rates measured by SR OS and observed via the **show qos hs-scheduler-hierarchy** command.

Some of the **output** columns in this table require additional explanation:

- Per stream output rates are L2 rates as measured by the traffic generator. The resulting analysis will be based on these rates.
- However, those rates can also be displayed via the **show qos hs-scheduler-hierarchy subscriber <sub-name> egress** command. RCF must be used to correctly interpret the results in respective rate domains (L1 versus L2).
- WRR rates at the subscriber level (second column under the Output Rate section) are populated by manually adding rates measured by the traffic generator for traffic streams 1 and 2. The 4:1 indication in

parentheses represents the weight ratio between the two streams. There are no means to observe the WRR group rates at the subscriber level directly within the system (via **show** commands).

- Aggregate rates per subscriber (third column under the Output Rate section) are populated by manually adding the rates measured by the traffic generator from all six traffic streams for each subscriber. There are no means to observe the aggregate subscriber rate directly within the system (via **show** commands).
- Aggregate rates per HS secondary shaper (fourth column under the Output Rate section) are on-the-wire rates displayed via the **show qos hs-scheduler-hierarchy port 3/1/1 hs-secondary-shapers** command. The displayed rates are rounded to the nearest Mb/s.
- Per scheduling class rates at the port level (fifth column under the Output Rate section) are displayed via the **show qos hs-scheduler-hierarchy port 3/1/1** command. The displayed rates are rounded by the system to the nearest Mb/s. The discrepancy between the actual rates for scheduling class 3 (60 Mb/s) and the displayed rate (61 Mb/s) is due to the measuring and rounding inaccuracy at display time. The actual rate at scheduling class 3 at the port level can be calculated by summing the measured rate of each traffic stream 3 for each subscriber and adjusting the sum for the on-the-wire rate (RCF).
- The same logic applies to scheduling class 1 (actual rate 60 Mb/s versus 59 Mb/s measured rate).
- The WRR 1 rate at the port level (sixth column under the Output Rate section) is manually calculated by adding the rates measured by the traffic generator of all streams mapped to scheduling classes 5 and 4 (6 traffic streams in total, two per each subscriber). At the port level, traffic on scheduling classes 5 and 4 is weighted in a 2:1 ratio (this is noted in parentheses).
- The port rate (last column under the Output Rate section) is collected from two places:
 - The top number is obtained via the **show qos hs-scheduler-hierarchy** command and represents the on-the-wire rate.
 - The bottom number with blue background is the number measured by the traffic generator, which only measures L2 rates. On-the-wire rates can be converted to L2 rates by multiplying the on-the-wire rates by the RCF.

Table 23: Input and output rates throughout the subscriber QoS hierarchy

Subscr	Strm	Input rate in [Mb/s]		Output rate in [Mb/s]						
		Per strm	Agg per subscr	Per strm (L2)	WRR on subscr level	Agg per subscr (L2)	Agg per second shaper (wire)	Per sch class on port level (wire)	WRR-1 on port level (wire)	Port
Sub-1	1-6	10	110	6.54	-	50 (50)	120 (120)	SC 6	-	200 (200) 196.08 (L2)
	1-5	15		13.07	-			20	-	
	1-4	15		6.54	-			(20)	-	
	1-3	20		20	-			SC 5	60	
	1-2	25		3.1	3.9			40	(60) (2:1)	
	1-1	25		0.8	(4:1)					

Subscr	Strm	Input rate in [Mb/s]		Output rate in [Mb/s]						
		Per strm	Agg per subscr	Per strm (L2)	WRR on subscr level	Agg per subscr (L2)	Agg per second shaper (wire)	Per sch class on port level (wire)	WRR-1 on port level (wire)	Port
Sub-2	2-6	10	130	6.54	-	67.65 (75)		SC 4 20		
	2-5	15		13.07	-					
	2-4	15		6.54	-					
	2-3	20		20	-			SC 3 61	-	
	2-2	35		17.2	21.5				-	
	2-1	35		4.3	(4:1)				-	
Sub-3	3-6	10	160	6.54	-	78.43 (100)	80 (150)	SC 2 0	-	
	3-5	15		13.07	-				-	
	3-4	15		6.54	-				-	
	3-3	20		20	-			SC 1 59	-	
	3-2	50		25.8	32.3				-	
	3-1	50		6.5	(4:1)				-	

Analysis of results

The analysis of results begins with the rates per stream measured by the traffic generator. The expected behavior is that those rates are in line with the theoretical rate calculations based on our understanding of QoS on HSQ IOM.

Considering the six strict priority classes in the HSQ scheduling mechanism, the expectation is that the traffic is serviced in the order of priority, from the highest scheduling class 6 to the lowest scheduling class 1. Consequently, the traffic analysis starts with the streams that are mapped to the highest priority scheduling class 6 (streams 1-6, 2-6, and 3-6), that is, stream 6 of sub-1, sub-2, and sub-3, respectively.

Scheduling class 6 (streams 1-6, 2-6, and 3-6) - the highest priority scheduling class

Due to the aggregate rate limit of 20 Mb/s for scheduling class 6 at the port level, it is expected that each subscriber receives an equal amount of traffic on scheduling class 6:

$$\frac{20 \text{ Mb/s (aggregate of SC 6)}}{3 \text{ streams (one per subscriber)}} = 6.6 \text{ Mb/s}$$

26900

The measured results on the traffic generator for traffic streams 6 in [Table 23: Input and output rates throughout the subscriber QoS hierarchy](#) show that each subscriber receives 6.54 Mb/s out of an offered 10 Mb/s. The slight difference between the expected (6.6 Mb/s) and the measured (6.54 Mb/s) rate is caused by the discrepancy between the L2 rates at the subscriber level (as measured by the traffic generator), and on-the-wire rate enforcement (20 Mb/s) at the scheduling class 6 at the port level. Multiplying 6.6 Mb/s by the RCF

$$\left(\frac{1000}{1020} \right)$$

will align the results.

This slight discrepancy between the L2 and on-the-wire rates is common throughout the remaining analysis, and will not be repeated.

Scheduling classes 5 and 4 (streams 1-5, 1-4, 2-5, 2-4, 3-5, and 3-4)

Scheduling classes 5 and 4 are collapsed at the port level into WRR group 1, which is at the next scheduling priority to be served. These two scheduling classes are served by WRR group 1 in a 2:1 ratio. The aggregate rate limit for WRR 1 at the port level is set to 60 Mb/s.

Expected rates for combined scheduling classes 5 and 4 of each subscriber are:

$$\frac{60 \text{ Mb/s}}{3} = 20 \text{ Mb/s}$$

When 20 Mb/s is distributed between the two scheduling classes in the 2:1 ratio, each subscriber should receive 13.34 Mb/s on scheduling class 5 and 6.66 Mb/s on scheduling class 4, for the total of 20 Mb/s (out of offered 30 Mb/s).

The measured results in [Table 23: Input and output rates throughout the subscriber QoS hierarchy](#) are 13 Mb/s on scheduling class 5 and 6.5 Mb/s on scheduling class 4 for traffic streams 4 and 5 of each individual subscriber. This is in line with the expected results (the slight difference is due to the measured L2 rates by the traffic generator and enforced on-the-wire rates at the port level).

Scheduling class 3 (streams 1-3, 2-3, and 3-3)

Traffic streams mapped to scheduling class 3 do not have any rate restriction at the scheduling class level. Those streams can be only limited by the congestion at the subscriber aggregate level, the HS secondary shaper aggregate level, or the port aggregate level. Because the total amount of traffic so far is below the congestion level at each point in the hierarchy, it is expected that traffic stream 3 flows unimpeded. Consequently, each subscriber should receive the full input rate of 20 Mb/s on scheduling class 3. The actual results in [Table 23: Input and output rates throughout the subscriber QoS hierarchy](#) are aligned with the expected results.

To confirm that there is no congestion at the aggregate level so far for subscribers, HS secondary shapers, and the port, a calculation shows that each subscriber has received an equal amount of bandwidth so far:

46 Mb/s. This is below the configured aggregate rate limits at the subscriber, HS secondary scheduler, and port levels:

- 46 Mb/s is below the configured limit of 50 Mb/s for sub-1.
- 46 Mb/s is below the configured limit of 75 Mb/s for sub-2.
- 46 Mb/s is below the configured limit of 100 Mb/s for sub-3.
- Sub-1 and sub-2 compete for the bandwidth at HS secondary shaper 1, and their combined rate of 92 Mb/s is below the configured aggregate rate of HS secondary shaper 1 (120 Mb/s).
- Sub-3 with its 46 Mb/s is below its configured aggregate subscriber rate limit (100 Mb/s) and that of the HS secondary shaper 2 (150 Mb/s).
- The total subscriber bandwidth of $3 \times 46 \text{ Mb/s} = 138 \text{ Mb/s}$ is below the 200 Mb/s aggregate rate limit at the port level.

Scheduling classes 2 and 1 (streams 1-2, 1-1, 2-2, 2-1, 3-2, 3-1) - the lowest priority scheduling classes

The total amount of offered (input) traffic for lowest priority scheduling classes 1 and 2 across all three subscribers is 230 Mb/s (50 Mb/s for sub-1, 70 Mb/s for sub-2, and 100 Mb/s for sub-3). This amount of traffic will cause congestion at the aggregate level for sub-1, at the aggregate level of HS secondary shaper 1, and at the aggregate port level. The amount of traffic that each subscriber will receive for those two scheduling classes will now differ (up to this point, each subscriber received an equal amount of traffic).

Sub-1 will become limited by 50 Mb/s (the L2 rate limit) of its configured aggregate rate limit, which has only ~4 Mb/s left (with ~46 Mb/s received for sub-1). Considering that the scheduling classes 2 and 1 are serviced in 4:1 fashion, the ~4 Mb/s should be split between the two scheduling classes. Scheduling class 2 should receive ~3.2 Mb/s while scheduling class 1 should receive ~0.8 Mb/s. The measured data in [Table 23: Input and output rates throughout the subscriber QoS hierarchy](#) shows that the expected and measured results on the traffic generator are aligned.

Up to this point, the spare capacity on the HS secondary shaper 1 is:

$$\begin{aligned}
 &120 \text{ Mb/s (HS secondary shaper 1 aggregate on-the-wire rate limit)} \\
 &- 51 \text{ Mb/s (on-the-wire traffic from sub-1 up to this point)} \\
 &- 47 \text{ Mb/s (on-the-wire traffic from sub-2 up to this point)} \\
 \hline
 &= 22 \text{ Mb/s}^{12} \text{ (capacity left on the HS secondary shaper 1 up to this point)}
 \end{aligned}$$

26901

while the spare capacity at the sub-2 aggregate level is:

$$\begin{aligned}
 &75 \text{ Mb/s (sub-2 L2 aggregate rate limit)} \\
 &- 46 \text{ Mb/s (L2 traffic from sub-2 up to this point)} \\
 \hline
 &= 29 \text{ Mb/s (L2 capacity left on the HS secondary shaper 1 up to this point)}
 \end{aligned}$$

26902

This means that the 70 Mb/s that sub-2 is sending on scheduling classes 2 and 1 will congest the HS secondary shaper 1 (~22 Mb/s left) before it congests the subscriber itself (~29 Mb/s left). The 22 Mb/s of on-the-wire bandwidth capacity left on the HS secondary shaper will be divided in 4:1 ratio between scheduling classes 2 and 1. Therefore, the scheduling class 2 should receive 17.6 Mb/s and the scheduling class 1 should receive 4.4 Mb/s. When converted to L2 rates, these expected results match the measured

rates by the traffic generator in [Table 23: Input and output rates throughout the subscriber QoS hierarchy](#) for sub-2.



Note:

1. At the subscriber level, rate enforcement is based on L2 packet size.
2. Spare capacity for HS secondary shaper 1 is calculated based on on-the-wire rates while the subscriber aggregate capacity is calculated in L2 rates, because the aggregate rate limit in the HS secondary shaper is configured in on-the-wire rates while the subscriber aggregate rates are configured in L2 rates.

Similar logic can be used for traffic on scheduling classes 2 and 1 for sub-3. The difference from the previous case is that this traffic will be limited by port congestion and not the HS secondary shaper congestion (or subscriber aggregate congestion).

The available bandwidth at the port level up to this point is:

$$\begin{array}{r}
 200 \text{ Mb/s (on-the-wire port aggregate rate cap)} \\
 - 120 \text{ Mb/s (on-the-wire traffic from the HS secondary shaper 1 to sub-1 and sub-2)} \\
 - \sim 47 \text{ Mb/s (on-the-wire traffic from sub-3)} \\
 \hline
 = 33 \text{ Mb/s (on-the-wire capacity left on the port)}
 \end{array}$$

26903

The 100 Mb/s sent to sub-3 on scheduling classes 2 and 1 will be reduced to 33 Mb/s on the output in 4:1 ratio between scheduling classes 2 and 1. Scheduling class 2 should receive 26.3 Mb/s and scheduling class 1 should receive 6.6 Mb/s. Converted to L2 rates, these numbers match the measured rates by the traffic generator for streams 3-2 and 3-1 in [Table 23: Input and output rates throughout the subscriber QoS hierarchy](#).

Conclusion

This chapter described traffic management capabilities and configuration of HSQ IOM with ESM in single SLA mode. Because HSQ performs egress traffic management functions, the focus in this chapter was on egress QoS. The ingress QoS remains the same as on other non-HSQ IOMs where a combination of queues and policers can be deployed. Ingress and egress queues on HSQ are allocated from different queue pools that are separated at the hardware level (by different chips).

Some configuration practices are summarized here:

- SAP queues in the ESM context should be removed in most subscriber/SAP (1:1) deployment scenarios (**profiled-only-traffic** command). This maximizes the number of ingress queues available for ESM, which is important because HSQ does not support shared queueing on ingress. The alternative to queueing is to use policers on ingress.
- Queue depth (buffer size) can be statically configured (**mbs** command). By default, this parameter (MBS) is set dynamically and is somewhat high. To reduce the amount of buffering and delay, a smaller buffer size can be provisioned per queue.
- The **low-burst-max-class** should be configured at the subscriber level and at the HS secondary scheduler level to ensure smaller buffering delays for higher priority traffic.
- The delta between the high and low thresholds (**hs-fixed-high-thresh-delta** command) should be configured to accommodate a reasonable number of packets on higher priority classes that are serviced after the packets on lower priority classes have stopped due to crossing of the lower burst threshold. In

this example, the value of 12 kbytes means that about 12 packets (each of 1000 byte size) on higher scheduling classes will be serviced unimpeded before the higher burst threshold is reached, which stops the service.

HSQ IOM provides scalable traffic management functions on egress. A high number of egress queues can be serviced in strict priority fashion and extensive shaping hierarchy provides protection against overrunning the link capacities at an access node level or at a subscriber level. HSQ IOM is a card for a scaled ESM (in single SLA mode) environment that requires a higher number of queues per subscriber with throughput demands in the range of 200 Gb/s per HSQ IOM.

Appendix A — Generic ESM configuration

Sub-profile, SLA-profile, and MSAP-policy configuration is omitted because it is already described in the [Configuration](#) section.

ESM configuration in this example starts with a **subscriber-interface** configured in an IES context. Subscriber hosts are instantiated in IES 3 service, under the **group-interface "group-int-1"**, which is created under the **subscriber-interface "sub-int-1"**. Authentication and address assignment of the subscriber hosts is performed via LUDB user-db "user-db-1". The IP addresses that are assigned to the hosts are statically configured in LUDB (DHCP server is not used in this setup).

```
configure
service
  ies 3 name "3" customer 1 create
  subscriber-interface "sub-int-1" create
  address 10.10.1.254/24
  group-interface "group-int-1" create
  dhcp
    proxy-server
      emulated-server 10.10.1.254
      no shutdown
    exit
    option
      action keep
      circuit-id
      remote-id
    exit
    trusted
    lease-populate 100
    gi-address 10.10.1.254
    user-db "user-db-1"
    no shutdown
  exit
exit
no shutdown
exit
exit
```

Subscriber SAPs are automatically created based on the VLAN tags carried in the initial control packets of the subscriber hosts. This VLAN auto-detection and SAP auto-creation is configured under the capture SAP hierarchy. The capture SAP is configured to support LUDB authentication for dynamic DHCPv4 host instantiation:

```
configure
service
  vpls 10 name "10" customer 1 create
  sap 3/1/1:*.* capture-sap create
```

```
        trigger-packet dhcp
        dhcp-user-db "user-db-1"
    exit
exit
```

Sub-ident-policy is a mandatory configuration in ESM. It determines the mapping method between the sub/SLA profiles and the corresponding strings obtained during the authentication phase for the subscriber. Subscribers strings obtained during the authentication phase point, in some form (determined by **sub-ident-policy**), to the configured sub/SLA profiles (in the SR OS node) that will be associated with the subscriber.

```
configure
  subscriber-mgmt
    sub-ident-policy "sub_ident_pol" create
    sub-profile-map
      use-direct-map-as-default
    exit
    sla-profile-map
      use-direct-map-as-default
    exit
  exit
exit
```

In this example, authentication of the subscriber hosts and IP address assignment is performed through LUDB. The hosts are identified based on the **circuit-id** and **remote-id** fields in DHCP control packets. Sub/SLA -profile strings in LUDB are directly mapped to the configured sub/SLA-profiles in the SR OS node. This direct mapping is implied by the **use-direct-map-as-default** command within the **sub-ident-policy**.

Service ID, group-interface name, and MSAP policy name for the subscriber is also determined during the authentication phase via LUDB. LUDB carries only ESM-specific configuration. There is no HSQ relevant configuration present in LUDB.

```
configure
  subscriber-mgmt
    local-user-db "user-db-1" create
    ipoe
      match-list circuit-id remote-id
      host "sub-1-host-1" create
        host-identification
          circuit-id string "sub-1"
          remote-id string "host-1"
        exit
      address 10.10.1.1
      identification-strings 254 create
        subscriber-id "sub-1"
        sla-profile-string "sla-1"
        sub-profile-string "sub-prof-1"
      exit
      msap-defaults
        group-interface "group-int-1"
        policy "msaps"
        service 3
      exit
      options
        subnet-mask 255.255.255.0
      exit
      no shutdown
    exit
    host "sub-2-host-1" create
      host-identification
        circuit-id string "sub-2"
        remote-id string "host-1"
```

```
    exit
    address 10.10.1.2
    identification-strings 254 create
        subscriber-id "sub-2"
        sla-profile-string "sla-2"
        sub-profile-string "sub-prof-2"
    exit
    msap-defaults
        group-interface "group-int-1"
        policy "msaps"
        service 3
    exit
    options
        subnet-mask 255.255.255.0
    exit
    no shutdown
exit
host "sub-3-host-1" create
    host-identification
        circuit-id string "sub-3"
        remote-id string "host-1"
    exit
    address 10.10.1.3
    identification-strings 254 create
        subscriber-id "sub-3"
        sla-profile-string "sla-3"
        sub-profile-string "sub-prof-3"
    exit
    msap-defaults
        group-interface "group-int-1"
        policy "msaps"
        service 3
    exit
    options
        subnet-mask 255.255.255.0
    exit
    no shutdown
exit
exit
no shutdown
exit
exit
```

Ingress Multicast Path Management

This chapter provides information about Ingress Multicast Path Management (IMPM).

Topics in this chapter include:

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

Applicability

The information and configuration in this chapter are based on SR OS Release 9.0.R6. There are no prerequisites for this configuration.

Overview

Ingress Multicast Path Management (IMPM) optimizes the IPv4 and IPv6 multicast capacity on the applicable systems with the goal of achieving the maximum system-wide IP multicast throughput. It controls the delivery of IPv4/IPv6 routed multicast groups and of VPLS (IGMP and PIM) snooped IPv4 multicast groups, which usually relate to the distribution of IP TV channels.

A description is also included of the use of IMPM resources by point-to-multipoint LSP IP multicast traffic, and policed ingress routed IP multicast or VPLS broadcast, unknown or multicast traffic. The system capacity for these traffic types can be increased even with IMPM disabled.

IMPM introduces the concept of paths on a line card (IOM/IMM) which connect to planes on the chassis switch fabric ([Figure 90: IOM/IMM Paths Connecting to Switch Fabric Planes](#)).

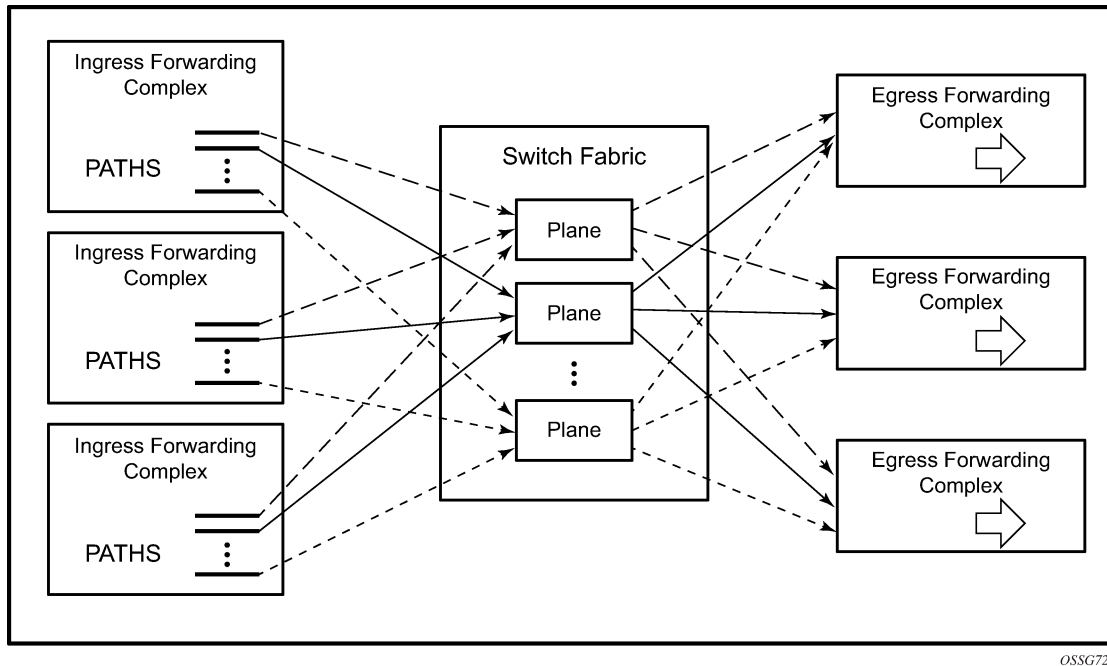
IMPM monitors the ingress rate of IP multicast channels (S,G multicast streams) on line card paths and optimizes the use of the capacity of each switch fabric plane. Its goal is to forward as many channels as possible through the system in order to make maximum use of the switch fabric planes without incurring multicast packet loss. IMPM achieves this by moving entire multicast channels between the line card paths, and therefore between switch fabric planes, to achieve an optimal packing of channels onto path/planes. These actions take into consideration the total ingress multicast traffic being received by all line cards with IMPM enabled and a configured preference of each channel.



Note:

S,G refers to an individual multicast stream by referencing the source (S) and multicast group (G) used by the stream.

Figure 90: IOM/IMM Paths Connecting to Switch Fabric Planes



OSSG725

There are three types of path: primary, secondary and ancillary paths (the ancillary path is specific to the IOM1/2 and is discussed in [Ancillary Path](#)).

When a new channel is received on a line card for which there is an egress join, its traffic is initially placed on to a secondary path by default. IMPM monitors the channel's traffic rate and, after an initial monitoring period, can move the channel to another path (usually a primary path) on which sufficient capacity exists for the channel. IMPM constantly monitors all of the ingress channels and therefore keeps a picture of the current usage of all the line card paths and switch fabric planes. As new channels arrive, IMPM assigns them onto available capacity, which may involve moving existing channels between paths (and planes). If a channel stops, IMPM will be aware that more capacity is now available. If the traffic rate of any channel(s) changes, IMPM will continue to optimize the use of the path/planes.

In the case where there is insufficient plane capacity available for all ingress channels, entire channel(s) are blackholed (dropped) rather than allowing a random degradation across all channels. This action is based on a configurable channel preference with the lowest preference channel being dropped first. If path/ plane capacity becomes available, then the blackholed channel(s) can be re-instated.

Paths and Planes

Each path connects to one plane on the switch fabric which is then used to replicate multicast traffic to the egress line cards. Further replication can occur on the egress line card but this is not related to IMPM so is not discussed.

Each plane has a physical connection to every line card and operates in full duplex mode allowing the possibility for traffic received on a plane from one line card path to be sent to every other line card, and back to the ingress line card (note that traffic is actually only sent to the line cards where it may exit). Therefore a given plane interconnects all line cards which allow ingress multipoint traffic from a line card with a path connected to this plane to be sent to multiple egress line cards.

Traffic could be sent by only one line card, or by multiple line cards simultaneously, on to a given plane. The total amount of traffic on a path or plane cannot exceed the capacity of that path or plane, respectively.

There could be more planes available on the switch fabrics than paths on the line cards. Conversely, there could be more total line card paths than planes available on the switch fabrics. In the latter case, the system distributes the paths as equally as possible over the available planes and multiple paths would be assigned to a given plane. Note that multiple paths of either type (primary or secondary) can terminate on a given plane.

The number of paths available per line card depends on the type of line card used whereas the number of planes on a system depends on the chassis type, the chassis mode (**a**, **b**, **c**, **d**) and the number of installed switch fabrics.

To clarify these concepts, consider a system with the following hardware installed.

```
A:PE-1# show card
=====
Card Summary
=====
```

Slot	Provisioned Card-type	Equipped Card-type	Admin State	Operational State	Comments
6	iom3-xp	iom3-xp	up	up	
7	imm8-10gb-xfp	imm8-10gb-xfp	up	up	
8	iom3-xp	iom3-xp	up	up	
A	sfm4-12	sfm4-12	up	up/active	
B	sfm4-12	sfm4-12	up	up/standby	

```
=====
A:PE-1#
```

Output 1 shows the mapping of paths to switch fabric planes.

```
A:PE-1# show system switch-fabric high-bandwidth-multicast
=====
Switch Fabric
=====
```

Slot/Mda	Cap:				Planes:															
	Min	Max	Hbm	Grp	Hi	Lo														
6/1	100%	100%	No	0	1	0	3	4	5	6	7	8	9	10	11	12	13	14	15	16
7/1	100%	100%	No	0	19	17	20	21	22	23	24	25	26	27	28	29	30	31	32	33
8/1	100%	100%	No	0	35	34	36	37	38	39	40	41	42	43	44	45	46	47	0	1
A	100%	100%	No	0	2	2														
B	100%	100%	No	0	2	2														

```
=====
A:PE-1#
```

Output 1: Paths and Planes in Chassis Mode d

This system has two SF/CPM4s and is using chassis mode **d**, this creates 24 planes per SF/CPM4 to give a total of 48 planes which are numbered 0-47. The IOM3-XP/IMMs have 16 paths each which are connected to different planes. The SF/CPM4s together use a single plane and an additional plane (18, which is not in the output above) is used by the system itself. As there are more paths (3x16=48) in this configuration than available planes (48-2[system planes 2,18]=46), some planes are shared by multiple paths, namely planes 0 and 1. Note that the path to plane mapping can change after a reboot or after changing hardware.

The following output shows the equivalent information if an IOM2 is added to this configuration in slot 5. In order for the IOM2 to be recognized, the system must be changed to use chassis mode **a**, **b** or **c**.

```
A:PE-1# show card
=====
Card Summary
=====
```

Slot	Provisioned Card-type	Equipped Card-type	Admin State	Operational State	Comments
5	iom2-20g	iom2-20g	up	up	
6	iom3-xp	iom3-xp	up	up	
7	imm8-10gb-xfp	imm8-10gb-xfp	up	up	
8	iom3-xp	iom3-xp	up	up	
A	sfm4-12	sfm4-12	up	up/active	
B	sfm4-12	sfm4-12	up	up/standby	

```
=====
A:PE-1#
```

The following output shows the mapping of the line card paths to the switch fabric planes with the IOM2 installed.

```
A:PE-1# show system switch-fabric high-bandwidth-multicast
=====
Switch Fabric
=====
```

Slot/Mda	Cap:				Planes:																	
	Min	Max	Hbm	Grp	Hi	Lo																
5/1	100%	100%	No	0	1	0																
5/2	100%	100%	No	0	4	3																
6/1	100%	100%	No	0	6	5	7	8	9	10	11	12	13	14	15	0	1	3	4	5		
7/1	100%	100%	No	0	7	6	8	9	10	11	12	13	14	15	0	1	3	4	5	6		
8/1	100%	100%	No	0	8	7	9	10	11	12	13	14	15	0	1	3	4	5	6	7		
A	100%	100%	No	0	2	2																
B	100%	100%	No	0	2	2																

```
=====
A:PE-1#
```

Output 2: Paths and Planes in Chassis Mode a/b/c

Now that the system is not in chassis mode **d**, in fact it is in mode **a** (but the output would be the same in modes **b** or **c**) the SF/CPM4s each create 8 planes giving a total of 16, numbered 0-15. One plane (2) is used by the SF/CPM4s, leaving 15 (0-1,3-15) planes for connectivity to the line card paths. Each IOM2 forwarding complex has 2 paths, so the paths of the IOM2 in slot 5 are using planes 0 and 1, and 3 and 4. Note that there are now fewer planes available and more paths, so there is more sharing of planes between paths than when chassis mode **d** was used.

IMPM Managed Traffic

IMPM manages IPv4/IPv6 routed multicast traffic and VPLS (IGMP and PIM) snooped IPv4 multicast traffic, traffic that matches a <*,G> or a <S,G> multicast record in the ingress forwarding table. It manages IP multicast traffic on a bud LSR when using point-to-multipoint (P2MP) LSPs but it does not manage IP protocol control traffic or traffic using multipoint-shared queuing. Traffic being managed by IMPM involves IMPM monitoring and potentially moving the related channels between paths/planes. The unmanaged traffic rates are also monitored and taken into account in the IMPM algorithm.

Care should be taken when using the mrouter-port configuration in a VPLS service. This creates a (*,*) multicast record and consequently all multicast channels that are not delivered locally to a non-mrouter port will be treated by IMPM as a single channel.

Configuration

This section covers:

- [IMPM on an IOM3-XP/IMM](#)
- [IMPM on an IOM1/2](#)
- [IMPM Not Enabled](#)

Prerequisites

As IMPM operates on IPv4/IPv6 routed or VPLS IGMP/PIM snooped IPv4 multicast traffic, some basic multicast configuration must be enabled. This section uses routed IP multicast in the global routing table which requires IP interfaces to be configured with PIM and IGMP. The configuration uses a PIM rendezvous point and static IGMP joins. The following is an example of the complete configuration of one interface.

```
configure
router
  interface "int-IOM3-1"
    address 172.16.6.254/24
    port 6/2/1
  exit
  igmp
    interface "int-IOM3-1"
      static
        group 239.255.0.1
        starg
      exit
    exit
  exit
  no shutdown
exit
pim
  interface "int-IOM3-1"
    exit
    rp
      static
        address 192.0.2.1
        group-prefix 239.255.0.0/16
      exit
    exit
  exit
  no shutdown
exit
exit
```

One interface is configured on each line card configured in the system, as shown in the following output, but omitting their IGMP and PIM configuration.

```
configure
```



```
router
  interface "int-IMM8"
    address 172.16.3.254/24
    port 7/2/1
  exit
  interface "int-IOM2"
    address 172.16.1.254/24
    port 5/2/1
  exit
  interface "int-IOM3-1"
    address 172.16.2.254/24
    port 6/2/1
  exit
  interface "int-IOM3-2"
    address 172.16.4.254/24
    port 8/2/1
  exit
exit
```

Configuring IMPM

The majority of the IMPM configuration is performed under the **mcast-management** CLI nodes and consists of:

1. The bandwidth-policy for characteristics relating to the IOM/IMM paths. This is applied on an IOM3-XP/IMM fp (fp is the system term for a forwarding complex on an IOM3-XP/IMM.), or an IOM1/2 MDA, under ingress mcast-path-management, with a bandwidth-policy named default being applied by default.

- IOM1/2

```
config# card slot-number mda mda-slot
  ingress
    mcast-path-management
      bandwidth-policy policy-name
```

- IOM3-XP/IMM

```
config# card slot-number fp [1]
  ingress
    mcast-path-management
      bandwidth-policy policy-name
```

2. The multicast-info policy for information related to the channels and how they are handled by the system. To facilitate provisioning, parameters can be configured under a three level hierarchy with each level overriding the configuration of its predecessor:

- Bundle: a group of channels
- Channel: a single channel or a non-overlapping range of channels
- Source-override: channels from a specific sender

```
config# mcast-management multicast-info-policy policy-name [create]
  bundle bundle-name [create]
    channel ip-address [ip-address] [create]
      source-override ip-address [create]
```

This policy is applied where the channel enters the system, so under router or service (vpls or vprn); the latter allows the handling of channels to be specific to a service, even if multiple services use overlapping channel addresses.

```
config# router multicast-info-policy policy-name

config# service vpls service-id multicast-info-policy policy-name

config# service vprn service-id multicast-info-policy policy-name
```

A default multicast-info-policy is applied to the above when IMPM is enabled.

3. The chassis-level node configures the information relating to the switch fabric planes.

```
config# mcast-management chassis-level
```

In addition, the command hi-bw-mcast-src (under an IOM3-XP/IMM fp or an IOM1/2 MDA) can be used to control the path to plane mapping among forwarding complexes.

IMPM on an IOM3-XP/IMM

IMPM is enabled on IOM3-XP/IMMs on under the card/fp CLI node as follows

```
config# card slot-number fp 1 ingress mcast-path-management no shutdown
```

IOM3-XP/IMM Paths

16 paths are available on an IOM3-XP/IMM when IMPM is enabled which can be either primary paths or secondary paths. By default the 16 paths are divided into 15 primary paths and 1 secondary path, as can be seen using the following command with IMPM enabled only on slot 6 (this corresponds to the plane assignment in Output 1):

```
*A:PE-1# tools dump mcast-path-mgr cpm
McPathMgr[6][0]: 0xf33b0a00
PATH:
PLANE:
Type SGs InUseBW AvailBW TotalBW ID SGs InUseBW AvailBW TotalBW
P 1 0 - - 1 1 0 2000000 2000000
P 1 0 - - 0 1 0 2000000 2000000
P 1 0 - - 3 1 0 2000000 2000000
P 1 0 - - 4 1 0 2000000 2000000
P 1 0 - - 5 1 0 2000000 2000000
P 1 0 - - 6 1 0 2000000 2000000
P 1 0 - - 7 1 0 2000000 2000000
P 1 0 - - 8 1 0 2000000 2000000
P 1 0 - - 9 1 0 2000000 2000000
P 1 0 - - 10 1 0 2000000 2000000
P 1 0 - - 11 1 0 2000000 2000000
P 1 0 - - 12 1 0 2000000 2000000
P 1 0 - - 13 1 0 2000000 2000000
P 1 0 - - 14 1 0 2000000 2000000
P 1 0 - - 15 1 0 2000000 2000000
S 1 0 - - 16 1 0 1800000 1800000
B 0 0 - - - - - - -
*A:PE-1#
```

Output 3: Paths/Planes on IOM3-XP/IMM

The left side of the output displays information about the paths (type {P=primary, s=secondary or B=blackholed}, number of "S,G"s and bandwidth in use (the path bandwidth cannot be set on an IOM3-XP/IMM, so the path available and total bandwidth always shows "-")) and the right side displays similar information about the associated planes (this will be a combination of the information for all paths connected to this plane). Note that one SG is always present on each path; this is used by the system and relates to the unmanaged traffic.

The primary/secondary paths are also highlighted in the planes section of Output 1, the primary paths being connected to the planes on the left of the "|" and the secondary paths to its right. There is a default primary path and a default secondary path; these correspond to the left-most plane and right-most plane for each line card, respectively.

Primary paths are used by:

- Expedited IES, VPLS and VPRN service ingress non-managed multipoint traffic (using the SAP based queues). This uses the default primary path.
- Expedited network ingress non-managed multipoint traffic (using the network interface queues). This uses the default primary path.
- Managed multicast explicit path primary channels (using the primary paths managed multipoint queue)
- All managed multicast dynamic path channels when the primary paths or multicast planes are not at their limit (using the primary paths managed multipoint queue)
- Highest preference managed multicast dynamic path channels when the primary paths or multicast planes are at their limit (using the primary paths managed multipoint queue)
- Non-managed P2MP LSP IP multicast traffic. This does not require IMPM to be enabled, so is discussed later in [IMPM Not Enabled](#).
- Non-managed expedited ingress policed multipoint traffic. This does not require IMPM to be enabled, so is discussed in [IMPM Not Enabled](#).

Secondary paths are used by:

- Best-Effort IES, VPLS and VPRN service ingress non-managed multipoint traffic (using the SAP based queues). This uses the default secondary path.
- Best-Effort network ingress non-managed traffic (using the network interface multipoint queues). This uses the default secondary path.
- Managed multicast explicit path secondary channels (using the secondary paths managed multipoint queue)
- Lower preference managed multicast dynamic path channels when the primary paths or multicast planes are at their limit (using the secondary paths managed multipoint queue)
- Non-managed best-effort ingress policed multipoint traffic. This does not require IMPM to be enabled, so is discussed in [IMPM Not Enabled](#).

When IMPM is enabled, the managed traffic does not use the standard multipoint queues but instead is placed onto a separate set of shared queues which are associated with the 16 paths. These queues are instantiated in an access ingress pool (called MC Path Mgmt, see "show output" section) which exists by default (this pool can be used even when IMPM is not enabled – see section "IMPM not enabled"). Statistics relating to traffic on these queues are reflected back to the standard ingress multipoint queues for accounting and troubleshooting purposes. Note that non-managed traffic continues to use the standard ingress multipoint queues, with the exception of P2MP LSP IP multicast traffic and policed multipoint traffic.

The size of the pool by default is 10% of the total ingress pool size, the reserved CBS is 50% of the pool and the default slope policy is applied. Care should be taken when changing the size of this pool as this would affect the size of other ingress pools on the line card.

```
config# mcast-management bandwidth-policy policy-name [create]
      mcast-pool percent-of-total percent-of-buffers
                  resv-cbs percent-of-pool
                  slope-policy policy-name
```

It is possible to configure the parameters for the queues associated with both the primary and secondary paths, and also the number of secondary paths available, within the bandwidth-policy.

```
config# mcast-management bandwidth-policy policy-name create
      t2-paths
        primary-path
          queue-parameters
            cbs percentage
            hi-priority-only percent-of-mbs
            mbs percentage
        secondary-path
          number-paths number-of-paths [dual-sfm number-of-paths]
          queue-parameters
            cbs percentage
            hi-priority-only percent-of-mbs
            mbs percentage
```

The number of primary paths is 16 minus the number of secondary paths (at least one of each type must exist). The number-paths parameter specifies the number of secondary paths when only one switch fabric is active, while the dual-sfm parameter specifies the same value when two switch fabrics are active.

Packets are scheduled out of the path/multicast queue as follows:

- Traffic sent on primary paths is scheduled at multicast high priority while that on secondary paths is scheduled at multicast low priority.
- For managed traffic, the standard ingress forwarding class/prioritization is not used, instead IMPM managed traffic prioritization is based on a channel's preference (described in [Channel Prioritization and Blackholing Control](#)). Egress scheduling is unchanged.

Congestion handling (packet acceptance into the path/multicast queue):

- For non-managed traffic, this is based on the standard mechanism, namely the packet's enqueueing priority is used to determine whether the packet is accepted into the path multipoint queue depending on the queue mbs/cbs and the pool shared-buffers/reserved-buffers/WRED.

For managed traffic, the congestion handling is based upon the channel's preference (described later) and the channel's cong-priority-threshold which is configured in the multicast-info-policy (here under a bundle).

```
config# mcast-management multicast-info-policy policy-name [create]
      bundle bundle-name [create]
        cong-priority-threshold preference-level
```

When the preference of a channel is lower than the cong-priority-threshold setting, the traffic is treated as low enqueueing priority, when it is equal to or higher than the cong-priority-threshold it is treated as high enqueueing priority. The default cong-priority-threshold is 4.

IOM3-XP/IMM Planes

The capacity per plane for managed traffic is by default 2Gbps for a primary path and 1.8Gbps for a secondary path. The logic behind a reduced default on the secondary is to leave capacity for new streams in case the default secondary is fully used by managed streams.

The plane capacities can be configured as follows, note that this command configures the plane bandwidth associated with primary/secondary paths as seen by each line card, the TotalBw on the right side of Output 3:

```
config# mcast-management chassis-level
      per-mcast-plane-limit megabits-per-second [secondary megabits-per-second]
                        [dual-sfm megabits-per-second [secondary-dual-sfm megabits-per-second]]
```

The first parameter defines the capacity for a primary path, the second a secondary path and the dual-sfm configures these capacities when two switch fabrics are active. The maximum plane capacity is 4Gbps but for the release used here it should only be configured on 7750 SR-12 or 7450 ESS-12 systems populated with SF/CPM4(s) and 100G FP IMMs; for all other hardware combinations the maximum should be 2Gbps. Note that secondary plane capacity cannot be higher than that of the primary plane.

These values can be tuned to constrain the amount of managed multicast traffic in favour of non-managed multicast and unicast traffic.

On the IOM3-XP/IMM line cards there is no separate control of the line card path capacity, the capacity is only constrained by the plane.

IOM3-XP/IMM Path to Plane Mapping

By default all fps (line cards for IOM3-XP/IMM) are configured into the default (zero) group as seen in Output 1 and the system distributes the paths as equally as possible over the available planes. This default works well if there is a low volume of multicast traffic (compared to the plane capacity), or if there is a higher volume multicast entering only one line card where the ingress capacity does not exceed that provided by the planes the line card is connected to.

If there are more paths than planes and, for example, there is a high bandwidth multicast channel entering two different line cards it could happen that both line cards select the same plane for two paths that are used. This would result in one of the channels being blackholed if the plane capacity is exceeded, effectively reducing the available multicast capacity from that line card. In order to avoid this situation, it is possible to configure the paths in to different groups and the system will attempt to use different planes for each group.

Output 1 and Output 2 show examples of how the paths are mapped to planes.

In both cases there are more paths than planes so some planes are shared by multiple paths. The following command provides control of this mapping.

```
config# card slot-number fp [1] hi-bw-mcast-src [alarm] [group group-id] [default-paths-only]
```

If an fp is configured into a non-zero group (range: 1 to 32), the system will attempt to assign dedicated planes to its paths compared to other line cards in different non-zero groups. This action is dependent on there being sufficient planes available. If two line cards are assigned to the same group, they will be assigned the same planes. The **default-paths-only** parameter performs the assignment optimization only for the default primary and secondary paths and is only applicable to IOM3-XP/IMMs. The **alarm** keyword causes an alarm to be generated if some planes are still shared with fps in a different group.

An example of the use of this command is shown later.

Note: When VPLS IGMP and PIM snooped traffic is forwarded to a spoke or mesh SDP, by default it is sent by the switch fabric to all line card forwarding complexes on which there is a network IP interface. This is due to the dynamic nature of the way that a spoke or mesh SDP is associated with one or more egress network IP interfaces. If there is an active spoke/mesh SDP for the VPLS service on the egress forwarding complex, the traffic will be flooded on that spoke/mesh SDP, otherwise it will be dropped on the egress forwarding complex. This can be optimized by configuring an inclusion list under the spoke or mesh SDP defining which MDAs this traffic should be flooded to.

```
config>service>vpls# [spoke-sdp|mesh-sdp] sdp-id:vc-id egress
    mfib-allowed-mda-destinations
        [no] mda mda-id
```

The switch fabric flooding domain for this spoke or mesh SDP is made up only of the MDAs that have been added to the list. An empty list implies the default behavior.

It is important to ensure that the spoke or mesh SDP can only be established across the MDAs listed, for example by using RSVP with an explicit path.

IMPM Operation on IOM3-XP/IMM

This section covers:

- [Principle of Operation](#)
- [Monitoring Traffic Rates](#)
- [Channel Prioritization and Blackholing Control](#)

Principle of Operation

Where IMPM is enabled, it constantly monitors the line cards for ingress managed traffic.

When a new channel arrives it will be placed by default on to the default secondary path. IMPM determines the ingress point for the channel and then monitors the traffic of the channel within its monitoring period in order to calculate the rate of the channel. The system then searches the multicast paths/planes attached to the line card for available bandwidth. If there is sufficient capacity on such a path/plane, the channel is moved to that plane. Planes corresponding to primary paths are used first, when there is no capacity available on any primary path/plane a secondary path/plane is used (unless the channel is explicitly configured onto a specific path type – see the following description).

If the required bandwidth is unavailable, the system will then look for any channels that ingress this or other line cards that could be moved to a different multicast plane in order to free up capacity for the new channel. Any channel that is currently mapped to a multicast plane available to the ingress line card is eligible to be moved to a different multicast plane.

If an eligible existing channel is found, whether on this or another line card, that existing channel is moved without packet loss to a new multicast plane. If necessary, this process can be repeated resulting in multiple channels being moved. The new multicast channel is then mapped to the multicast plane previously occupied by the moved channels, again this normally is using a primary path.

If no movable channel is found, then lower preference channel(s) on any ingress line card that share multicast planes with the ingress line card of the new channel can be blackholed to free up capacity for the new channel. It is also possible to both blackhole some channels and move other channels in order to

free up the required capacity. If no lower preference channel is found and no suitable channel moves are possible, the new channel will be blackholed.

If required, channels can be explicitly configured to be on either a primary or secondary path. This can be done for a bundle of channels, for example

```
config# mcast-management multicast-info-policy policy-name [create]
      bundle bundle-name [create]
      explicit-sf-path {primary|secondary|ancillary}
```

Note that the ancillary path is not applicable to the IOM3-XP/IMM line cards, however, it is discussed in the section relating to the IOM1/2. If a channel on an IOM3-XP/IMM is configured onto the ancillary path it will use a primary path instead.

One secondary path on an IOM3-XP/IMM is used as a default startup path for new incoming channels. If a large amount of new channel traffic could be received within the monitoring period, it is possible that the plane associated with the default secondary path is over loaded before IMPM has time to monitor the channels' traffic rate and move the channels to a primary path (and a plane with available capacity). This can be alleviated by configuring the following command:

```
config# mcast-management chassis-level
      round-robin-inactive-records
```

When round-robin-inactive-records is enabled, the system redistributes new channels (which are referenced by inactive S,G records) among all available line card multicast (primary, secondary) paths and their switch fabric planes.

Monitoring Traffic Rates

The monitored traffic rate is the averaged traffic rate measured over a monitoring period. The monitoring period used depends on the total number of channels seen by IMPM, the minimum is a 1 second interval and the maximum a 10 seconds interval.

The way in which the system reacts to the measured rate can be tuned using the following command:

```
config# mcast-management multicast-info-policy policy-name [create]
      bundle bundle-name [create]
      bw-activity {use-admin-bw|dynamic [falling-delay seconds]}
                  [black-hole-rate kbps]
```


The default is to use the dynamic bandwidth rate, in which case a channel's active traffic rate is determined based on the measured monitored rates. IMPM then makes a decision of how to process the channel as follows.

If the channel was un-managed, IMPM will attempt to place the channel on a path/plane with sufficient available bandwidth.

If the channel was already managed, IMPM determines the highest monitored traffic rate (within a given monitoring period) in a sliding window defined by the falling-delay. This highest chosen monitored rate is then used to re-assess the placement of the channel on the path/planes; this may cause IMPM to move the channel. This mechanism prevents the active rate for a channel being reduced due to a momentarily drop in traffic rate. The default value for falling-delay is 30 seconds, with a range of 10-3600 seconds.

The above logic is shown in [Figure 91: Dynamic Bandwidth Rate Management](#) (for simplicity, the falling-delay is exactly twice the monitoring period). It can be seen that the active rate used when the traffic rate decreases follows the highest monitored rate in any falling-delay period.

Figure 91: Dynamic Bandwidth Rate Management




By using the sliding window of monitored rate measurements in the dynamic bandwidth measurement mode, IMPM delays releasing capacity for a channel in its calculations when the channel's rate has been reduced. This allows IMPM to ignore temporary fluctuations in a channel's rate. It is possible to tune this for cases where the reduction in a channel's rate is large by using the falling-percent-reset parameter. The default for the falling-percent-reset is 50%. Setting this to 100% effectively disables it.

```
config# mcast-management bandwidth-policy policy-name create
       falling-percent-reset percent-of-highest
```

When the monitored rate falls by a percentage which is greater or equal to falling-percent-reset, the rate used by IMPM is immediately set to this new monitored rate. This allows IMPM to react faster to significant reductions in a channel's rate while at the same time avoiding too frequent reallocations due to normal rate fluctuations. An example of the falling-percent-reset is shown in [Figure 92: Falling-Percent-Reset](#). In the last two monitoring periods, it can be seen that the active rate used is equal to the monitored rate in the previous periods, and not the higher rate in the previous falling-delay window.

Figure 92: Falling-Percent-Reset



OSSG727


The rate management can be further tuned based on the expectation that the channel bandwidth will fluctuate around a given rate. When the bw-activity is set to use-admin-bw within the multicast-info-policy, the following parameters come into play.

```
config# mcast-management multicast-info-policy policy-name [create]
      bundle bundle-name [create]
      admin-bw kbps

config# mcast-management bandwidth-policy policy-name create
      admin-bw-threshold kilo-bits-per-second
```

IMPM will use the rate configured for the admin-bw if the monitored rate is above the admin-bw-threshold but below or equal to the admin-bw in the sliding window of the falling-delay. Whenever the monitored rate is below the admin-bw-threshold or above the admin-bw, IMPM uses the dynamic rate management mechanism. The admin-bw-threshold needs to be smaller than the admin-bw, with the latter being non-zero. This is shown in [Figure 93: Admin-Bw Rate Management](#) (for simplicity, the falling-delay is exactly twice the monitoring period). It can be seen that while the monitored rate stays between the admin-bw-threshold and the admin-bw, the active rate used is set to the admin-bw.

Figure 93: Admin-Bw Rate Management



OSSG728

Finally, IMPM also takes into consideration the unmanaged traffic rate on the primary and secondary paths associated with SAP/network interface queues when determining the available capacity on these paths/planes. This is achieved by constantly monitoring the rate of this traffic on these queues and including this rate in the path/plane capacity usage. IMPM must be enabled on the ingress card of the unmanaged traffic otherwise it will not be monitored.

Channel Prioritization and Blackholing Control

IMPM decides which channels will be forwarded and which will be dropped based on a configured channel preference. The preference value can be in the range 0-7, with 7 being the most preferred and the default value being 0.

When there is insufficient capacity on the paths/planes to support all ingress multipoint traffic, IMPM uses the channel preferences to determine which channels should be forwarded and which should be blackholed (dropped).

This is an important distinction compared to the standard forwarding class packet prioritization; by using a channel preference, an entire channel is either forwarded or blackholed, this allows IMPM to avoid congestion having a small impact on multiple channels at the cost of entire channels being lost.

The channel preference is set within the multicast-info-policy, for example at the bundle level, with the settable values being 1-7:

```
config# mcast-management multicast-info-policy policy-name [create]
      bundle bundle-name [create]
      preference preference-level
```

The channel preference is also used for congestion control in the line card path queues – see "congestion handling" in the section on "IOM3-XP/IMM Paths" above.

Blackhole protection can also be enabled using the `bw-activity` command, shown above in the “Monitoring traffic rates” section. Regardless of which rate monitoring mechanism is used, a channel can be blackholed if the monitored rate exceeds the black-hole-rate, in which case the channel will be put immediately on the blackhole list and its packets dropped at ingress. This channel will no longer consume line card path or switch fabric capacity. The intention of this parameter is to provide a protection mechanism in case channels consume more bandwidth than expected which could adversely affect other channels.

The black-hole-rate can range from 0 to 40000000kbps, with no black-hole-rate by default. This protection is shown in the last monitoring period of both Figure 2 and Figure 4. Note that it will take a falling-delay period in which the channel's rate is always below the black-hole-rate in order for the channel to be re-instated unless the reduction in the rate is above the falling-percent-reset.

IMPM on an IOM1/2

As most of the principles when using IMPM on an IOM1/2 compared to on an IOM3-XP/IMM are the same and are described above, this section focuses only on the difference between the two.

Note that an IOM1 and IOM2 have two independent 10G forwarding complexes; in both cases there is a single MDA per forwarding complex, consequently some aspects of IMPM are configured under the `mda` CLI node.

IMPM is enabled on an IOM1/2 under the MDA CLI node as follows:

```
config# card slot-number mda mda-slot ingress mcast-path-management no shutdown
```

IOM1/2 Paths

Each forwarding complex has three paths: one primary and one secondary path, and another type called the ancillary path which is IOM1/2 specific. The paths can be seen using the following command with IMPM enabled only on slot 5 MDA 1 and MDA 2 (referenced as [0] and [1] respectively):

```
A:PE-1# tools dump mcast-path-mgr cpm
McPathMgr[5][0]: 0xf33b0a00
PATH:
Type SGs InUseBW AvailBW TotalBw ID SGs InUseBW AvailBW TotalBw
P 1 0 2000000 2000000 1 1 0 2000000 2000000
S 1 0 1500000 1500000 0 1 0 1800000 1800000
A 0 0 5000000 5000000 - - - - -
B 0 0 - - - - - - -
McPathMgr[5][1]: 0xf33b3198
PATH:
Type SGs InUseBW AvailBW TotalBw ID SGs InUseBW AvailBW TotalBw
P 1 0 2000000 2000000 4 1 0 2000000 2000000
S 1 0 1500000 1500000 3 1 0 1800000 1800000
A 0 0 5000000 5000000 - - - - -
B 0 0 - - - - - - -
A:PE-1#
```

Output 4: Paths/Planes on IOM1/2

The primary and secondary paths function as on the IOM3-XP/IMM, specifically for:

- Traffic usage.
- Associated queues instantiated in the ingress “MC Path Mgmt” ingress pool.

- Packet scheduling.
- Congestion handling.

The queue parameters can be configured within the bandwidth-policy in a similar way to the IOM3-XP/IMM (note that the IOM3-XP/IMM equivalent for this is under the t2-paths CLI node). The bandwidth-policy is then applied under the respective MDA.

```
config# mcast-management bandwidth-policy policy-name create
  primary-path
    queue-parameters
      cbs percentage
      hi-priority-only percent-of-mbs
      mbs percentage
  secondary-path
    queue-parameters
      cbs percentage
      hi-priority-only percent-of-mbs
      mbs percentage
```

The IOM1/2 allows capacity control on the paths themselves, which is not possible on the IOM3-XP/IMM. This is achieved using the following commands.

```
config# mcast-management bandwidth-policy policy-name create
  primary-path
    path-limit megabits-per-second
  secondary-path
    path-limit megabits-per-second
```

The maximum path-limit for both the primary and secondary path is 2Gbps with a default of 2Gbps for the primary path and 1.5Gbps for the secondary path. The capability to set a path limit for the IOM1/2 can be seen when comparing Output 3 with Output 4; in the latter the "AvailBW" and "TotalBw" for the "PATH" shows the path limit.

In addition to setting the path limits in the bandwidth-policy, they can also be overridden on a given MDA.

```
config# card slot-number mda mda-slot
  ingress
    mcast-path-management
      primary-override
        path-limit megabits-per-second
      secondary-override
        path-limit megabits-per-second
```

The achievable capacity will be the minimum of the path's path-limit and the plane's per-mcast-plane-limit.

Ancillary Path

The ancillary path

The ancillary path allows managed multicast traffic to be forwarded through the switch fabric as unicast and so is not constrained to the path or plane capacities. This is achieved using ingress replication, in order to send a channel to multiple destination forwarding complexes (DFCs), the ingress forwarding complex creates and forwards one copy of each packet to each DFC connected to the switch fabric.

However, the total replication capacity available for the ancillary path is constrained to 5G to prevent it impacting the unicast or primary/secondary path capacities. This means that the total amount of ancillary

capacity usable can be calculated from (note that the first copy sent is not included in this capacity, hence the "-1"):

$5\text{Gbps}/(\text{number_of_switch_fabric_DFCs} - 1)$

Taking an example shown later, if some channels ingress an IOM2 and egress 2 IOM3-XP (1 DFC each) and 1 IMM8 (1 DFC) to give a total of 3 egress DFCs, then total ancillary capacity available is

$5\text{Gbps}/(3-1) = 2.5\text{Gbps}$.

This would allow, for example, approximately 250 channels at 10Mbps each to use the ancillary path.

Due to the relationship between ancillary capacity and number of DFCs, the system will prefer the ancillary path as default whenever a channel enters an IOM1/2 and egresses on up to 3 DFCs. If there are 4 or more egress DFCs for the channel, then the primary path is preferred. The determination is performed on a per channel basis.

The configuration parameters relating to the primary and secondary paths are also available for the ancillary path.

```
config# mcast-management bandwidth-policy policy-name create
  ancillary-path
    queue-parameters
      cbs percentage
      hi-priority-only percent-of-mbs
      mbs percentage

config# mcast-management bandwidth-policy policy-name create
  ancillary-path
    path-limit megabits-per-second

config# card slot-number mda mda-slot
  ingress
    mcast-path-management
      ancillary-override
        path-limit megabits-per-second
```

IOM1/2 Planes

The capacity per plane for managed traffic is by default 2Gbps for a primary path and 1.8Gbps for a secondary path. Note that the default IOM1/2 secondary path limit is 1.5Gbps. A maximum of 2Gbps should be configured for either path type using the per-mcast-plane-limit (as shown for the IOM3-XP/IMM) when an IOM1/2 is being used with IMPM.

As the ancillary path does not use the switch fabric planes, there is no associated plane limit.

IOM1/2 Path to Plane Mapping

The hi-bw-mcast-src command function is the same for IOM1/2 line cards as for IOM3-XP/IMM line cards, as described above.

IMPM Operation on IOM1/2

This is exactly the same as the operation as for the IOM3-XP/IMM, see above.

IMPM Not Enabled

When IMPM is not enabled most multipoint traffic on an IOM1/2 and IOM3-XP/IMM can use only one primary path and one secondary path per forwarding complex. When ingress multipoint arrives it is placed on a multipoint queue and these queues are connected either to a primary path (if the queue is expedited) or a secondary path (if the queue is best-effort) depending on the ingress QOS classification applied. Standard ingress forwarding class/scheduling prioritization is used.

The capacity of the primary and secondary paths is 2Gbps, unless the system is a 7750 SR-12 or 7450 ESS-12 populated with SF/CPM4(s) and 100G FP IMM in which case the capacity is 4Gbps.

In Output 1 and Output 2, the primary path is associated with the left-most plane and the secondary path is associated with the right-most plane for each line card.

There are exceptions to the above on the IOM3-XP/IMM line cards for

- Point-to-multipoint LSP IP multicast traffic
- Policed ingress routed IP multicast or VPLS broadcast, unknown or multicast traffic

Point-to-Multipoint (P2MP) LSP IP Multicast Traffic

IMPM will manage traffic on a P2MP LSP for any IP multicast channel that is delivered locally, for example, the system is a bud LSR. However, non-managed P2MP LSP IP multicast traffic will also make use of the primary paths, regardless of whether IMPM is enabled or not.

For each primary queue created in the MC IMPM Path pool, an additional queue is created to carry non-managed P2MP LSP IP multicast traffic. The non-managed P2MP LSP IP multicast traffic is automatically distributed across all primary paths based on a modulo N function of the 10 least significant bits of channel destination group address, where N is the number of primary paths. Note that the number of primary paths can be changed with IMPM enabled or disabled by applying a bandwidth-policy which sets the number of secondary paths.

Policed Ingress Routed IP Multicast or VPLS Broadcast, Unknown Or Multicast Traffic

Routed IP multicast or VPLS broadcast, unknown or multicast traffic passing through ingress hardware policers on the IOM3-XP/IMM can also use the IMPM managed queues, with IMPM enabled or disabled. If this traffic is best-effort (forwarding classes BE, L2, AF, L1) it will use the secondary paths, if it is expedited (forwarding classes H2, EF, H1, NC) it will use the primary paths. Note that this traffic uses the shared ingress policer-output-queues which have a fixed forwarding class to queue mapping).

When IMPM is not enabled, this traffic is non-managed and 1 secondary path plus 15 primary paths are available (the default). Consequently, extra capacity is only available for the expedited traffic, which could use up to 15 planes worth of switch fabric capacity. If extra capacity is required for best-effort traffic, a bandwidth-policy allocating more secondary paths can be applied to the line card even without IMPM being enabled.

The policed ingress routed IP or VPLS broadcast, unknown or multicast traffic is distributed across the paths using a standard LAG hash algorithm (as described in the Traffic Load Balancing Options section in the *7450 ESS, 7750 SR, 7950 XRS, and VSR Interface Configuration Guide*).

For both of these exceptions, it is recommended to reduce the managed traffic primary/secondary plane limits (using per-mcast-plane-limit) in order to allow for the non-managed traffic.

Show Output

This section includes the show output related to IMPM. The first part covers generic output and uses IOM3-XP/IMMs and chassis mode **d**. The second part includes an IOM2 and so uses chassis mode **a**.

IOM3-XP/IMM and Generic Output

The system has an IOM3-XP in slots 6 and 8, with an IMM8 is slot 7.

```
A:PE-1# show card
=====
Card Summary
=====
```

Slot	Provisioned Card-type	Equipped Card-type	Admin State	Operational State	Comments
6	iom3-xp	iom3-xp	up	up	
7	imm8-10gb-xfp	imm8-10gb-xfp	up	up	
8	iom3-xp	iom3-xp	up	up	
A	sfm4-12	sfm4-12	up	up/active	
B	sfm4-12	sfm4-12	up	up/standby	

```
=====
A:PE-1#
```

The status of IMPM on a given card can be shown as follows:

```
*A:PE-1# show card 6 detail
=====
Card 6
=====
```

Slot	Provisioned Card-type	Equipped Card-type	Admin State	Operational State	Comments
6	iom3-xp	iom3-xp	up	up	

```
FP 1 Specific Data
  hi-bw-mc-srcEgress Alarm      : 2
  hi-bw-mc-srcEgress Group     : 0
  mc-path-mgmt Admin State      : In Service
  Ingress Bandwidth Policy      : default
```

IMPM is enabled on the fp, it is using the default bandwidth-policy and is using the default hi-bw-mcast-src group (0).

The MC Path Mgmt pool is created by default with the default settings.

```
*A:PE-1# show pools 6/1
=====
```

Type	Id	App.	Pool Name	Actual ResvCBS Admin ResvCBS	PoolSize
MDA	6/1	Acc-Ing	MC Path Mgmt	18816 50%	37632

```
=====
*A:PE-1#
```

The default bandwidth-policy can be shown, giving the default parameters for the MC Path Pool and the associated queues.

```
*A:PE-1# show mcast-management bandwidth-policy "default" detail
=====
Bandwidth Policy Details
=====
-----
Policy                : default
-----
Admin BW Thd          : 10 kbps          Falling Percent RST: 50
Mcast Pool Total      : 10              Mcast Pool Resv Cbs: 50
Slope Policy          : default
Primary
Limit                 : 2000 mbps        Cbs                 : 5.00
Mbs                   : 7.00            High Priority         : 10
Secondary
Limit                 : 1500 mbps        Cbs                 : 30.00
Mbs                   : 40.00           High Priority         : 10
Ancillary
Limit                 : 5000 mbps        Cbs                 : 65.00
Mbs                   : 80.00           High Priority         : 10
T2-Primary
Cbs                   : 5.00            Mbs                 : 7.00
High Priority          : 10
T2-Secondary
Cbs                   : 30.00           Mbs                 : 40.00
High Priority          : 10            Paths(Single/Dual) : 1/1
=====
Bandwidth Policies : 1
=====
*A:PE-1#
```

The defaults for the multicast-info-policy can be seen in configuration mode.

```
*A:PE-1# configure mcast-management
*A:PE-1>config>mcast-mgmt# info detail
-----
multicast-info-policy "default" create
no description
bundle "default" create
no cong-priority-threshold
no description
no ecmp-opt-threshold
no admin-bw
no preference
no keepalive-override
no explicit-sf-path
bw-activity dynamic falling-delay 30
no primary-tunnel-interface
exit
exit
```

The paths/planes on an IOM3-XP/IMM can be shown here for card 6.

```
*A:PE-1# tools dump mcast-path-mgr cpm
McPathMgr[6][0]: 0xf33b0a00
PATH:
PLANE:
Type SGs      InUseBW  AvailBW  TotalBW  ID  SGs  InUseBW  AvailBW  TotalBW
P      1         0         -         -   4    1         0  2000000  2000000
P      1         0         -         -   3    1         0  2000000  2000000
P      1         0         -         -   5    1         0  2000000  2000000
```



```

P      1      0      -      -      6      1      0      2000000      2000000
P      1      0      -      -      7      1      0      2000000      2000000
P      1      0      -      -      8      1      0      2000000      2000000
P      1      0      -      -      9      1      0      2000000      2000000
P      1      0      -      -     10      1      0      2000000      2000000
P      1      0      -      -     11      1      0      2000000      2000000
P      1      0      -      -     12      1      0      2000000      2000000
P      1      0      -      -     13      1      0      2000000      2000000
P      1      0      -      -     14      1      0      2000000      2000000
P      1      0      -      -     15      1      0      2000000      2000000
P      1      0      -      -     16      1      0      2000000      2000000
P      1      0      -      -      0      1      0      2000000      2000000
S      1      0      -      -      1      1      0      1800000      1800000
B      0      0      -      -      -      -      -      -      -
*A:PE-1#

```

Notice the plane total bandwidth is by default 2000Mbps for the primary paths and 1800Mbps for the secondary path, as can also be seen using this output.

```

*A:PE-1# show mcast-management chassis
=====
Chassis Information
=====
BW per MC plane          Single SFM   Dual SFM
-----
Primary Path              2000        2000
Secondary Path            1800        1800
-----
MMRP Admin Mode           Disabled
MMRP Oper Mode             Disabled
Round Robin Inactive Records Disabled
=====
*A:PE-1#

```

The Round Robin Inactive Records is disabled. The MMRP (Multiple MAC Registration Protocol) modes relate to the use of the MC Path Mgmt queues for MMRP traffic. When this is enabled, normal IMPM behavior is suspended so it is not in the scope of this configuration note.

A single channel (239.255.0.2) is now sent into interface int-IOM3-1 on port 6/2/1 with static IGMP joins on interfaces int-IMM8, int-IOM3-1 and int-IOM3-2. The current forwarding rate can be seen.

```

*A:PE-1# show router pim group detail
=====
PIM Source Group ipv4
=====
Group Address       : 239.255.0.2
Source Address      : 172.16.2.1
RP Address          : 192.0.2.1
Flags               : spt, rpt-prn-des   Type           : (S,G)
MRIB Next Hop       : 172.16.2.1
MRIB Src Flags      : direct              Keepalive Timer Exp: 0d 00:03:14
Up Time             : 0d 00:00:16         Resolved By       : rtable-u

Up JP State         : Joined               Up JP Expiry       : 0d 00:00:00
Up JP Rpt           : Pruned              Up JP Rpt Override : 0d 00:00:00

Register State      : Pruned              Register Stop Exp  : 0d 00:00:59
Reg From Anycast RP: No

Rpf Neighbor        : 172.16.2.1
Incoming Intf       : int-IOM3-1
Outgoing Intf List  : int-IMM8, int-IOM3-1, int-IOM3-2

```

```

Curr Fwding Rate   : 9873.0 kbps
Forwarded Packets  : 18017
Forwarded Octets   : 24899494
Spt threshold      : 0 kbps
Admin bandwidth    : 1 kbps
Discarded Packets  : 0
RPF Mismatches     : 0
ECMP opt threshold : 7

```

From the two sets of output below it can be seen that this is using the default primary path and switch fabric plane 4.

```

*A:PE-1# tools dump mcast-path-mgr cpm
McPathMgr[6][0]: 0xf33b0a00
PATH:
Type SGs      InUseBW AvailBW TotalBW ID SGs      InUseBW AvailBW TotalBW
P    2        9895      -      -    4    2        9895 1990105 2000000
P    1         0      -      -    3    1         0 2000000 2000000
P    1         0      -      -    5    1         0 2000000 2000000
P    1         0      -      -    6    1         0 2000000 2000000
P    1         0      -      -    7    1         0 2000000 2000000
P    1         0      -      -    8    1         0 2000000 2000000
P    1         0      -      -    9    1         0 2000000 2000000
P    1         0      -      -   10    1         0 2000000 2000000
P    1         0      -      -   11    1         0 2000000 2000000
P    1         0      -      -   12    1         0 2000000 2000000
P    1         0      -      -   13    1         0 2000000 2000000
P    1         0      -      -   14    1         0 2000000 2000000
P    1         0      -      -   15    1         0 2000000 2000000
P    1         0      -      -   16    1         0 2000000 2000000
S    1         0      -      -    1    1         0 1800000 1800000
B    0         0      -      -    -    -         -      -      -

```

```

*A:PE-1#

```

```

*A:PE-1# show system switch-fabric high-bandwidth-multicast
=====
Switch Fabric
=====
Cap:      Planes:
Slot/Mda  Min  Max  Hbm Grp  Hi | Lo
-----
6/1      100% 100% No  0    4 3 5 6 7 8 9 10 11 12 13 14 15 16 0 | 1
7/1      100% 100% No  0   19 17 20 21 22 23 24 25 26 27 28 29 30 31 32 | 33
8/1      100% 100% No  0   35 34 36 37 38 39 40 41 42 43 44 45 46 47 0 | 1
A         100% 100% No  0    2 | 2
B         100% 100% No  0    2 | 2
=====

```

```

*A:PE-1#

```

The information about the channel can be seen using this command.

```

*A:PE-1# show mcast-management
channel [router router-instance|vpls service-id|service-name service-name]
[mda slot[/mda]]
[group ip-address [source ip-address]]
[path path-type]
[detail]

```

The output for the channel being sent is as follows.

```

*A:PE-1# show mcast-management channel

```

```
=====
Multicast Channels
=====
Legend : D - Dynamic E - Explicit
=====
Source Address      Slot/Cpx  Current-Bw  Path      D/E
Group Address      Highest-Bw Plane
-----
172.16.2.1         6/1      9873        Primary   D
239.255.0.2         9873        4
=====
Multicast Channels : 1
=====
*A:PE-1#
```

```
*A:PE-1# show mcast-management channel detail
=====
Multicast Channels
=====
-----
Source Address      : 172.16.2.1
Group Address       : 239.255.0.2
-----
Slot/Complex        : 6/1          Current Bw       : 9873 kbps
Dynamic/Explicit     : Dynamic      Current Path     : Primary
Oper Admin Bw       : 0 kbps       Current Plane    : 4
Ing last highest     : 9873        Preference      : 0
Black-hole rate      : None          Ing sec highest  : 9873
Time remaining       : 30 seconds  Blackhole        : No
=====
Multicast Channels : 1
=====
*A:PE-1#
```

The channel is using the dynamic bandwidth activity measurement and the current bandwidth, last highest and second last highest rates are shown (which are the same as this traffic is from a traffic generator).

The Time remaining is the time remaining in the current falling-delay period. This is reset to the falling-delay every time the last highest bandwidth gets updated, when it reaches zero the value of last highest bandwidth will be replaced with second highest bandwidth and the second highest bandwidth will be set to the value of current bandwidth.

The Oper Admin Bw displays the value used for the admin-bw for this channel.

A subset of this information can be seen using this tools command.

```
*A:PE-1# tools dump mcast-path-mgr channels slot 6
=====
Slot: 6 Complex: 0
=====
Source address      CurrBw  Plane  PathType  Path Pref
Group address      PathBw  Repl   Exp       BlkHoleBw
-----
172.16.2.1         9873    4      primary   0      0
239.255.0.2        9873    2      none      0
Unmanaged traffic  0       4      primary   0      8
slot: 6 cmplx: 0 path: 0  0       0      none      0
Unmanaged traffic  0       3      primary   1      8
slot: 6 cmplx: 0 path: 1  0       0      none      0
Unmanaged traffic  0       5      primary   2      8
slot: 6 cmplx: 0 path: 2  0       0      none      0
Unmanaged traffic  0       6      primary   3      8
```

```

slot: 6 cplx: 0 path: 3      0      0      none      0
Unmanaged traffic          0      7      primary    4      8
slot: 6 cplx: 0 path: 4      0      0      none      0
Unmanaged traffic          0      8      primary    5      8
slot: 6 cplx: 0 path: 5      0      0      none      0
Unmanaged traffic          0      9      primary    6      8
slot: 6 cplx: 0 path: 6      0      0      none      0
Unmanaged traffic          0     10      primary    7      8
slot: 6 cplx: 0 path: 7      0      0      none      0
Unmanaged traffic          0     11      primary    8      8
slot: 6 cplx: 0 path: 8      0      0      none      0
Unmanaged traffic          0     12      primary    9      8
slot: 6 cplx: 0 path: 9      0      0      none      0
Unmanaged traffic          0     13      primary   10      8
slot: 6 cplx: 0 path: 10     0      0      none      0
Unmanaged traffic          0     14      primary   11      8
slot: 6 cplx: 0 path: 11     0      0      none      0
Unmanaged traffic          0     15      primary   12      8
slot: 6 cplx: 0 path: 12     0      0      none      0
Unmanaged traffic          0     16      primary   13      8
slot: 6 cplx: 0 path: 13     0      0      none      0
Unmanaged traffic          0      0      primary   14      8
slot: 6 cplx: 0 path: 14     0      0      none      0
Unmanaged traffic          0      1      secondary  15      8
slot: 6 cplx: 0 path: 15     0      0      none      0
*A:PE-1#

```

The bandwidth activity monitoring is now changed to use an admin-bw of 12Mbps with a blackhole rate of 15Mbps.

```

*A:PE-1# configure mcast-management
*A:PE-1>config>mcast-mgmt# info
-----
bandwidth-policy "bandwidth-policy-1" create
  admin-bw-threshold 8000
exit
multicast-info-policy "multicast-info-policy-1" create
  bundle "default" create
  exit
  bundle "bundle-1" create
    channel "239.255.0.1" "239.255.0.16" create
    admin-bw 12000
    bw-activity use-admin-bw black-hole-rate 15000
  exit
exit
exit
-----
*A:PE-1>config>mcast-mgmt# exit all

*A:PE-1# show mcast-management channel group 239.255.0.2 detail
=====
Multicast Channels
=====
-----
Source Address      : 172.16.2.1
Group Address       : 239.255.0.2
-----
Slot/Complex        : 6/1          Current Bw         : 9873 kbps
Dynamic/Explicit    : Dynamic      Current Path        : Primary
Oper Admin Bw       : 12000 kbps   Current Plane       : 4
Ing last highest    : 12000        Preference          : 0
Black-hole rate     : 15000 kbps   Ing sec highest    : 12000
Time remaining      : 30 seconds   Blackhole           : No

```

```
=====
Multicast Channels : 1
=====
*A:PE-1#

*A:PE-1# tools dump mcast-path-mgr cpm
McPathMgr[6][0]: 0xf33b0a00
PATH:
Type SGs      InUseBW   AvailBW   TotalBW   PLANE:
P      2      12000    -         -      4      2      12000   1988000   2000000
P      1         0         -         -      3      1         0   2000000   2000000
P      1         0         -         -      5      1         0   2000000   2000000
P      1         0         -         -      6      1         0   2000000   2000000
P      1         0         -         -      7      1         0   2000000   2000000
P      1         0         -         -      8      1         0   2000000   2000000
P      1         0         -         -      9      1         0   2000000   2000000
P      1         0         -         -     10      1         0   2000000   2000000
P      1         0         -         -     11      1         0   2000000   2000000
P      1         0         -         -     12      1         0   2000000   2000000
P      1         0         -         -     13      1         0   2000000   2000000
P      1         0         -         -     14      1         0   2000000   2000000
P      1         0         -         -     15      1         0   2000000   2000000
P      1         0         -         -     16      1         0   2000000   2000000
P      1         0         -         -      0      1         0   2000000   2000000
S      1         0         -         -      1      1         0   1800000   1800000
B      0         0         -         -      -      -         -     -         -
*A:PE-1#
```

Now the system treats the channel as though it is using 12Mbps capacity even though its current rate has not changed.

If the rate is increased above the blackhole rate, the channel is blackholed and an alarm is generated.

```
*A:PE-1#
11 2011/10/21 01:40:13.21 UTC MINOR: MCPATH #2001 Base Black-hole-rate is reached
"Channel (172.16.2.1,239.255.0.2) for vRtr instance 1 slot/cplx 6/1 has been blackholed."
*A:PE-1# show mcast-management channel group 239.255.0.2 detail
=====
Multicast Channels
=====
-----
Source Address      : 172.16.2.1
Group Address       : 239.255.0.2
-----
Slot/Complex        : 6/1          Current Bw       : 19458 kbps
Dynamic/Explicit     : Dynamic      Current Path      : Blackhole
Oper Admin Bw        : 12000 kbps   Current Plane     : N/A
Ing last highest     : 19480        Preference       : 0
Black-hole rate      : 15000 kbps   Ing sec highest  : 19469
Time remaining       : 23 seconds   Blackhole        : Yes
=====
Multicast Channels : 1
=====
*A:PE-1#

*A:PE-1# tools dump mcast-path-mgr cpm
McPathMgr[6][0]: 0xf33b0a00
PATH:
Type SGs      InUseBW   AvailBW   TotalBW   PLANE:
P      1         0         -         -      4      1         0   2000000   2000000
P      1         0         -         -      3      1         0   2000000   2000000
P      1         0         -         -      5      1         0   2000000   2000000
P      1         0         -         -      6      1         0   2000000   2000000
P      1         0         -         -      7      1         0   2000000   2000000
```

```

P      1      0      -      -      8      1      0      2000000      2000000
P      1      0      -      -      9      1      0      2000000      2000000
P      1      0      -      -     10      1      0      2000000      2000000
P      1      0      -      -     11      1      0      2000000      2000000
P      1      0      -      -     12      1      0      2000000      2000000
P      1      0      -      -     13      1      0      2000000      2000000
P      1      0      -      -     14      1      0      2000000      2000000
P      1      0      -      -     15      1      0      2000000      2000000
P      1      0      -      -     16      1      0      2000000      2000000
P      1      0      -      -      0      1      0      2000000      2000000
S      1      0      -      -      1      1      0      1800000      1800000
B      1      19480      -      -      -      -      -      -      -
*A:PE-1#

```

The output displayed above shows an alarm generated for a channel being blackholed due to the channel rate reaching the configured black-hole-rate. The example output below is an alarm for a channel being blackholed due to insufficient bandwidth being available to it.

```

7 2011/10/22 21:53:43.54 UTC MINOR: MCPATH #2001 Base No bandwidth available
"Channel (172.16.2.1,239.255.0.2) for vRtr instance 1 slot/cplx 6/1 has been blackholed."

```

Note that the following alarm relates to a dummy channel used to account for the unmanaged traffic. However, this traffic is never actually blackholed.

```

6 2011/10/21 00:27:58.00 UTC MINOR: MCPATH #2002 Base
"Channel (0.0.0.0,0.6.0.0) for unknown value (2) instance 0 slot/cplx 6/1 is no longer being
blackholed."

```

Alarms are also generated when all paths of a given type reach certain thresholds.

For primary and secondary paths, the two path range limit thresholds are

- Full: less than 5% capacity is available

```

9 2011/10/24 22:53:27.02 UTC MINOR: MCPATH #2003 Base
"The available bandwidth on secondary path on slot/cplx 6/1 has reached its maximum limit."

```

- Not full: more than 10% of the path capacity is available.

```

10 2011/10/24 22:53:48.02 UTC MINOR: MCPATH #2004 Base
"The available bandwidth on secondary path on slot/cplx 6/1 is within range limits."

```

A maximum of one alarm is generated for each event (blackhole start/stop, path full/not full) within a 3 second period. So, for example, if multiple channels are blackholed within the same 3 second period only one alarm will be generated (for the first event).

The effect of using the hi-bw-mcast-src command is illustrated below. Firstly, line card 6 is configured into group 1.

```

*A:PE-1# configure card 6 fp hi-bw-mcast-src group 1 alarm
*A:PE-1# show system switch-fabric high-bandwidth-multicast
=====
Switch Fabric
=====
          Cap:          Planes:
Slot/Mda Min  Max  Hbm Grp  Hi | Lo
-----
6/1      100% 100% Yes  1   4 3 5 6 7 8 9 10 11 12 13 14 15 16 0 | 1
7/1      100% 100% No   0  19 17 20 21 22 23 24 25 26 27 28 29 30 31 32 | 33
8/1      100% 100% No   0  35 34 36 37 38 39 40 41 42 43 44 45 46 47 0 | 1

```

```
A      100% 100% No  0    2 | 2
B      100% 100% No  0    2 | 2
```

```
=====
*A:PE-1#
```

The plane assignment has not changed (though it is possible that the system could re-arrange the planes used by card 6) and there are still planes (0,1) shared between card 6 and card 8.

Now card 8 is configured into group 2 (note that IMPM is only enabled on card 6 here).

```
*A:PE-1# configure card 8 fp hi-bw-mcast-src group 2 alarm
*A:PE-1# show system switch-fabric high-bandwidth-multicast
```

```
=====
Switch Fabric
```

```
=====
Cap:          Planes:
Slot/Mda Min  Max  Hbm Grp Hi | Lo
-----
6/1      100% 100% Yes  1    4  3  5  6  7  8  9 10 11 12 13 14 15 16  0 |  1
7/1      100% 100% No   0   19 17 20 21 22 23 24 25 26 27 28 29 30 31 32 | 33
8/1      100% 100% Yes  2   35 34 36 37 38 39 40 41 42 43 44 45 46 47 17 | 19
A         100% 100% No   0    2 |  2
B         100% 100% No   0    2 |  2
=====
```

```
*A:PE-1#
```

There are no longer planes shared between cards 6 and 8.

If card 7 is configured into group 3, the following is seen.

```
*A:PE-1# configure card 7 fp hi-bw-mcast-src group 3 alarm
*A:PE-1#
7 2011/10/21 00:35:50.95 UTC MINOR: CHASSIS #2052 Base Mda 6/1
"Class MDA Module : Plane shared by multiple multicast high bandwidth taps"

8 2011/10/21 00:35:50.95 UTC MINOR: CHASSIS #2052 Base Mda 6/2
"Class MDA Module : Plane shared by multiple multicast high bandwidth taps"

9 2011/10/21 00:35:50.97 UTC MINOR: CHASSIS #2052 Base Mda 7/1
"Class MDA Module : Plane shared by multiple multicast high bandwidth taps"

10 2011/10/21 00:35:50.97 UTC MINOR: CHASSIS #2052 Base Mda 7/2
"Class MDA Module : Plane shared by multiple multicast high bandwidth taps"
```

```
*A:PE-1# show system switch-fabric high-bandwidth-multicast
```

```
=====
Switch Fabric
```

```
=====
Cap:          Planes:
Slot/Mda Min  Max  Hbm Grp Hi | Lo
-----
6/1      100% 100% Yes  1    4  3  5  6  7  8  9 10 11 12 13 14 15 16  0 |  1
7/1      100% 100% Yes  3   21 20 22 23 24 25 26 27 28 29 30 31 32 33  0 |  1
8/1      100% 100% Yes  2   35 34 36 37 38 39 40 41 42 43 44 45 46 47 17 | 19
A         100% 100% No   0    2 |  2
B         100% 100% No   0    2 |  2
=====
```

```
*A:PE-1#
```

There are insufficient planes to allow each card/group to have dedicated planes. Planes 0 and 1 are still shared between cards 6 and 7, generating the associated alarms.

A common example of the use of the **hi-bw-mcast-src** command would be when cards 6 and 8 have uplink ports on which high bandwidth multicast channels could be received. It would be desired to have these cards use different planes. To achieve this, card 7 could be configured into group 1, as follows.

```
*A:PE-1# configure card 7 fp hi-bw-mcast-src group 1 alarm
*A:PE-1# show system switch-fabric high-bandwidth-multicast
=====
Switch Fabric
=====
          Cap:          Planes:
Slot/Mda Min  Max  Hbm Grp  Hi | Lo
-----
6/1      100% 100% Yes  1    4 3 5 6 7 8 9 10 11 12 13 14 15 16 0 | 1
7/1      100% 100% Yes  1    4 3 5 6 7 8 9 10 11 12 13 14 15 16 0 | 1
8/1      100% 100% Yes  2   35 34 36 37 38 39 40 41 42 43 44 45 46 47 17 | 19
A         100% 100% No   0    2 | 2
B         100% 100% No   0    2 | 2
=====
*A:PE-1#
```

Now it can be seen that card 7 shares the same planes as card 6, but more importantly card 6 has no planes in common with card 8.

Note that when traffic is received on card 6, it will also be seen on the same plane (not path) on card 7. In the example below, traffic can be seen on plane 4 which is used by both cards 6 and 7, but only card 6 has non-zero InUseBW path capacity.

```
*A:PE-1# tools dump mcast-path-mgr cpm
McPathMgr[6][0]: 0xf33b0a00
PATH:
Type SGs      InUseBW  AvailBW  TotalBw  ID  SGs      InUseBW  AvailBW  TotalBw
P     2        9707      -        -    4    3        9707    1990293  2000000
P     1         0      -        -    3    2         0    2000000  2000000
...
McPathMgr[7][0]: 0xf33b3198
PATH:
Type SGs      InUseBW  AvailBW  TotalBw  ID  SGs      InUseBW  AvailBW  TotalBw
P     1         0      -        -    4    3        9707    1990293  2000000
P     1         0      -        -    3    2         0    2000000  2000000
```

When IMPM managed traffic is received on SAPs (in an IES, VPLS or VPRN service) it can be seen against a specific queue counter (Off. Managed) in the SAP stats. The following output shows where sap 7/2/1:3 belongs to a VPLS service using igmp-snooping. A similar counter is not available for policer statistics.

```
*A:PE-1# show service id 2 sap 7/2/1:3 stats
=====
Service Access Points(SAP)
=====
-----
Sap per Queue stats
-----
          Packets          Octets
Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio      : 0          0
Off. LoPrio      : 0          0
Dro. HiPrio      : 0          0
Dro. LoPrio      : 0          0
For. InProf      : 0          0
For. OutProf     : 0          0
```



```
Ingress Queue 11 (Multipoint) (Priority)
Off. HiPrio      : 0
Off. LoPrio      : 0
Off. Managed     : 149410      209771640
Dro. HiPrio      : 0
Dro. LoPrio      : 0
For. InProf      : 149410      209771640
For. OutProf     : 0
Egress Queue 1
For. InProf      : 0
For. OutProf     : 0
Dro. InProf      : 0
Dro. OutProf     : 0
```

*A:PE-1#

IOM1/2 Specific Output

The system is configured with the following cards and is in chassis mode **a**. As can be seen, an IOM2 is in slot 5.

A:PE-1# show card

Card Summary

Slot	Provisioned Card-type	Equipped Card-type	Admin State	Operational State	Comments
5	iom2-20g	iom2-20g	up	up	
6	iom3-xp	iom3-xp	up	up	
7	imm8-10gb-xfp	imm8-10gb-xfp	up	up	
8	iom3-xp	iom3-xp	up	up	
A	sfm4-12	sfm4-12	up	up/active	
B	sfm4-12	sfm4-12	up	up/standby	

A:PE-1#

IMPM is enabled on MDA 1 and 2 of the IOM2 in slot 5, with a primary, secondary and ancillary path.

A:PE-1# show mcast-management mda

MDA Summary

S/C	Policy	Type	In-use-Bw	Admin
5/1	default	Primary	0 Kbps	up
	default	Secondary	0 Kbps	up
	default	Ancillary	0 Kbps	up
5/2	default	Primary	0 Kbps	up
	default	Secondary	0 Kbps	up
	default	Ancillary	0 Kbps	up
6/2	default	Primary	0 Kbps	down
	default	Secondary	0 Kbps	down
	default	Ancillary	0 Kbps	down
7/1	default	Primary	0 Kbps	down
	default	Secondary	0 Kbps	down
	default	Ancillary	0 Kbps	down
7/2	default	Primary	0 Kbps	down

	default	Secondary	0 Kbps	down
8/2	default	Ancillary	0 Kbps	down
	default	Primary	0 Kbps	down
	default	Secondary	0 Kbps	down
	default	Ancillary	0 Kbps	down
=====				
A:PE-1#				

The path/plane usage can be shown.

```
*A:PE-1# show system switch-fabric high-bandwidth-multicast
=====
Switch Fabric
=====
Cap:          Planes:
Slot/Mda Min  Max  Hbm Grp  Hi | Lo
-----
5/1      100% 100% No  0    1 | 0
5/2      100% 100% No  0    4 | 3
6/1      100% 100% No  0    6 5 7 8 9 10 11 12 13 14 15 0 1 3 4 | 5
7/1      100% 100% No  0    7 6 8 9 10 11 12 13 14 15 0 1 3 4 5 | 6
8/1      100% 100% No  0    8 7 9 10 11 12 13 14 15 0 1 3 4 5 6 | 7
A         100% 100% No  0    2 | 2
B         100% 100% No  0    2 | 2
=====

*A:PE-1#

*A:PE-1# tools dump mcast-path-mgr cpm
McPathMgr[5][0]: 0xf33b0a00
PATH:
Type SGs      InUseBW  AvailBW  TotalBw  PLANE:
P      1          0    2000000    2000000  ID SGs      InUseBW  AvailBW  TotalBw
S      1          0    1500000    1500000  0  1          0    1800000    1800000
A      0          0    5000000    5000000  -  -          -          -          -
B      0          0          -          -          -  -          -          -          -
McPathMgr[5][1]: 0xf33b3198
PATH:
Type SGs      InUseBW  AvailBW  TotalBw  PLANE:
P      1          0    2000000    2000000  ID SGs      InUseBW  AvailBW  TotalBw
S      1          0    1500000    1500000  3  1          0    1800000    1800000
A      0          0    5000000    5000000  -  -          -          -          -
B      0          0          -          -          -  -          -          -          -
*A:PE-1#
```

The path range limit alarm thresholds for the ancillary path are

- Full: less than 2% capacity is available
- Not full: more than 4% of the path capacity is available.

A single channel (239.255.0.1) is now sent into interface int-IOM2 on port 5/2/1 with static IGMP joins on interfaces int-IMM8, int-IOM3-1 and int-IOM3-2. The current forwarding rate can be seen.

```
*A:PE-1# show router pim group 239.255.0.1 detail
=====
PIM Source Group ipv4
=====
Group Address      : 239.255.0.1
Source Address     : 172.16.1.1
RP Address         : 192.0.2.1
Flags              : spt, rpt-prn-des  Type          : (S,G)
MRIB Next Hop      : 172.16.1.1
MRIB Src Flags     : direct             Keepalive Timer Exp: 0d 00:02:44
```

```

Up Time           : 0d 00:07:45      Resolved By      : rtable-u
Up JP State       : Joined            Up JP Expiry     : 0d 00:00:00
Up JP Rpt         : Pruned            Up JP Rpt Override : 0d 00:00:00

Register State    : Pruned            Register Stop Exp : 0d 00:00:32
Reg From Anycast RP: No

Rpf Neighbor      : 172.16.1.1
Incoming Intf     : int-IOM2
Outgoing Intf List : int-IMM8, int-IOM3-1, int-IOM3-2

Curr Fwding Rate  : 9734.8 kbps
Forwarded Packets : 591874            Discarded Packets : 0
Forwarded Octets  : 817969868        RPF Mismatches    : 0
Spt threshold     : 0 kbps            ECMP opt threshold : 7
Admin bandwidth   : 1 kbps
-----

```

As there are only 3 (<4) DFCs, the ancillary path is used.

```

*A:PE-1# show mcast-management channel
=====
Multicast Channels
=====
Legend : D - Dynamic E - Explicit
=====
Source Address          Slot/Cpx  Current-Bw  Path      D/E
Group Address           Highest-Bw Plane
-----
172.16.1.1              5/2      9729       Ancillary D
239.255.0.1             9740     -
=====
Multicast Channels : 1
=====
*A:PE-1#

*A:PE-1# show mcast-management channel detail
=====
Multicast Channels
=====
-----
Source Address          : 172.16.1.1
Group Address           : 239.255.0.1
-----
Slot/Complex            : 5/2      Current Bw      : 9729 kbps
Dynamic/Explicit         : Dynamic   Current Path    : Ancillary
Oper Admin Bw            : 0 kbps    Current Plane   : N/A
Ing last highest         : 9740     Preference      : 0
Black-hole rate          : None      Ing sec highest : 9740
Time remaining           : 27 seconds Blackhole       : No
=====
Multicast Channels : 1
=====
*A:PE-1#

```

If another join caused this traffic to be switched via an additional DFC, the system would place the channel on the primary path.

The ancillary path being used can also be seen as follows.

```

*A:PE-1# tools dump mcast-path-mgr cpm
McPathMgr[5][0]: 0xf33b0a00

```

```

PATH:
Type SGs      InUseBW AvailBW TotalBW
P    1         0    2000000 2000000
S    1         0    1500000 1500000
A    0         0    5000000 5000000
B    0         0         -      -
McPathMgr[5][1]: 0xf33b3198
PATH:
Type SGs      InUseBW AvailBW TotalBW
P    1         0    2000000 2000000
S    1         0    1500000 1500000
A    1       19480  4980520 5000000
B    0         0         -      -
* A: PE-1#
PLANE:
ID SGs      InUseBW AvailBW TotalBW
1  1         0    2000000 2000000
0  1         0    1800000 1800000
-  -         -         -      -
-  -         -         -      -
PLANE:
ID SGs      InUseBW AvailBW TotalBW
4  1         0    2000000 2000000
3  1         0    1800000 1800000
-  -         -         -      -
-  -         -         -      -

```

Note that the bandwidth shown on the ancillary path on the second MDA is approximately two times that of the ingress traffic, this matches the algorithm described earlier for the ancillary path. This can also be seen in the next output, there is the original channel traffic plus two replications (Repl).

```

* A: PE-1# tools dump mcast-path-mgr channels
=====
Slot: 5 Complex: 0
=====
Source address      CurrBw  Plane PathType  Path Pref
Group address      PathBw  Repl  Exp      BlkHoleBw
-----
Unmanaged traffic   0       1     primary  0      8
slot: 5 cmplx: 0 path: 0    0       0     none     0
Unmanaged traffic   0       0     secondary 1      8
slot: 5 cmplx: 0 path: 1    0       0     none     0
=====
Slot: 5 Complex: 1
=====
Source address      CurrBw  Plane PathType  Path Pref
Group address      PathBw  Repl  Exp      BlkHoleBw
-----
172.16.1.1          9740    48    ancillary 16     0
239.255.0.1         19480   2     none      0
Unmanaged traffic   0       4     primary  0      8
slot: 5 cmplx: 1 path: 0    0       0     none     0
Unmanaged traffic   0       3     secondary 1      8
slot: 5 cmplx: 1 path: 1    0       0     none     0
* A: PE-1#

```

Conclusion

This chapter has described the configuration of Ingress Multicast Path Management which optimizes IPv4 and IPv6 multicast capacity to achieve the maximum system-wide IP multicast throughput. It can be used for both routed IPv4/IPv6 and VPLS (IGMP and PIM) snooped IPv4 multicast groups, which usually relate to the distribution of IP TV channels.

IPoE Sessions

This chapter describes IPoE Sessions.

Topics in this chapter include:

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

Applicability

The information and configuration in this chapter are based on SR OS Release 13.0.R7.

IPoE sessions require a Routed CO environment with Enhanced Subscriber Management (ESM) enabled.


Overview

Definitions

Grouping a number of subscriber hosts with different IP stacks belonging to the same end user device in a single IPoE session simplifies operations; see [Figure 94: IPoE Session](#).

In this way, IPoE sessions provide behavior similar to PPPoE sessions for authentication, mid-session changes, and accounting.

Figure 94: IPoE Session



25641

The hosts (IP stacks) associated with a single IPoE session share ESM data, such as the subscriber ID, the sub-profile, the SLA profile, and so on. The shared ESM data is fetched and cached when the first host for that session connects: only a single authentication is needed. This requires ESM to be enabled; see the ESM basics chapter for more information.

An IPoE session can have one IPoEv4 host, up to two IPoEv6 wan hosts (one DHCPv6 host and a SLAAC host), and one IPoEv6 pd host. Hosts with the same SAP, MAC address, and optionally the same circuit ID (CID) or remote ID (RID), are grouped in a single IPoE session, using that combination as a key to the IPoE session data.

Authentication occurs when the first host for that IPoE session is created. Subsequent hosts belonging to the same IPoE session do not require additional authentication. For instantiating these subsequent hosts, SR OS uses the cached ESM data that was fetched while authenticating and instantiating the first host.

Mid-session changes are typically triggered by RADIUS CoA or Diameter Gx RAR messages, and automatically apply to all hosts associated with the IPoE session. A re-authentication could also lead to policy changes, but the changes are triggered through host renewal messages.

As well as queue instance and host accounting, RADIUS session accounting can be enabled for IPoE sessions. An accounting session identity (ASID) is created when an IPoE session is started.

An IPoE session is created when the first host is created, and the IPoE session is deleted when the last host is deleted. IPoE session creation is always protocol triggered. IPoE session deletion is triggered through protocol (DHCPv4, DHCPv6 release, or expiration of a lease) or through policy (idle-timeout, session-timeout, clear command, and so on).

IPoE sessions require the Routed CO model, and are supported on regular as well as on capture and managed SAPs.

Trigger Packets

Unlike PPP, where PPP sessions have clear and unique triggers that start (PADS) and stop (PADT) a PPP session, IPoE does not have unique triggers that start and stop an IPoE session.

IPoE sessions are created when the following trigger packets on ESM-enabled SAPs are received:

- DHCPv4 discover/request
- DHCPv6 solicit/request (native)
- DHCPv6 solicit/request (single relay)
- DHCPv6 solicit/request (double relay)
- Router Solicitation

No IPoE sessions are created on reception of ARP requests or CoA messages.

IPoE Session Key

The key to the IPoE session data is a combination of the SAP ID, the MAC address, and optionally the CID or the RID, as defined by the ipoe-session-policy:


```
configure
  subscriber-mgmt
    ipoe-session-policy <pol-name>
      session-key sap mac [cid] [rid]
```

The ipoe-session-policy name *default* is reserved for future use.

The CID and RID are mutually exclusive; see [Figure 95: IPoE Session Key](#):

- The CID corresponds to DHCPv4 option 82, sub-option 1 (Circuit-ID) and to DHCPv6 option 18 (Interface-ID).
- The RID corresponds to DHCPv4 option 82, sub-option 2 (Remote-ID) and to DHCPv6 option 37 (the remote-id field of the Relay Agent Remote-ID, excluding the enterprise-number field).

Figure 95: IPoE Session Key



When an IPoE session trigger packet is received, the IPoE session key is validated, ensuring that no field is missing. For example, if the key requires the CID or RID, and a device connects without CID or RID, the IPoE session setup and the host setup fail. Therefore, the CID or RID should only be part of the key when all devices include this parameter in the trigger packets.

If no IPoE session exists for a session key derived from a trigger packet, an IPoE session is created. If an IPoE session exists, a new host is created and added to the existing IPoE session, on condition that the host type is compatible with the already associated host types.

IPoE Session Authentication

Authenticating IPoE sessions requires generic identification parameters, which must be supported in both IPv4 and IPv6, so some restrictions apply.

LUDB Authentication

When using an LUDB for IPoE session authentication, all of the host-identification criteria can be used, except for the following:

- Option 60 - DHCPv4 only

LUDB entries containing option 60 are ignored while scanning the LUDB for a matching entry.

AAA/RADIUS Authentication

When using AAA/RADIUS for IPoE session authentication, all username formats can be used, except for the following:

- dhcp-client-vendor-opts - DHCPv4 only
- mac-giaddr - DHCPv4 only
- ppp-user-name - PPP only

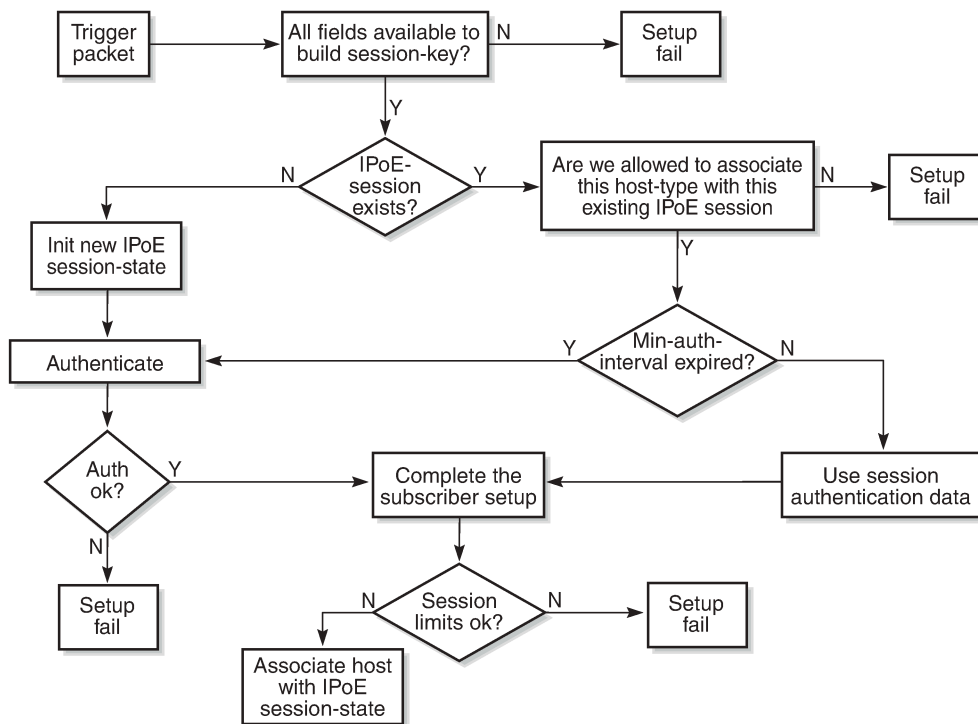
IPoE Session Creation

To ensure successful creation of an IPoE session, the following conditions must be met:

- the IPoE session key must be valid
- the session limits should not be exceeded
- an Accounting Session Identity (ASID) must be allocated
- the first ESM host must be successfully authenticated

Figure 96: IPoE Session Creation Flow shows the high-level flow of the IPoE session creation process. When a new trigger is received for an existing IPoE session, a new host is created and added to the session, on condition that the new host is compatible with the already associated hosts.


Figure 96: IPoE Session Creation Flow



25643

Figure 97: IPoE Session Creation via AAA/RADIUS shows an example where a device initiates DHCPv6 first. The data fetched while authenticating this device through AAA/RADIUS is cached. When the same device (leading to the same key) later initiates DHCPv4, the cached IPoE session data can be used, so no further AAA/RADIUS access is needed. The sequence could also first initiate DHCPv4, then DHCPv6. The result would be the same.

Figure 97: IPoE Session Creation via AAA/RADIUS



25680

An Access-Request message is sent to the AAA/RADIUS server for authentication. On successful authentication, the AAA/RADIUS returns the ESM data (subscriber ID, sub-profile, SLA profile, and so on) in an Access-Accept message.

The Access-Request message optionally includes the host ID or ASID.

The behavior when authenticating through LUDB is similar.

IPoE Session Re-Authentication

IPoE sessions can be re-authenticated, meaning that the ESM data is fetched again from the ESM data source (LUDB, RADIUS, and so on), and controlled through the min-auth-interval timer.

If a host renewal packet is received while the min-auth-interval timer is running for the corresponding IPoE session, the cached ESM data is used (with one exception: see forced authentication). If a host renewal packet is received when this timer has expired, re-authentication is performed. If re-authentication fails, the host renewal packet is dropped; see [Figure 96: IPoE Session Creation Flow](#).

IPoE session re-authentication is configured in the ipoe-session context using following command:

```

min-auth-interval ?
- min-auth-interval [days <days>] [hrs <hours>] [min <minutes>] [sec <seconds>]
- min-auth-interval infinite
- no min-auth-interval
<days>           : [0..365]
<hours>           : [0..23]
<minutes>         : [0..59]
<seconds>         : [0..59]
<infinite>        : keyword
    
```

By default, re-authentication is disabled by having the min-auth-interval set to infinite. Setting the min-auth-interval to zero will lead to every single message (DORA, SARR) triggering re-authentication, but that is not recommended.

IPoE session re-authentication can be used to implement dynamic policy changes. For alternatives also implementing dynamic changes, see the [Mid-Session Changes](#) section.

For IPoE sessions, the re-authentication option in the RADIUS authentication-policy context is ignored.

IPoE Session Forced Authentication

Forced authentication means that the ESM data is fetched again from the ESM data source, regardless of the value of the re-authentication timer.

By default, forced authentication occurs when the CID or RID in the trigger packet has changed value, but this behavior can be disabled.

An absent or empty CID or RID is not considered as a change.

Forced authentication is configured in the ipoe-session context using following command:

```
force-auth ?
- force-auth [cid-change] [rid-change]
- force-auth disabled
- no force-auth
<cid-change>      : keyword - ignore min-auth-interval when cid changed
<rid_change>      : keyword - ignore min-auth-interval when rid changed
<disabled>        : keyword - never ignore min-auth-interval
```

IPoE Session Deletion

When the last host associated with an IPoE session is deleted, the IPoE session is deleted.

IPoE sessions are forcibly deleted in following situations:

- group-interface ipoe-session shutdown
- clear service <id> ipoe-session
- session timeout
- RADIUS disconnect or Diameter Gx RAR
- Credit Control (out-of-credit)

In all these cases, all hosts belonging to that session are deleted, with one exception. When the SLAAC inactivity-timer expires, only the corresponding SLAAC host is deleted, not the remaining hosts. When this SLAAC host is the last host of the IPoE session, the IPoE session is deleted.

The IPoE session-timeout is configured in the ipoe-session-policy:

```
configure
  subscriber-mgmt
    ipoe-session-policy <pol-name>
      session-timeout <timeout>
```

The timeout value ranges from 1 to 31104000 seconds (360 days). By default, no session-timeout is specified.

When RADIUS or Diameter Gx returns the Session-Timeout [27] or the Alc-Relative-Session-Timeout [26-6527-160] attributes, these values are used and the behavior is the same as for PPP sessions.

When no Session-Timeout or Alc-Relative-Session-Timeout attribute is returned by RADIUS, the session-timeout as configured in the ipoe-session-policy is used.

A RADIUS disconnect message, even when targeted at a single host, will also lead to the deletion of the entire IPoE session including all associated hosts.

A shutdown in the ipoe-session context of the group interface results in the deletion of all its IPoE sessions and associated hosts.

Session and Host limits

The number of IPoE sessions on a group interface and on a SAP can be limited:

```
configure service ies|vprn <service-id>
  subscriber-interface <ip-int-name>
    group-interface <ip-int-name>
      ipoe-session
        session-limit      [1..max*]
        sap-session-limit [1..max*]
```

The default values for the session-limit and the sap-session-limit are unlimited (no session-limit) and 1, respectively.

For retail services, the IPoE session-limit is configured at the linked subscriber interface level:

```
configure service vprn <retail-service-id>
  subscriber-interface <RT-ip-int-name> fwd-service <WS-service-id>
    fwd-subscriber-interface <WS-ip-int-name>
      session-limit      [1..max*]
```

The default session-limit is unlimited.

Additionally, host limits can be imposed through the SLA profile:

```
configure subscriber-mgmt
  sla-profile <subscriber-profile-name>
    host-limits
      ipv4-arp          - Maximum number of IPv4 ARP hosts
      ipv4-dhcp          - Maximum number of IPv4 DHCP hosts
      ipv4-overall       - Maximum number of IPv4 hosts
      ipv4-ppp          - Maximum number of IPv4 PPP hosts
      ipv6-overall       - Maximum number of IPv6 hosts
      ipv6-pd-ipoe-dhcp  - Maximum number of IPv6-PD IPOE DHCP hosts
      ipv6-pd-overall    - Maximum number of IPv6-PD hosts
      ipv6-pd-ppp-dhcp   - Maximum number of IPv6-PD PPP DHCP hosts
      ipv6-wan-ipoe-dhcp - Maximum number of IPv6-Wan IPOE DHCP hosts
      ipv6-wan-ipoe-slaac - Maximum number of IPv6-Wan IPOE SLAAC hosts
      ipv6-wan-overall   - Maximum number of IPv6-Wan hosts
      ipv6-wan-ppp-dhcp  - Maximum number of IPv6-Wan PPP DHCP hosts
      ipv6-wan-ppp-slaac - Maximum number of IPv6-Wan PPP SLAAC hosts
      lac-overall        - Maximum number of L2TP LAC hosts
      overall            - Maximum number of hosts
      remove-oldest      - Remove oldest
    exit
```

See the [Wholesale/Retail](#) section for more information about limits in a wholesale/retail configuration.

RADIUS Session Accounting

As well as queue instance and host accounting, RADIUS session accounting can be enabled for IPoE sessions.

Usually, a RADIUS Accounting-Start message is sent when the first host is associated with an IPoE session. Regular and triggered accounting Interim-Update messages are sent during the IPoE session. An Accounting-Stop message is sent when the last host is deleted from the session.

Session accounting is configured in the RADIUS accounting policy, and can be set to the following values:

- session-accounting
- session-accounting interim-update
- session-accounting host-update
- session-accounting interim-update host-update

Plain session accounting sends start and stop messages. The RADIUS accounting server is informed about the start and the stop time of the session, but no counters are maintained. This implements time-based accounting.

The interim-update option additionally sends interim-update messages, so that the RADIUS accounting server maintains counters. This implements volume accounting.

The host-update option additionally sends host up/down event messages, so that the RADIUS accounting server keeps track of host creation and deletion events.

The combination of the interim-update and host-update options allows the RADIUS accounting server to track all changes.

Mid-Session Changes

Mid-session changes, such as those initiated via RADIUS CoA or Diameter Gx RAR messages, are applied to all hosts associated with the IPoE session. There is no way to update a single host of an IPoE session.

A RADIUS CoA message targeting any host of an IPoE session has the same effect as a RADIUS CoA message targeting the IPoE session using the IPoE session Acct-Session-Id as key. All hosts of the session are targeted and the session state data is updated.

Mid-session changes also can be applied manually, using the following command:

```
# tools perform subscriber-mgmt edit-ipoe-session sap <sap-id> mac <mac-address> [subscriber  
<sub-ident-string>] [sub-profile-string <sub-profile-string>] [sla-profile-string <sla-  
profile-string>] [inter-dest-id <intermediate-destination-id>] [ancp-string <ancp-string>]  
[app-profile-string <app-profile-string>] [circuit-id <circuit-id>] [remote-id <remote-id>]  
  
# tools perform subscriber-mgmt eval-ipoe-session [svc-id <service-id>] [sap <sap-id>] [mac  
<mac-address>] [circuit-id <circuit-id>] [remote-id <remote-id>] [subscriber <sub-ident-  
string>]
```

The tools commands eval-lease-state and eval-slaac-host are blocked when the host is part of a session.

IPoE session re-authentication can also lead to dynamic policy changes.

Subscriber Host Connectivity Verification

Subscriber host connectivity verification (SHCV) can be enabled for hosts associated with an IPoE session.

When a single host fails and stops responding to the SHCV messages, that host is deleted without affecting the other hosts that are part of the session. When the last host fails, the session is deleted.

IA-PD managed routes are not subject to SHCV, and cannot be removed because of SHCV.

Dual Homing

IPoE sessions are supported in a dual-homed environment, where Multichassis Synchronization (MCS) and Subscriber Routed Redundancy Protocol (SRRP) are active.

MCS ensures that the IPoE session data is synchronized between the BNG pair.

Wholesale/Retail

IPoE sessions are supported in single-homed and dual-homed wholesale/retail environments.

The wholesale IPoE session limit is configured at the group interface level, and the retail IPoE session limit is configured at the linked subscriber interface level:

```
configure service vprn <retail-service-id>
  subscriber-interface <RT-ip-int-name> fwd-service <WS-service-id>
    fwd-subscriber-interface <WS-ip-int-name>
      ipoe-session
        session-limit      [1..max*] #default unlimited
```

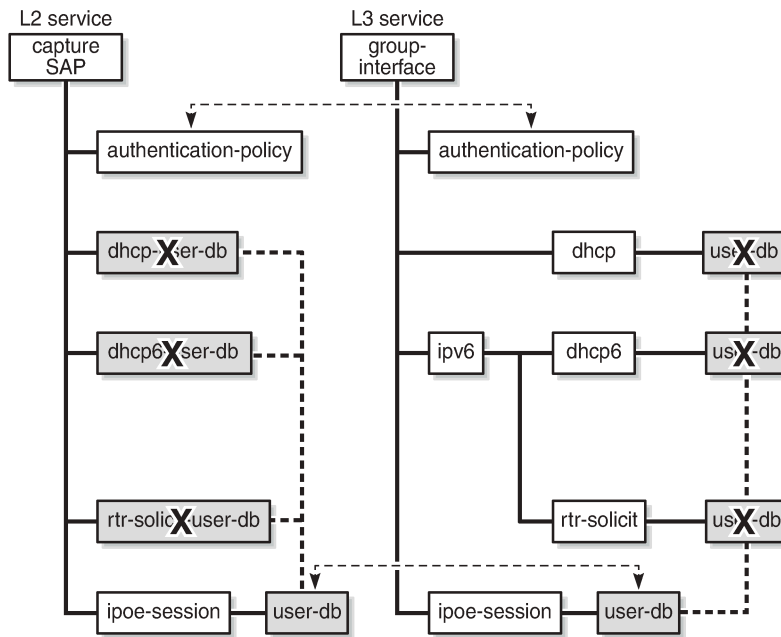
In IPoE, private-retail-subnets only apply to IPoEv6 in single-homed scenarios (no MCS). Therefore, the use cases for private-retail-subnets in combination with IPoE sessions are limited.

Practical Considerations

The rules for configuring authentication for regular, capture, and managed SAPs, also apply to IPoE sessions; see [Figure 98: Configuring IPoE session authentication](#):

- If an authentication policy is applied at capture-SAP or group-interface level, this policy has priority, regardless of whether, or in which other sub-contexts, an LUDB is assigned. Therefore, for an LUDB to provide ESM data, no authentication policy may be applied at capture-SAP or group interface level.
- If an LUDB is applied in the ipoe-session context of a group interface or capture SAP, the LUDBs assigned in the dhcp, dhcp6, and router-solicit related contexts of the same group interface or capture SAP are ignored.

Figure 98: Configuring IPoE session authentication



25644

When the AAA/RADIUS server referenced from the authentication policy is not available, SR OS can rely on a fallback LUDB, if configured; see the LUDB for ESM chapter for more information.

Be aware of the following:

- Static hosts can be configured on a group interface with IPoE sessions enabled. A static host will not be associated with an IPoE session.
- ARP hosts are not supported in an IPoE session and cannot be instantiated on a group interface with IPoE sessions enabled.
- Up to sixteen framed-routes and sixteen framed-IPv6-routes can be associated with an IPoE session.
- Python-based subscriber identification based on the DHCPv4 Ack message is ignored when IPoE sessions are enabled.

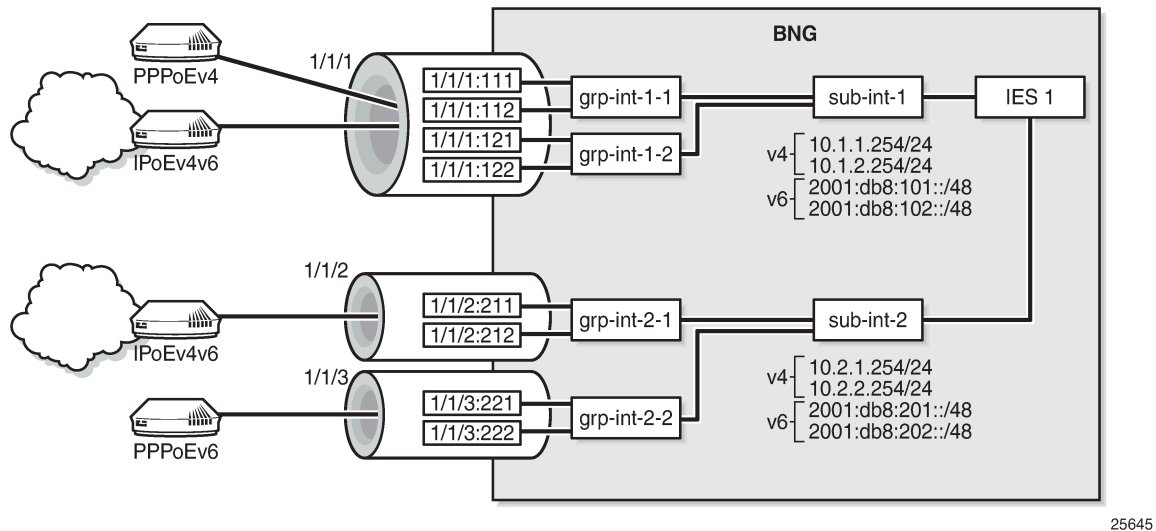
Configuration

Baseline configuration

Figure 99: Baseline configuration shows the baseline configuration used for the examples in this chapter, excluding the ipoe-session configurations. These will be added later in this chapter.

The first example uses LUDB authentication, and the second example uses AAA/RADIUS authentication. As alternatives, AAA/NASREQ authentication or AAA/Gx authentication can be used.

Figure 99: Baseline configuration



The following partial configuration applies to IES-1. This service is provisioned with ESM on all of its SAPs, and supports proxy and relay scenarios on all group interfaces for both IPv4 and IPv6. Only the part relevant to subscriber interface *sub-int-1* and group interface *grp-int-1-1* is shown. The configurations for the other subscriber and group interfaces are similar. See the [ESM Basics](#) and [Routed CO](#) chapters for more information.

```
configure
service
  ies 1 customer 1 create
    subscriber-interface "sub-int-1" create
      address 10.1.1.254/24
      ---snip---
      ipv6
        delegated-prefix-len 56
        link-local-address fe80::ea:4b:ff
        subscriber-prefixes
          prefix 2001:db8:101::/48 wan-host
          prefix 2001:db8:f101::/48 pd
          ---snip---
        exit
      exit
    exit
  group-interface grp-int-1-1
    ipv6
      router-advertisements
        prefix-options
          autonomous
        exit
      no shutdown
    exit
  dhcp6
    proxy-server
      client-applications dhcp ppp
      no shutdown
    exit
  relay
    link-address 2001:db8:101::1
    server 2001:db8::11
```

```

        client-applications dhcp ppp
        no shutdown
    exit
exit
router-solicit
no shutdown
exit
exit
local-address-assignment
ipv6
    client-application ipoe-slaac
    server "dhcp6-srv"
    exit
    no shutdown
exit
arp-populate
dhcp
    proxy-server
        emulated-server 10.1.1.254
        no shutdown
    exit
    server 10.11.11.1
    trusted
    lease-populate 1000
    client-applications dhcp ppp
    gi-address 10.1.1.254
    no shutdown
exit
---snip---
sap 1/1/1:111 create
    sub-sla-mgmt
        def-sub-profile "sub-prof-1"
        def-sla-profile "sla-prof-1"
        sub-ident-policy "sub-id-pol-1"
        multi-sub-sap
        no shutdown
    exit
exit
---snip---
```

No DHCPv4 or DHCPv6 relay options are defined.

Troubleshooting

The syntax to show the active IPoE sessions is as follows:

```

show service id <service-id> ipoe session ?
- session [sap <sap-id>] [mac <ieee-address>] [circuit-id <circuit-id>] [remote-id <remote-id>] [interface <ip-int-name|ip-address>] [inter-dest-id <intermediate-destination-id>] [no-inter-dest-id] [ip-address <ip-prefix[/prefix-length]>] [port <port-id>] [subscriber <sub-ident-string>] [sap-session-id <sap-session-index>] [wholesaler <service-id>]
- session [sap <sap-id>] [mac <ieee-address>] [circuit-id <circuit-id>] [remote-id <remote-id>] [interface <ip-int-name|ip-address>] [inter-dest-id <intermediate-destination-id>] [no-inter-dest-id] [ip-address <ip-prefix[/prefix-length]>] [port <port-id>] [subscriber <sub-ident-string>] [sap-session-id <sap-session-index>] detail [wholesaler <service-id>]
```

The following show commands have been extended, so that session filtering is available:

```

show service id <svc-id> dhcp lease-state ?
```



```
- lease-state [wholesaler <service-id>] [sap <sap-id>|sdp <sdp-id:vc-id>|interface
<interface-name>|ip-address <ip-address[/mask]>|chaddr <ieee-address>|mac <ieee-address>|
{[port <port-id>] [no-inter-dest-id | inter-dest-id <inter-dest-id>]]} [session {none|ipoe}}
[detail]

show service id <svc-id> dhcp6 lease-state ?
- lease-state [detail] [wholesaler <service-id>] [session {none|ipoe|ppp}]
- lease-state [detail] interface <interface-name> [wholesaler <service-id>] [session {none|
ipoe|ppp}]
- lease-state [detail] ipv6-address <ipv6-prefix[/prefix-length]> [wholesaler <service-id>]
[session {none|ipoe|ppp}]
- lease-state [detail] mac <ieee-address> [wholesaler <service-id>] [session {none|ipoe|ppp}]

show service id <svc-id> slaac host ?
- host [detail] [wholesaler <service-id>] [session {none|ipoe|ppp}]
- host interface <interface-name> [detail] [wholesaler <service-id>] [session {none|ipoe|
ppp}]
- host mac <ieee-address> [detail] [wholesaler <service-id>] [session {none|ipoe|ppp}]
- host ipv6-address <ipv6-prefix> [detail] [wholesaler <service-id>] [session {none|ipoe|
ppp}]
- host sap <sap-id> [detail] [wholesaler <service-id>] [session {none|ipoe|ppp}]
```

The following debug configuration is used for demonstration and troubleshooting purposes:

```
debug
router "Base"
  ip
    dhcp
      detail-level medium
      mode egr-ingr-and-dropped
    exit
    dhcp6
      mode egr-ingr-and-dropped
      detail-level medium
    exit
  exit
  radius
    packet-type authentication accounting coa
    detail-level medium
  exit
  exit
  subscriber-mgmt
    local-user-db "ludb-1"
    detail all
  exit
  exit
exit
```

IPoE session failure events are also issued to log-id 99:

```
*A:BNG-1# show log event-control "svcmgr"
=====
Log Events
=====
Application
ID#      Event Name                                P   g/s    Logged    Dropped
-----
 2011 svcTlsMacPinningViolation                WA thr         0         0
---snip---
 2554 tmnxSubIpoeInvalidSessionKey              WA thr        30         0
 2555 tmnxSubIpoeInvalidCidRidChange            WA thr         0         0
 2556 tmnxSubIpoeSessionLimitReached            WA thr         0         0
 2557 tmnxSubIpoePersistenceRecovery            WA thr         0         0
```

```
2559 tmnxSubIpoemigrHostDeleted      WA thr      0      0
=====
*A:BNG-1#
```

IPoE Session Authentication through LUDB

The LUDB *ludb-1* uses the MAC address for host matching, and is defined as follows:

```
configure
  subscriber-mgmt
    local-user-db "ludb-1" create
      description "example user-db"
      ipoe
        match-list mac
        host "entry-01" create
          host-identification
            mac 00:00:00:00:00:01
        exit
        address pool "pool4-1"
        identification-strings 254 create
          subscriber-id "sub-11"
          sla-profile-string "sla-profile-1"
          sub-profile-string "sub-profile-1"
        exit
        ipv6-slaac-prefix-pool "pool6-1"
        ipv6-wan-address-pool "pool6-1"
        no shutdown
      exit
    ---snip---
```

This LUDB is then applied to the group interface in the ipoe-session context:

```
configure
  service
    ies 1 customer 1 create
      subscriber-interface "sub-int-1"
      group-interface "grp-int-1-1"
      ipoe-session
        description "ipoe-sessions with LUDB"
        ipoe-session-policy "sespol-sap-mac"
        sap-session-limit 100
        session-limit 500
        user-db "ludb-1"
        no shutdown
      exit
```

The LUDB applied in the ipoe-session context takes priority over LUDBs applied in the **dhcp router-solicit**, **ipv6 dhcpv6 router-solicit**, and **ipv6 router-solicit** contexts for a Layer 3 service.

The IPoE session policy *sespol-sap-mac* used in this example is defined as follows:

```
configure
  subscriber-mgmt
    ipoe-session-policy "sespol-sap-mac" create
      description "plain ipoe session policy, sap-mac key"
      session-key sap mac
      no session-timeout
    exit
```

Debug

The following debug trace appears when the user with MAC address 00:00:00:00:00:01 first connects using DHCPv4 and subsequently connects using SLAAC and DHCPv6 (IANA), without disconnecting DHCPv4.

Messages 1 through 9 show the message sequence for DHCPv4 (DORA). Messages 11 through 22 show the message sequence for DHCPv6 (SARR). Message 10 and 23 are the router solicitation and advertisement messages. Therefore, three hosts are created.

The LUDB is accessed just once; immediately after the DHCPv4 Discover message:

```
1 2016/02/03 13:51:28.15 CET MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 6 (grp-int-1-1),
  received DHCP Boot Request on Interface grp-int-1-1 (1/1/1:111) Port 67
  H/W Type: Ethernet(10Mb)  H/W Address Length: 6
  ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
  siaddr: 0.0.0.0          giaddr: 0.0.0.0
  chaddr: 00:00:00:00:00:01  xid: 0x1
  DHCP options:
  [53] Message type: Discover
  [255] End
"

2 2016/02/03 13:51:28.15 CET MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
  mac: 00:00:00:00:00:01
  Host entry-01 found in user data base ladb-1"
3 2016/02/03 13:51:28.15 CET MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base),
  transmitted DHCP Boot Request to 10.11.11.1 Port 67
  H/W Type: Ethernet(10Mb)  H/W Address Length: 6
  ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
  siaddr: 0.0.0.0          giaddr: 10.1.1.254
  chaddr: 00:00:00:00:00:01  xid: 0x1
  DHCP options:
  [53] Message type: Discover
  [255] End
"

---snip---
9 2016/02/03 13:51:28.16 CET MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 6 (grp-int-1-1),
  transmitted DHCP Boot Reply to Interface grp-int-1-1 (1/1/1:111) Port 68
  H/W Type: Ethernet(10Mb)  H/W Address Length: 6
  ciaddr: 0.0.0.0          yiaddr: 10.1.1.60
  siaddr: 10.11.11.1       giaddr: 10.1.1.254
  chaddr: 00:00:00:00:00:01  xid: 0x1
  DHCP options:
  [53] Message type: Ack
  [54] DHCP server addr: 10.11.11.1
  [51] Lease time: 900
  [1] Subnet mask: 255.255.255.0
  [3] Router: 10.1.1.254
  [255] End
"

10 2016/02/03 13:51:40.77 CET MINOR: DEBUG #2001 Base TIP
"TIP: ICMP6_PKT
ICMP6 ingressing on grp-int-1-1 (Base):
  fe80::200:ff:fe00:1 -> ff02::2
  Type: Router Solicitation (133)
```

```
Code: No Code (0)
  Option : Src Link Layer Addr 00:00:00:00:00:01
"
11 2016/02/03 13:51:40.78 CET MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
Incoming DHCP6 Msg : SOLICIT (1)
on itf grp-int-1-1
  Trans Id : 0x411fc8
  Option : CLIENTID (1), Length : 14
    LLT : HwTyp=0001,T=507311564,LL=000000000001
    000100011e3cf5cc000000000001
  Option : IA_NA (3), Length : 12
    IAID : 2
    Time1: 0 seconds
    Time2: 0 seconds
  Option : OR0 (6), Length : 2
    Requested Option : DNS_NAME_SRVR (23)
"
---snip---
22 2016/02/03 13:51:40.92 CET MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
Outgoing DHCP6 Msg : REPLY (7)
to itf grp-int-1-1
  Trans Id : 0xe15ddf
  Option : SERVERID (2), Length : 10
    LL : HwTyp=0001,LL=ea4bff000000
    00030001ea4bff000000
  Option : CLIENTID (1), Length : 14
    LLT : HwTyp=0001,T=507311564,LL=000000000001
    000100011e3cf5cc000000000001
  Option : IA_NA (3), Length : 40
    IAID : 2
    Time1: 1800 seconds
    Time2: 2880 seconds
  Option : IAADDR (5), Length : 24
    Address : 2001:db8:101:c::1
    Preferred Lifetime : 3600 seconds
    Valid Lifetime : 86400 seconds
  Option : DNS_NAME_SRVR (23), Length : 16
    Server : 2001:db8::1:1:1:1
"
23 2016/02/03 13:51:41.91 CET MINOR: DEBUG #2001 Base TIP
"TIP: ICMP6_PKT
ICMP6 egressing on grp-int-1-1 (Base):
fe80::ea:4b:ff -> fe80::200:ff:fe00:1
Type: Router Advertisement (134)
Code: No Code (0)
  Hop Limit : 64
  Flags :
  Retrans Time : 0
  Def Life Time : 4500
  Reachable Time: 0
  Option : Src Link Layer Addr 00:00:5e:00:01:01
  Option : Prefix : 2001:db8:101:d::/64
    Flags : On Link Autoconfig
    Valid Life Time: 86400
    Pref Life Time: 3600
"
```

Verification

The following shows the IPoE session for MAC address 00:00:00:00:00:01:

```
*A:BNG-1# show service id 1 ipoe session mac 00:00:00:00:00:01
=====
IPoE sessions for svc-id 1
=====
Sap Id          Mac Address      Up Time      MC-Stdbby
  Subscriber-Id
    [CircuitID] | [RemoteID]
-----
1/1/1:111      00:00:00:00:00:01  0d 00:10:18
  sub-11
-----
CID | RID displayed when included in session-key
Number of sessions : 1
=====
*A:BNG-1#
```

The IPoE session details for MAC address 00:00:00:00:00:01 are shown using the following command. The session time left is undefined (N/A), because no IPoE session-timeout is defined in the IPoE session policy. Re-authentication does not apply, so the minimum authentication interval is infinite (N/A).

```
*A:BNG-1# show service id 1 ipoe session mac 00:00:00:00:00:01 detail
=====
IPoE sessions for service 1
=====
SAP                : 1/1/1:111
Mac Address        : 00:00:00:00:00:01
Circuit-Id         :
Remote-Id          :
Session Key        : sap-mac
MC-Standby         : No
Subscriber-interface : sub-int-1
Group-interface    : grp-int-1-1
Termination Type   : local
Up Time            : 0d 00:09:42
Session Time Left   : N/A
Last Auth Time     : 02/03/2016 13:51:29
Min Auth Intvl (left) : infinite (N/A)
Persistence Key     : N/A
Subscriber         : "sub-11"
Sub-Profile-String : "sub-profile-1"
SLA-Profile-String : "sla-profile-1"
ANCP-String        : ""
Int-Dest-Id        : ""
App-Profile-String : ""
Category-Map-Name  : ""
Acct-Session-Id    : "EA4BFF000001A656B1F7D0"
Sap-Session-Index  : 1
IP Address         : 10.1.1.60/24
IP Origin          : DHCP
Primary DNS        : N/A
Secondary DNS      : N/A
Primary NBNS       : N/A
Secondary NBNS     : N/A
Address-Pool       : pool4-1
IPv6 Prefix        : 2001:db8:101:d::/64
IPv6 Prefix Origin : LclPool
IPv6 Prefix Pool   : "pool6-1"
```

```
IPv6 Del.Pfx.      : N/A
IPv6 Del.Pfx. Origin : None
IPv6 Del.Pfx. Pool  : ""
IPv6 Address       : 2001:db8:101:c::1
IPv6 Address Origin : DHCP
IPv6 Address Pool   : "pool6-1"
Primary IPv6 DNS    : 2001:db8::1:1:1:1
Secondary IPv6 DNS   : N/A
Radius Session-T0    : N/A
Radius Class        :
Radius User-Name     :
```

```
-----
Number of sessions : 1
=====
```

```
*A:BNG-1#
```

The following command shows the subscriber hosts for this MAC address:

```
*A:BNG-1# show service id 1 subscriber-hosts mac 00:00:00:00:00:01
```

```
=====
Subscriber Host table
```

```
=====
Sap      Subscriber
IP Address
MAC Address  PPPoE-SID Origin  Fwding State
-----
1/1/1:111    sub-11
10.1.1.60
00:00:00:00:00:01  N/A      DHCP      Fwding
1/1/1:111    sub-11
2001:db8:101:c::1/128
00:00:00:00:00:01  N/A      IPoE-DHCP6  Fwding
1/1/1:111    sub-11
2001:db8:101:d::/64
00:00:00:00:00:01  N/A      IPoE-SLAAC  Fwding
-----
```

```
Number of subscriber hosts : 3
=====
```

```
*A:BNG-1#
```

The following commands show the corresponding dhcp and dhcp6 lease-states:

```
*A:BNG-1# show service id 1 dhcp lease-state session ipoe
```

```
=====
DHCP lease state table, service 1
```

```
=====
IP Address      Mac Address      Sap/Sdp Id      Remaining Lease   MC
LeaseTime      Origin      Stdby
-----
10.1.1.60       00:00:00:00:00:01 1/1/1:111       00h13m17s  DHCP
-----
```

```
Number of lease states : 1
=====
```

```
*A:BNG-1#
```

```
*A:BNG-1# show service id 1 dhcp6 lease-state session ipoe
```

```
=====
DHCP lease state table, service 1
```

```
=====
IP Address      Mac Address      Sap/Sdp Id      Remaining Lease   MC
LeaseTime      Origin      Stdby
-----
2001:db8:101:c::1/128
00:00:00:00:00:01 1/1/1:111       23h59m02s  DHCP
-----
```

```
-----  
Number of lease states : 1  
=====
```

```
*A:BNG-1#
```

IPoE Session Authentication through AAA/RADIUS

The FreeRADIUS server users file contains following data for the connecting device:

```
00:00:00:00:00:01  Cleartext-Password := "spasswd"  
                   Alc-Subsc-ID-Str = "ipoe-%{User-name}",  
                   Alc-Subsc-Prof-Str = "sub-prof-1",  
                   Alc-SLA-Prof-Str = "sla-prof-1",  
                   Framed-Pool = "pool4-1",  
                   Framed-Ipv6-Pool = "pool6-1",  
                   Alc-Delegated-IPv6-Pool = "pool6-1",  
                   Alc-Relative-Session-Timeout = "300"
```

The authentication policy *radius-pol* used in this example is defined as follows:

```
configure  
  subscriber-mgmt  
    authentication-policy "radius-pol"  
    password letmein  
    include-radius-attribute  
      acct-session-id  
      circuit-id  
      sap-session-index  
    exit  
    radius-server-policy "rsp-1"  
  exit
```

The IPoE session policy used in this example is defined as follows. The key now also includes the circuit ID.

```
configure  
  subscriber-mgmt  
    ipoe-session-policy "sespol-sap-mac-cid" create  
    description "key also including cid now"  
    session-key sap mac cid  
    no session-timeout  
  exit
```

The authentication and IPoE session policies are then applied to the group interface *grp-int-1-1*:

```
configure  
  service  
    ies 1 customer 1 create  
      subscriber-interface "sub-int-1"  
      group-interface "grp-int-1-1"  
        authentication-policy "radius-pol"  
        ipoe-session  
          ipoe-session-policy "sespol-sap-mac-cid"  
          sap-session-limit 100  
          session-limit 100  
          no shutdown  
        exit
```

The authentication policy takes precedence over any LUDB applied in one of the subcontexts of that group interface.

Debug

The following debug trace appears when the user with MAC address 00:00:00:00:00:01 first connects using DHCPv6 (IA_NA and IA_PD) and subsequently connects using DHCPv4.

The ESM data source is accessed just once, immediately after the DHCPv6 Solicit message. Because LDRA is active, the Solicit message is embedded in the Relay Forward message. The RADIUS server is sent an Access-Request message including the circuit ID, and returns an Access-Accept message including the Alc-Relative-Session-Timeout attribute (message 3). Message 14 is the final IPv6 Reply message containing both the IA_NA address and the IA_PD prefix used by the CPE. Messages 15 through 22 are the DHCPv4 DORA messages.

The initial DHCPv6 Solicit message and the initial DHCPv4 Discover message contain the same interface ID/CID (11), which is why no re-authentication is triggered. The IPoE session is deleted when the RADIUS-provided session timer expires, so the BNG releases both the IPv4 and the IPv6 address (messages 23 through 31).

```
1 2016/02/03 14:51:27.54 CET MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
  Incoming DHCP6 Msg : RELAY_FORW (12)
  on itf grp-int-1-1
  Hop Count : 0
  Link Addr : ::
  Peer Addr : fe80::200:ff:fe00:1
  Option : RELAY_MSG (9), Length : 60
  Msg Type : SOLICIT (1)
  Trans Id : 0x89ade3
  Option : CLIENTID (1), Length : 14
  LLT : HwTyp=0001,T=507311564,LL=000000000001
  000100011e3cf5cc000000000001
  Option : IA_PD (25), Length : 12
  IAID : 1
  Time1: 0 seconds
  Time2: 0 seconds
  Option : IA_NA (3), Length : 12
  IAID : 2
  Time1: 0 seconds
  Time2: 0 seconds
  Option : OR0 (6), Length : 2
  Requested Option : DNS_NAME_SRVR (23)
  Option : INTERFACE_ID (18), Length : 2
  Interface Id : 3131 (11)
"
2 2016/02/03 14:51:27.54 CET MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Access-Request(1) 172.16.1.1:1812 id 165 len 109 vrid 1 pol rsp-1
  USER NAME [1] 17 00:00:00:00:00:01
  PASSWORD [2] 16 8DF.2ZKk.XRvmbLExcKEOk
  NAS IP ADDRESS [4] 4 192.0.2.1
  VSA [26] 4 DSL(3561)
  AGENT CIRCUIT ID [1] 2 11
  SESSION ID [44] 22 EA4BFF000001C356B205DF
  VSA [26] 6 Alcatel(6527)
  SAP SESSION INDEX [180] 4 1
"
3 2016/02/03 14:51:27.54 CET MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
```



```
Access-Accept(2) id 165 len 131 from 172.16.1.1:1812 vrid 1 pol rsp-1
VSA [26] 24 Alcatel(6527)
  SUBSC ID STR [11] 22 ipoe-00:00:00:00:00:01
VSA [26] 12 Alcatel(6527)
  SUBSC PROF STR [12] 10 sub-prof-1
VSA [26] 12 Alcatel(6527)
  SLA PROF STR [13] 10 sla-prof-1
FRAMED POOL [88] 7 pool4-1
FRAMED IPV6 POOL [100] 7 pool6-1
VSA [26] 9 Alcatel(6527)
  DELEGATED IPV6 POOL [131] 7 pool6-1
VSA [26] 6 Alcatel(6527)
  RELATIVE SESSION TIMEOUT [160] 4 300
"
4 2016/02/03 14:51:27.55 CET MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
  Outgoing DHCP6 Msg : RELAY_FORW (12)
  to itf int-DHCP
    Hop Count : 1
    Link Addr : 2001:db8:101::1
    Peer Addr : fe80::200:ff:fe00:1
    Option : RELAY_MSG (9), Length : 104
      Msg Type : RELAY_FORW (12)
      Hop Count : 0
      Link Addr : ::
      Peer Addr : fe80::200:ff:fe00:1
      Option : RELAY_MSG (9), Length : 60
        Msg Type : SOLICIT (1)
        Trans Id : 0x89ade3
        ---snip---
      Option : VENDOR_OPTS (17), Length : 36
        Enterprise : 0000197f
        Option : WAN_POOL (1), Length : 7
          pool6-1
        Option : PFX_POOL (2), Length : 7
          pool6-1
        Option : PFX_LEN (3), Length : 1
          56
        Option : RESERVED_NA_LEN (4), Length : 1
          64
"
---snip---
14 2016/02/03 14:51:27.67 CET MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
  Outgoing DHCP6 Msg : RELAY_REPLY (13)
  to itf grp-int-1-1
    Hop Count : 0
    Link Addr : ::
    Peer Addr : fe80::200:ff:fe00:1
    Option : RELAY_MSG (9), Length : 145
      Msg Type : REPLY (7)
      Trans Id : 0x4c546f
      ---snip---
    Option : IA_PD (25), Length : 41
      IAID : 1
      Time1: 1800 seconds
      Time2: 2880 seconds
    Option : IAPREFIX (26), Length : 25
      Prefix : 2001:db8:f101:700::/56
      Preferred Lifetime : 3600 seconds
      Valid Lifetime : 86400 seconds
    Option : IA_NA (3), Length : 40
      IAID : 2
      Time1: 1800 seconds
```

```
Time2: 2880 seconds
Option : IAADDR (5), Length : 24
  Address : 2001:db8:101:c::1
  Preferred Lifetime : 3600 seconds
  Valid Lifetime : 86400 seconds
Option : DNS_NAME_SRVR (23), Length : 16
  Server : 2001:db8::1:1:1:1
Option : INTERFACE_ID (18), Length : 2
  Interface Id : 3131 (11)
"
15 2016/02/03 14:51:34.59 CET MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 6 (grp-int-1-1),
  received DHCP Boot Request on Interface grp-int-1-1 (1/1/1:111) Port 67
  H/W Type: Ethernet(10Mb) H/W Address Length: 6
  ciaddr: 0.0.0.0 yiaddr: 0.0.0.0
  siaddr: 0.0.0.0 giaddr: 0.0.0.0
  chaddr: 00:00:00:00:00:01 xid: 0x1
  DHCP options:
  [82] Relay agent information: len = 4
    [1] Circuit-id: 11
  [53] Message type: Discover
  [255] End
"
---snip---
22 2016/02/03 14:51:34.73 CET MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 6 (grp-int-1-1),
  transmitted DHCP Boot Reply to Interface grp-int-1-1 (1/1/1:111) Port 68
  H/W Type: Ethernet(10Mb) H/W Address Length: 6
  ciaddr: 0.0.0.0 yiaddr: 10.1.1.64
  siaddr: 10.11.11.1 giaddr: 10.1.1.254
  chaddr: 00:00:00:00:00:01 xid: 0x1
  DHCP options:
  [82] Relay agent information: len = 4
    [1] Circuit-id: 11
  [53] Message type: Ack
  [54] DHCP server addr: 10.11.11.1
  [51] Lease time: 900
  [1] Subnet mask: 255.255.255.0
  [3] Router: 10.1.1.254
  [255] End
"
23 2016/02/03 14:56:26.55 CET MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
  Outgoing DHCP6 Msg : RELAY_FORW (12)
  to itf int-DHCP
  Hop Count : 0
  Link Addr : 2001:db8:101::1
  Peer Addr : fe80::200:ff:fe00:1
  Option : RELAY_MSG (9), Length : 80
  Msg Type : RELEASE (8)
  Trans Id : 0x000000
  ---snip---
  Option : INTERFACE_ID (18), Length : 22
  Interface Id : 5f746d6e785f696e7465726e616c5f636c65616e7570 (snipped)
"
---snip---
26 2016/02/03 14:56:26.55 CET MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base),
  transmitted DHCP Boot Request to 10.11.11.1 Port 68
  H/W Type: Ethernet(10Mb) H/W Address Length: 6
  ciaddr: 10.1.1.64 yiaddr: 0.0.0.0
```

```

siaddr: 0.0.0.0      giaddr: 0.0.0.0
chaddr: 00:00:00:00:00:01  xid: 0x0
DHCP options:
[53] Message type: Release
[54] DHCP server addr: 10.11.11.1
[255] End
"
---snip---
31 2016/02/03 14:56:26.56 CET MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
  Incoming DHCP6 Msg : RELAY_REPLY (13)
  on itf int-DHCP
    Hop Count : 0
    Link Addr : 2001:db8:101::1
    Peer Addr : fe80::200:ff:fe00:1
    Option : RELAY_MSG (9), Length : 89
    Msg Type : REPLY (7)
    ---snip---
    Option : IA_PD (25), Length : 49
      IAID : 1
      Time1: 0 seconds
      Time2: 0 seconds
    Option : STATUS_CODE (13), Length : 33
      Status : SUCCESS (0)
      All prefixes have been released
    Option : INTERFACE_ID (18), Length : 22
      Interface Id : 5f746d6e785f696e7465726e616c5f636c65616e7570 (snipped)
"

```

Verification

The following command shows the session details for MAC address 00:00:00:00:00:01. The key now includes the circuit ID (11), and the session timer is running. The RADIUS-provided session timeout is 5 minutes.

```

*A:BNG-1# show service id 1 ipoe session mac 00:00:00:00:00:01 detail
=====
IPoE sessions for service 1
=====
SAP                : 1/1/1:111
Mac Address        : 00:00:00:00:00:01
Circuit-Id         : 11
Remote-Id          :
Session Key        : sap-mac-cid
MC-Standby         : No
Subscriber-interface : sub-int-1
Group-interface    : grp-int-1-1
Termination Type   : local
Up Time            : 0d 00:03:13
Session Time Left  : 0d 00:01:47
Last Auth Time     : 02/03/2016 14:51:28
Min Auth Intvl (left) : infinite (N/A)
Persistence Key    : N/A
Subscriber         : "ipoe-00:00:00:00:00:01"
Sub-Profile-String : "sub-prof-1"
SLA-Profile-String : "sla-prof-1"
ANCP-String        : ""
Int-Dest-Id        : ""
App-Profile-String : ""
Category-Map-Name  : ""
Acct-Session-Id    : "EA4BFF000001C256B205DF"

```

```
Sap-Session-Index      : 1
IP Address              : 10.1.1.64/24
IP Origin               : DHCP
Primary DNS             : N/A
Secondary DNS           : N/A
Primary NBNS            : N/A
Secondary NBNS          : N/A
Address-Pool            : pool4-1
IPv6 Prefix             : N/A
IPv6 Prefix Origin      : None
IPv6 Prefix Pool        : ""
IPv6 Del.Pfx.           : 2001:db8:f101:700::/56
IPv6 Del.Pfx. Origin    : DHCP
IPv6 Del.Pfx. Pool      : "pool6-1"
IPv6 Address            : 2001:db8:101:c::1
IPv6 Address Origin     : DHCP
IPv6 Address Pool       : "pool6-1"
Primary IPv6 DNS        : N/A
Secondary IPv6 DNS      : N/A
Radius Session-T0       : 0d 00:05:00
Radius Class             :
Radius User-Name        : 00:00:00:00:00:01
-----
Number of sessions : 1
=====
*A:BNG-1#
```

Three hosts are created, as the following command shows:

```
*A:BNG-1# show service id 1 subscriber-hosts mac 00:00:00:00:00:01

=====
Subscriber Host table
=====
Sap      Subscriber
  IP Address
  MAC Address      PPPoE-SID Origin      Fwding State
-----
1/1/1:111      ipoe-00:00:00:00:00:01
  10.1.1.64
  00:00:00:00:00:01      N/A      DHCP      Fwding
1/1/1:111      ipoe-00:00:00:00:00:01
  2001:db8:101:c::1/128
  00:00:00:00:00:01      N/A      IPoE-DHCP6      Fwding
1/1/1:111      ipoe-00:00:00:00:00:01
  2001:db8:f101:700::/56
  00:00:00:00:00:01      N/A      IPoE-DHCP6      Fwding
-----
Number of subscriber hosts : 3
=====
*A:BNG-1#
```

After 300 seconds, the session is deleted:

```
*A:BNG-1# show service id 1 ipoe session detail
No entries found.
*A:BNG-1#
```

Conclusion

IPoE sessions offer ISPs a simplified way to manage dual-stack IPoE devices. IPoE sessions have features similar to PPP sessions, in terms of authentication, mid-session changes, and accounting. IPoE sessions can be used on regular, capture, and managed SAPs, and are supported in single- and dual-homed scenarios, including wholesale and retail configurations.

IPv4 DHCP Hosts

This chapter provides information about IPv4 DHCP host configurations.

Topics in this chapter include:

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

Applicability

This chapter is related to the use of IPv4 only, and was originally written for and tested on Release 7.0.R6. The CLI now corresponds to Release 16.0.R6.

Configuration and troubleshooting commands are given for Bridged CO and Routed CO scenarios.

In the Triple Play Service Delivery Architecture (TPSDA), a subscriber is defined as a collection of hosts pertaining to a single access connection (such as a DSL line) and identified by a subscriber identifier. A subscriber host is an end user terminal within the subscriber home (for example, a PC, set-top box, home gateway) that is identified in the network with a unique (IP address; MAC address) tuple for IPoE or (PPPoE session ID; MAC address) tuple for PPPoE.

Following IPv4 host types are distinguished:

- Static hosts
 - ip-mac
 - ip-only
- Dynamic hosts
 - ARP-host
 - DHCP-host
 - PPPoE-host

This chapter provides configuration and troubleshooting commands for DHCP-hosts.

Overview

Knowledge of the Triple Play Service Delivery Architecture (TPSDA) concepts is assumed throughout this document.

The network topology for a Bridged CO environment is displayed in [Figure 100: Bridged CO Network Topology](#) and for a Routed CO environment in [Figure 101: Routed CO Network Topology](#).

Figure 100: Bridged CO Network Topology


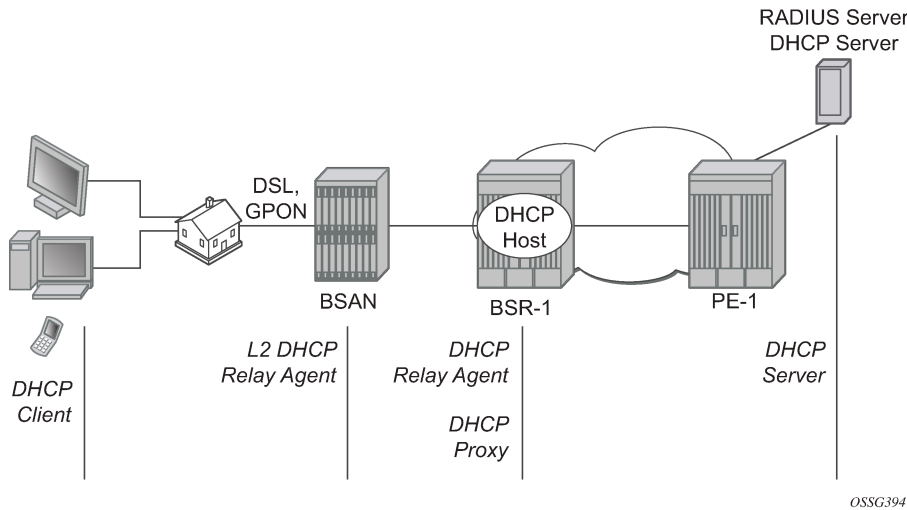


Figure 101: Routed CO Network Topology



Following configuration tasks should be done first and are not detailed in this configuration note:

- Basic service router configurations such as system interface, IGP (ISIS or OSPF), MPLS, BGP.
- Bridged CO service topology: VPLS on BSA-1, terminated in a VPRN or IES service on BSR-1.
- Routed CO service topology: VPRN or IES service with subscriber and group interface on BSR-1.
- External DHCP server: server configuration and connectivity in the VPRN or base router instance.
- External RADIUS server: server configuration and connectivity in the VPRN or base router instance (Enhanced Subscriber Management (ESM) only).

This chapter focuses on DHCP hosts instantiated in a VPLS service on BSA-1 (Bridged CO) or in a VPRN service subscriber interface on BSR-1 (Routed CO). Note that in case of Routed CO, it is also possible to instantiate the DHCP hosts in the base routing instance using an IES service.

Most of the DHCP host functionality is available with Basic Subscriber Management (BSM). When ESM is required, it is explicitly stated.

Review of the DHCP Protocol

The DHCP protocol is used by a DHCP server to dynamically assign IP addresses and other optional configuration parameters on request of DHCP clients. These parameters are leased by the DHCP server for a duration specified by the lease time.

The DHCP lease process is outlined in [Figure 102: DHCP Lease Process](#).

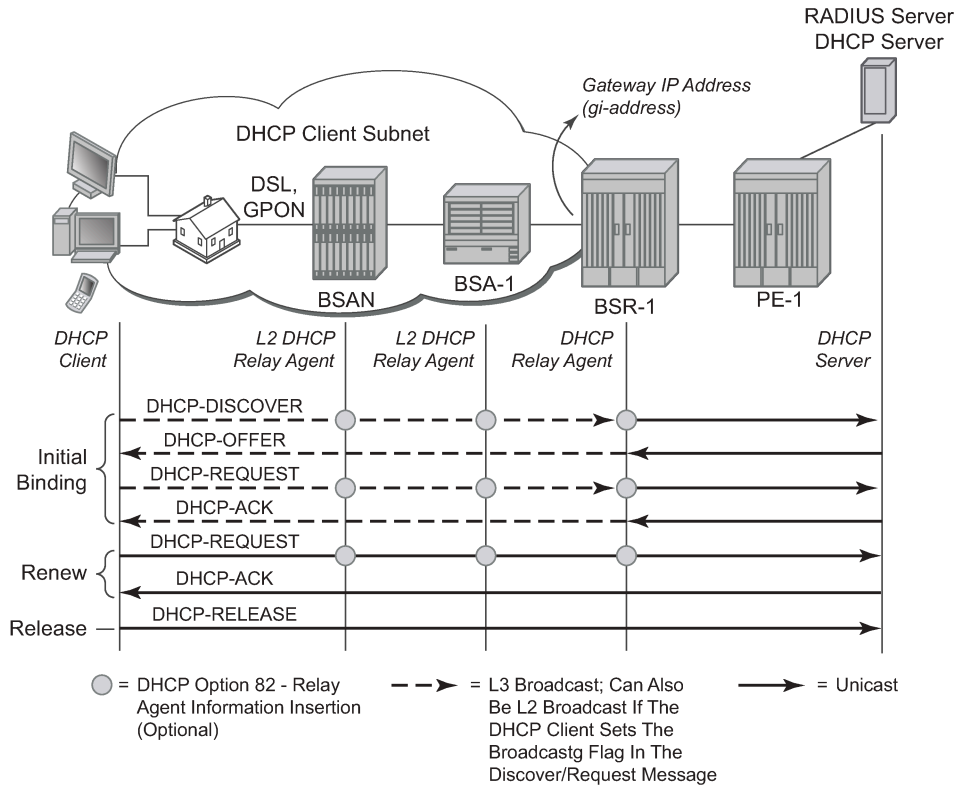
When a DHCP client boots, a DHCP discover message is broadcast on the local subnet (dest-ip = 255.255.255.255).

A DHCP server in the local subnet responds with a unicast DHCP offer message containing the *your ip address* field as well as other configuration parameters in the option fields (such as subnet mask, default gateway, DNS server IP addresses, lease time, etc.).

The DHCP client responds with a DHCP request message to accept the parameters specified in the DHCP offer. The DHCP request is also broadcast on the local subnet.

The DHCP server acknowledges the DHCP request with a unicast DHCP ack message.

Figure 102: DHCP Lease Process



OSSG389

When the DHCP client receives a DHCP ack from the server, it is said to be in the bound state.

When half of the lease time has expired, the DHCP client tries to renew the lease. It will send a unicast DHCP request message to the DHCP server. The DHCP server will reply to the request with a unicast DHCP ack to the client.

If the renew failed, a rebind is attempted by default at 7/8 of the lease time. It will send a broadcast DHCP request message.

Before disconnecting from the local subnet, a DHCP client may return its lease by sending a DHCP release message to the DHCP server.

In case no DHCP server is present in the subnet of the DHCP client, a DHCP relay agent is needed to forward the broadcast DHCP discover/request messages on behalf of the DHCP client to a DHCP server located on a different subnet. The DHCP relay agent will add the gateway IP address field to the messages and send them as unicast to the DHCP server IP address. The DHCP server in this case will respond to the DHCP relay agent using a unicast frame. The DHCP relay agent forwards the DHCP server messages in broadcast frames on the DHCP client subnet.

Configuration

DHCP Snooping

DHCP client originated messages (discover, request, release) must be snooped (intercepted and sent to the control plane for further processing) to enable DHCP Option 82 insertion, authentication through local user database (LUDB), AAA/RADIUS or AAA/Diameter, and releasing the DHCP host session state.

For Bridged CO, DHCP snooping must be enabled explicitly on the subscriber SAP:

```
# Bridged CO @ BSA-1
configure
  service
    vpls 1
      ---snip---
      sap 1/1/3:1 split-horizon-group "rshg-1" create
        description "sub-1"
        dhcp
          snoop
          no shutdown
        exit
      exit
    exit
  exit
exit
```

DHCP server originated messages (offer, ack, nak, etc.) must be snooped to enable DHCP Option 82 removal, lease state population and/or ESM functions.

For Bridged CO, DHCP snooping must be enabled explicitly on all SDPs and/or SAPs that should provide connectivity to the DHCP server:

```
# Bridged CO @ BSA-1
configure
  service
    vpls 1
      ---snip---
      spoke-sdp 1:1 create
        dhcp
          snoop
        exit
      no shutdown
    exit
  exit
exit
```

For Routed CO, DHCP snooping is implicitly enabled by configuring a DHCP relay agent ([DHCP Relay Agent](#)): All DHCP messages received on a routed network interface will be snooped, that is, they are intercepted and sent to the control plane for further processing.

DHCP Relay Agent

For Bridged CO, the DHCP relay agent function is configured in the IP edge (BSR), at the regular interface level:

```
# Bridged CO @ BSR-1
configure
service
  vprn 1
  ---snip---
  interface "int-BSA1-p2mp-1" create
    description "Bridged CO"
    address 10.1.0.254/16
    dhcp
      server 172.16.0.1
      trusted
      gi-address 10.1.0.254
      no shutdown
    exit
  ---snip---
  ip-mtu 1500
  spoke-sdp 1:1 create
    no shutdown
  exit
exit
exit
exit
exit
```

For Routed CO, the DHCP relay agent function must be configured at BSR-1 group-interface level where the DHCP host will be instantiated:

```
# Routed CO @ BSR-1
configure
service
  vprn 1
  ---snip---
  subscriber-interface "sub-int-1" create
    description "Routed CO"
    address 10.2.0.254/16
    group-interface "group-int-1" create
      ---snip---
      dhcp
        server 172.16.0.1
        trusted
        ---snip---
        gi-address 10.2.0.254
        no shutdown
      exit
    exit
  exit
exit
exit
exit
exit
```

The **server** command defines the IP address of the DHCP server and must be reachable in the same routing instance as where the (subscriber-)interface is defined.

The **trusted** command makes the interface a trusted interface and enables Option 82 insertion by a Layer 2 DHCP relay agent (see [DHCP Options \(Relay Agent Information\)](#)).

The **gi-address** must be a locally configured IP address on the (subscriber-) interface. By default the DHCP messages relayed to the DHCP server use the outgoing interface IP address as source IP address. By specifying the optional **src-ip-addr** flag, the configured gi-address is used as the source IP address:

CLI Syntax:

```
- gi-address 10.2.0.254 src-ip-addr
```

A Layer 2 DHCP relay agent (such as BSAN or BSA) can add DHCP Option 82 information and leave the gi-address field to 0.0.0.0. The gi-address is the gateway IP address, filled in by the DHCP relay agent. An incoming DHCP discover with Option 82 present and gi-address field = 0.0.0.0 will be dropped by the DHCP relay agent according the RFC. The Rx Untrusted Packets and client Packets Discarded counters are increased in the DHCP statistics.

Output from DHCP debug log on BSR-1:

```
158 2019/04/09 14:44:32.78 CET MINOR: DEBUG #2001 vprn1 PIP
"PIP: DHCP
instance 2 (1), interface index 4 (group-int-1),
DROPPED DHCP Boot Request on Interface group-int-1 (1/1/3:1) Port 67
Problem: message is received from an untrusted client
```

Therefore, the DHCP relay agent should be configured as trusted to allow DHCP Option 82 insertion by a Layer 2 DHCP relay agent.

DHCP Options (Relay Agent Information)

In Bridged CO, when DHCP snooping is enabled on a VPLS SAP, DHCP Option 82 relay agent information can be altered or added on an incoming DHCP discover/request. This is referred to as a Layer 2 DHCP relay agent function.

In Routed CO, a DHCP relay agent can alter or add the DHCP Option 82 relay agent information on an incoming DHCP discover/request.

Supported DHCP Option 82 sub-options and their format are listed in [Table 24: Supported DHCP Option 82 Sub-Options](#):

Table 24: Supported DHCP Option 82 Sub-Options

Option 82 Sub-Option	Format	Example
Opt82 [1] Circuit ID (Routed CO)	ifindex — 32 bit virtual router ID followed by a 32 bit ifindex in hex	00 00 00 02 00 00 00 04
	sap-id [sap id in ascii]	1/1/3:1
	ascii-tuple [system-name service-id group-interface sap-id]	
	vlan-ascii-tuple [system-name service-id group-interface dot1p vlan-id]	"BSR-1 1 group-int-1 0 1"
Opt82 [1] Circuit ID (Bridged CO)	ascii-tuple [system-name service-id sap-id]	"BSA-1 1 1/1/2:1"

Option 82 Sub-Option	Format	Example
	vlan-ascii-tuple [system-name service-id sap-id dot1p vlan-id]	"BSA-1 1 1/1/2:1 0 1"
Opt82 [2] Remote ID (Bridged and Routed CO)	MAC [client hw address in hex]	fe fd 00 02 45 00
	string (max. 32 chars)	"Opt-82 [2] – Remote ID"
Opt82 [9] Vendor Specific (Bridged and Routed CO)	[1] system-id [hostname in ascii]	"BSA-1" or "BSR-1"
	[2] client-mac-address [client hw address in hex]	fe fd 00 02 45 00
	[3] service-id	1
	[4] sap-id [sap id in ascii]	"1/1/2:1"
	[5] string (max. 32 chars)	"Opt-82 [9] [5] – string"
Opt82 [9] Vendor Specific (Routed CO)	[13] pool-name [dhcp pool name from Radius/Local User DB in ascii]	"dhcp-pool-1"



Note:

The application for the Option 82 Circuit-ID vlan-ascii-tuple format is to preserve the Dot1p marking of DHCP packets in the downstream direction (DHCP server to client). The dot1p value of the incoming DHCP discover/request is recorded as part of the Option 82 Circuit ID. The outgoing DHCP offer/ack packets are then marked with the Dot1p value found as part of the Circuit ID echoed by the DHCP server.

Following actions can be taken on incoming DHCP discover/request:

- replace
- drop
- keep (default)

Replace

At ingress:

If present, remove all the Option 82 information from the incoming DHCP discover/request. Insert the configured DHCP options before forwarding to the DHCP relay agent or DHCP server.

At egress:

Remove all Option 82 information from the incoming DHCP offer/ack before forwarding to the client.

```
# Bridged CO @ BSA-1
configure
service
vpls 1
---snip---
sap 1/1/3:1 split-horizon-group "rshg-1" create
description "sub-1"
dhcp
snoop
option
action replace
```



```
    ---snip---
        sap 1/1/3:1 split-horizon-group "rshg-1" create
        description "sub-1"
        dhcp
            snoop
            option
                action drop
            exit
        exit
    exit
exit
exit
exit
exit
```

```
# Routed C0 @ BSR-1
configure
service
    vprn 1
        ---snip---
        subscriber-interface "sub-int-1" create
        description "Routed C0"
        address 10.2.0.254/16
        group-interface "group-int-1" create
        ---snip---
        dhcp
            option
                action drop
            exit
            server 172.16.0.1
            trusted
            gi-address 10.2.0.254
            no shutdown
        exit
    exit
exit
exit
exit
exit
```

The output from the DHCP debug log on BSA-1 and BSR-1 is as follows:

```
# Bridged C0 @ BSA-1

343 2019/04/10 10:39:28.811 CEST MINOR: DEBUG #2001 Base SVCMMGR
"SVCMMGR: Dropped DHCP Packet
  VPLS 1, SAP 1/1/3:1

  Problem: port config doesn't allow BOOTP/DHCP packets with option 82
```

```
# Routed C0 @ BSR-1

730 2019/04/10 10:42:58.978 CEST MINOR: DEBUG #2001 vprn1 PIP
"PIP: DHCP
instance 2 (1), interface index 6 (group-int-1),
  DROPPED DHCP Boot Request on Interface group-int-1 (1/1/3:1) Port 67
  Problem: action drop is configured and packet contains option 82
```

The Clients Packets Dropped counter is increased in the DHCP statistics:

```
*A:BSA-1# show service id 1 dhcp statistics
```

```
=====
DHCP Statistics, service 1
=====
Client Packets Snooped           : 85
Client Packets Forwarded         : 52
Client Packets Dropped           : 9
Client Packets Proxied (RADIUS)  : 0
Client Packets Proxied (Diameter): 0
Client Packets Proxied (User-Db) : 0
Client Packets Proxied (Lease-Split) : 24
Server Packets Snooped           : 43
Server Packets Forwarded         : 19
Server Packets Dropped           : 24
DHCP RELEASEs Spoofed           : 24
DHCP FORCERENEWs Spoofed        : 0
=====
*A:BSA-1#
```

```
*A:BSR-1# show service id 1 dhcp statistics

=====
DHCP Global Statistics, service 1
=====
Rx Packets           : 287
Tx Packets           : 251
Rx Malformed Packets : 0
Rx Untrusted Packets : 0
Client Packets Discarded : 36
Client Packets Relayed : 124
Client Packets Snooped : 6
Client Packets Proxied (RADIUS) : 0
Client Packets Proxied (Diameter) : 0
Client Packets Proxied (User-Db) : 0
Client Packets Proxied (Lease-Split) : 0
Server Packets Discarded : 52
Server Packets Relayed : 69
Server Packets Snooped : 0
DHCP RELEASEs Spoofed : 52
DHCP FORCERENEWs Spoofed : 0
Client packets streamed : 0
=====
*A:BSR-1#
```

Keep (Default)

At ingress: Incoming DHCP discover/request without Option 82 information will be forwarded to (Bridged CO) or processed by (Routed CO) the DHCP relay agent as is, ignoring any configured option.

At ingress for incoming DHCP discover/request with Option 82 information present. Configured vendor specific options will be merged with the existing Option 82 information before sending to (Routed CO) or processing by (Routed CO) the DHCP relay agent. Configured Circuit ID and Remote ID options will be ignored.

At egress: Remove Option 82 vendor specific information from the incoming DHCP offer/ack before forwarding to the client. Other existing DHCP Option 82 information is retained.

```
# Bridged CO @ BSA-1
configure
  service
    vpls 1
    ---snip---
    sap 1/1/3:1 split-horizon-group "rshg-1" create
```



```

        description "sub-1"
        dhcp
            snoop
            option
                action keep
            exit
        exit
    exit
exit
exit

# Routed C0 @ BSR-1
configure
service
    vprn 1
        ---snip---
        subscriber-interface "sub-int-1" create
            description "Routed C0"
            address 10.2.0.254/16
            group-interface "group-int-1" create
                ---snip---
                dhcp
                    option
                        action keep
                    exit
                    server 172.16.0.1
                    trusted
                    gi-address 10.2.0.254
                    no shutdown
                exit
            exit
        exit
    exit
exit
exit
```

DHCP Lease State

The DHCP lease state table keeps track of the DHCP host states. The DHCP lease state table enables subscriber management functions (per-subscriber QoS and accounting) and security functions (dynamic anti-spoof filtering) on the DHCP host.

The DHCP lease information for a specific host is extracted from the DHCP ack message.

[Table 25: Information in DHCP Lease State](#) displays some information stored in the DHCP lease state. The table does not display all information: additional data is added for managed SAPs, DHCPv6, etc.

Table 25: Information in DHCP Lease State

Parameter	Comment
Service ID	Service where the DHCP host is connected
IP Address	IP address of the DHCP host
Client HW Address	Ethernet MAC address of the DHCP host
Subscriber-interface	Subscriber interface name where the DHCP host is instantiated

Parameter	Comment
(Routed CO only)	
Group-interface (Routed CO only)	Group interface name where the DHCP host is instantiated
SAP	SAP where the DHCP hosts is connected
Up Time	The DHCPv4 host uptime
Remaining Lease Time	The time remaining before the lease expires
Remaining SessionTime	The time remaining before the DHCPv4 host is deleted from the system (updated each time a DHCP renew/rebind occurs)
Persistence Key	Lookup key for this host in the persistency file (see further)
Sub-Ident	ESM: Subscriber ID of the DHCP host
Sub-Profile-String	ESM: Subscriber profile string of the DHCP host
SLA-Profile-String	ESM: SLA profile string of the DHCP host
App-Profile-String	ESM: Application profile string of the DHCP host
Lease ANCP-String	ESM: ANCP string for this DHCP host
Lease Int Dest Id	ESM: Internal destination ID for this DHCP host
Category-Map-Name	ESM: Volume and Time based accounting
Lease Info origin	ESM: Origin for the IP configuration for this host (None, DHCP, RADIUS, etc.)
Ip-Netmask	The IP netmask for this DHCP host
Broadcast-Ip-Addr	The broadcast IP address for this host
Default-Router	The default gateway for this host
Primary-Dns	The primary DNS server for this host
Secondary-Dns	The secondary DNS server for this host
Primary-Nbns	The primary NetBIOS name server for this host
Secondary-Nbns	The secondary NetBIOS name server for this host
ServerLeaseStart	Time and date that the lease for this host started (first DHCP ack received)
ServerLastRenew	Time and date that the lease for this host was last renewed
ServerLeaseEnd	Time and date that the lease for this host will expire
Session-Timeout	The DHCPv4 is deleted when its uptime reaches this value

Parameter	Comment
IPoE PPP session	Indication if this lease belongs to an IPoE or PPP session, or to no session
Lease-Time	The lease time specified by the DHCPv4 server
DHCP Server Addr	IP address of the DHCP server that allocated the lease for this host
Circuit Id	DHCP Relay Agent information Option 82 Circuit ID content
Remote Id	DHCP Relay Agent information Option 82 Remote ID content
RADIUS User-Name	ESM: Username used in the RADIUS authentication access request

For Bridged CO, the DHCP lease state table can only be populated through explicit configuration with the **lease-populate** command. The number of leases allowed on the VPLS SAP must be specified. When omitted, a single DHCP host is allowed per SAP.

```
# Bridged CO @ BSA-1
configure
  service
    vpls 1
      ---snip---
      sap 1/1/3:1 split-horizon-group "rshg-1" create
        description "sub-1"
        dhcp
          snoop
          lease-populate 10
          no shutdown
        exit
      exit
    exit
  exit
exit
```

For Routed CO, DHCP lease state table population is enabled by default on a group interface with DHCP configured as **no shutdown**. The number of leases allowed on each SAP of the group-interface must be configured (by default a single DHCP host is allowed on each SAP):

```
# Routed CO @ BSR-1
configure
  service
    vprn 1
      ---snip---
      subscriber-interface "sub-int-1" create
        description "Routed CO"
        address 10.2.0.254/16
        group-interface "group-int-1" create
          dhcp
            server 172.16.0.1
            trusted
            lease-populate 10
            gi-address 10.2.0.254
            no shutdown
          exit
        exit
      exit
    exit
  exit
exit
```

```
exit
exit
```

To check the DHCP lease state for a particular service, use the **show service id <service-id> dhcp lease-state** command. Detailed output as well as additional output filtering is available:

```
*A:BSA-1# show service id 1 dhcp lease-state ?
- lease-state [wholesaler <service-id>] [sap <sap-id>|sdp <sdp-id:vc-id>|interface
<interface-name>|ip-address <ip-address[/mask]>|chaddr <ieee-address>|mac <ieee-address>|
{[port <port-id>] [no-inter-dest-id | inter-dest-id <inter-dest-id>]]} [session {none|ipoe}]
[detail]
```

```
A:BSR-1# show service id 1 dhcp lease-state ?
- lease-state [wholesaler <service-id>] [sap <sap-id>|sdp <sdp-id:vc-id>|interface
<interface-name>|ip-address <ip-address[/mask]>|chaddr <ieee-address>|mac <ieee-address>|
{[port <port-id>] [no-inter-dest-id | inter-dest-id <inter-dest-id>]]} [session {none|ipoe}]
[detail]
```

```
*A:BSA-1# show service id 1 dhcp lease-state mac 00:00:00:11:11:11 detail
```

```
=====
DHCP lease states for service 1
=====
```

```
Service ID      : 1
IP Address      : 10.1.0.100
Client HW Address : 00:00:00:11:11:11
SAP             : 1/1/3:1
Termination Type : local
Up Time         : 0d 01:46:22
Remaining Lease Time : 0d 10:13:37
Remaining SessionTime: N/A
Persistence Key  : N/A

Sub-Ident       : "sub-11"
Sub-Profile-String : "sub-profile-1"
SLA-Profile-String : "sla-profile-1"
App-Profile-String : ""
Lease ANCP-String : ""
Lease Int Dest Id : ""
Category-Map-Name : ""

Lease Info origin : DHCP

Ip-Netmask      : 255.255.0.0
Broadcast-Ip-Addr : 10.1.255.255
Default-Router   : 10.1.0.254
Primary-Dns      : N/A
Secondary-Dns    : N/A
Primary-Nbns     : N/A
Secondary-Nbns   : N/A

ServerLeaseStart : 04/09/2019 13:47:36
ServerLastRenew  : 04/09/2019 13:47:36
ServerLeaseEnd   : 04/10/2019 01:47:36
Session-Timeout  : N/A
IPoE|PPP session : No
Lease-Time       : 0d 12:00:00
DHCP Server Addr : 172.16.0.1
Radius User-Name : "00:00:00:11:11:11"
```

```
-----
Number of lease states : 1
```

```
=====
*A:BSA-1#

*A:BSR-1# show service id 1 dhcp lease-state mac 00:00:00:33:33:33 detail

=====
DHCP lease states for service 1
=====
Service ID          : 1
IP Address          : 10.2.0.102
Client HW Address   : 00:00:00:33:33:33
Subscriber-interface : sub-int-1
Group-interface     : group-int-1
SAP                 : 1/1/3:1
Termination Type    : local
Up Time             : 0d 01:43:48
Remaining Lease Time : 0d 10:16:12
Remaining SessionTime : N/A
Persistence Key      : N/A

Sub-Ident           : "sub-33"
Sub-Profile-String  : "sub-profile-1"
SLA-Profile-String  : "sla-profile-1"
App-Profile-String  : ""
Lease ANCP-String   : ""
Lease Int Dest Id   : ""
Category-Map-Name   : ""

Lease Info origin    : DHCP

Ip-Netmask           : 255.255.0.0
Broadcast-Ip-Addr    : 10.2.255.255
Default-Router       : 10.2.0.254
Primary-Dns          : N/A
Secondary-Dns        : N/A
Primary-Nbns         : N/A
Secondary-Nbns       : N/A

ServerLeaseStart     : 04/09/2019 13:53:58
ServerLastRenew      : 04/09/2019 13:53:58
ServerLeaseEnd       : 04/10/2019 01:53:58
Session-Timeout      : N/A
IPoE|PPP session     : No
Lease-Time            : 0d 12:00:00
DHCP Server Addr     : 172.16.0.1
Radius User-Name      : "00:00:00:33:33:33"
-----
Number of lease states : 1
=====
*A:BSR-1#
```

DHCP Host Session: Set-up, Operation and Release

Snooping the DHCP communication between a DHCP client and a DHCP relay agent/server facilitates the DHCP host instantiation: Upon the reception of a DHCP ack message from the server, the DHCP lease state table is populated. With ESM enabled, a DHCP host is also instantiated. The DHCP host will appear in the subscriber-host table for the service with origin set to DHCP.

```
*A:BSA-1# show service id 1 subscriber-hosts
```

```
=====
Subscriber Host table
=====
Sap          Subscriber
  IP Address
  MAC Address  PPPoE-SID Origin
-----
1/1/3:1      sub-11
  10.1.0.100
   00:00:00:11:11:11  N/A      DHCP
-----
Number of subscriber hosts : 1
=====
*A:BSA-1#
```

```
*A:BSR-1# show service id 1 subscriber-hosts
=====
Subscriber Host table
=====
Sap          Subscriber
  IP Address
  MAC Address  PPPoE-SID Origin  Fwding State
-----
1/1/3:1      sub-33
  10.2.0.102
   00:00:00:33:33:33  N/A      DHCP      Fwding
-----
Number of subscriber hosts : 1
=====
*A:BSR-1#
```

If ESM is enabled, the subscriber-host will also appear in the active subscriber table:

```
*A:BSR-1# show service active-subscribers
=====
Active Subscribers
=====
Subscriber sub-33 (sub-profile-1)
-----
(1) SLA Profile Instance sap:1/1/3:1 - sla:sla-profile-1
-----
IP Address          MAC Address          Session          Origin          Svc          Fwd
-----
10.2.0.102          00:00:00:33:33:33    N/A              DHCP              1              Y
-----
Number of active subscribers : 1
=====
*A:BSR-1#
```

Troubleshooting the DHCP session set-up is done with DHCP debugging:

```
*A:BSA-1# debug service id 1 dhcp ?
- dhcp
- no dhcp
```

```
[no] detail-level - Configure the DHCP tracing detail level
[no] mac          - Show DHCP packets for a particular MAC address
[no] mode         - Configure the DHCP tracing mode
[no] sap          - Show DHCP packets for a particular SAP
[no] sdp          - Show DHCP packets for a particular SDP
```

```
*A:BSA-1#
```

```
*A:BSR-1# debug router 1 ip dhcp ?
- dhcp [interface <ip-int-name>]
- dhcp mac <ieee-address>
- dhcp sap <sap-id>
- no dhcp [interface <ip-int-name>]
- no dhcp mac <ieee-address>
- no dhcp sap <sap-id>
```

```
---snip---
```

```
*A:BSR-1#
```

For example:

```
*A:BSA-1# show debug
debug
  service
    id 1
      dhcp
        mode egr-ingr-and-dropped
        detail-level medium
      exit
    exit
  exit
exit
*A:BSA-1#
```

```
*A:BSR-1# show debug
debug
  router "1"
    ip
      dhcp
        detail-level medium
        mode egr-ingr-and-dropped
      exit
    exit
  exit
exit
*A:BSR-1#
```

The example above will log all DHCP packets on the service. When thousands of DHCP hosts are active, fine grained filtering is required: for example look only to dropped packets or look only to packets from a particular MAC address.

To display the debugging information, a dedicated log should be created:

```
configure
  log
    log-id 1
      description "Send debug log to a buffer in memory"
      from debug-trace
      to memory
```

```

        no shutdown
    exit
exit
exit

```

The following shows a sample DHCP debug log output (detail-level medium) on BSA-1:

```

77 2019/04/09 13:58:20.022 CEST MINOR: DEBUG #2001 Base SVCNMR
"SVCNMR: RX DHCP Packet
  VPLS 1, spoke-sdp 1:1

  BootReply to UDP port 68
  ciaddr: 0.0.0.0          yiaddr: 10.1.0.102
  siaddr: 172.16.0.1       giaddr: 10.1.0.254
  chaddr: 00:00:00:22:22:22  xid: 0x7

  DHCP options:
  [53] Message type: Ack
  [54] DHCP server addr: 172.16.0.1
  [51] Lease time: 43200
  [1] Subnet mask: 255.255.0.0
  [3] Router: 10.1.0.254
  [58] Renew timeout: 21600
  [59] Rebind timeout: 37800
  [28] Broadcast addr: 10.1.255.255
  [255] End
"

```

During the lifespan of a DHCP host, the DHCP lease state is updated in the system: for example, the remaining lifetime changes after a DHCP renew. To check the lease details from the DHCP host during its lifespan, consult the DHCP lease state details:

```

*A:BSA-1# show service id 1 dhcp lease-state detail

=====
DHCP lease states for service 1
=====
Service ID       : 1
IP Address       : 10.1.0.100
Client HW Address : 00:00:00:11:11:11
SAP              : 1/1/3:1
Termination Type : local
Up Time          : 0d 02:09:01
Remaining Lease Time : 0d 09:50:58
Remaining SessionTime: N/A
Persistence Key   : N/A

---snip---

ServerLeaseStart   : 04/09/2019 13:47:36
ServerLastRenew    : 04/09/2019 13:47:36
ServerLeaseEnd     : 04/10/2019 01:47:36
Session-Timeout    : N/A
IPoE|PPP session   : No
Lease-Time         : 0d 12:00:00
DHCP Server Addr   : 172.16.0.1
Radius User-Name   : "00:00:00:11:11:11"
-----
Number of lease states : 1
=====
*A:BSA-1#

```


If the remaining lifetime timer expires before the DHCP session is renewed or rebound, the DHCP lease state is cleared. If ESM is enabled, the DHCP host is removed from the system.

A DHCP host can be manually deleted from the system using following clear command:

```
*A:BSA-1# clear service id 1 dhcp lease-state ?
- lease-state all [no-dhcp-release]
- lease-state [port <port-id>] inter-dest-id <intermediate-destination-id> [no-dhcp-release]
- lease-state [port <port-id>] no-inter-dest-id [no-dhcp-release]
- lease-state ip-address <ip-address[/mask]> [no-dhcp-release]
- lease-state mac <ieee-address> [no-dhcp-release]
- lease-state port <port-id> [no-dhcp-release]
- lease-state sap <sap-id> [no-dhcp-release]
- lease-state sdp <sdp-id:vc-id> [no-dhcp-release]

---snip---

*A:BSA-1# clear service id 1 dhcp lease-state ip-address 10.1.0.100
```

The DHCP lease state is deleted with all related state (such as, anti-spoof filter, ARP table entry). If ESM is enabled, the DHCP host is removed from the system. Optionally, a DHCP release is sent to the DHCP server to notify that the IP address can be released. This is reflected in the DHCP statistics in the DHCP RELEASES Spoofed counter. Use the **no-dhcp-release** flag in the clear command if no DHCP release is to be sent when issuing the **clear** command.

To display a summary overview of the DHCP configuration on a particular service:

```
*A:BSA-1# show service id 1 dhcp summary

=====
DHCP Summary, service 1
=====
Sap/Sdp           Snoop   Used/   Arp Reply   Info   Admin
                  Provided Agent   Option   State
-----
sap:1/1/3:1       Yes     0/10    Yes         Keep   Up
sap:1/1/3:2       Yes     0/10    Yes         Keep   Up
sdp:1:1           Yes     N/A     N/A         N/A    N/A
-----
Number of Entries : 3
-----
=====
*A:BSA-1#
```

```
*A:BSR-1# show service id 1 dhcp summary

=====
DHCP Summary, service 1
=====
Interface Name      Arp      Leases Per Interface/ Info   Admin
SapId/Sdp           Populate Per Sap Limit   Option State
-----
group-int-1         Yes      1/10             Keep   Up
int-BSA1-p2mp-1     No       0/0              Keep   Up
-----
Interfaces: 2
=====
*A:BSR-1#
```

The Leases Per Interface/Per Sap Limit field indicates the number of active versus the number of allowed DHCP leases on the SAP, SDP or interface.

To check the DHCP statistics, use the following command:

```
*A:BSA-1# show service id 1 dhcp statistics

=====
DHCP Statistics, service 1
=====
Client Packets Snooped           : 33
Client Packets Forwarded         : 33
Client Packets Dropped           : 0
Client Packets Proxied (RADIUS)  : 0
Client Packets Proxied (Diameter): 0
Client Packets Proxied (User-Db) : 0
Client Packets Proxied (Lease-Split) : 0
Server Packets Snooped           : 12
Server Packets Forwarded         : 12
Server Packets Dropped           : 0
DHCP RELEASEs Spoofed           : 1
DHCP FORCERENEWs Spoofed        : 0
=====
*A:BSA-1#
```

```
*A:BSR-1# show service id 1 dhcp statistics

=====
DHCP Global Statistics, service 1
=====
Rx Packets                       : 59
Tx Packets                       : 38
Rx Malformed Packets             : 0
Rx Untrusted Packets             : 0
Client Packets Discarded         : 21
Client Packets Relayed           : 18
Client Packets Snooped           : 2
Client Packets Proxied (RADIUS)  : 0
Client Packets Proxied (Diameter): 0
Client Packets Proxied (User-Db) : 0
Client Packets Proxied (Lease-Split) : 0
Server Packets Discarded         : 0
Server Packets Relayed           : 18
Server Packets Snooped           : 0
DHCP RELEASEs Spoofed           : 0
DHCP FORCERENEWs Spoofed        : 0
Client packets streamed          : 0
=====
*A:BSR-1#
```



Note:

Additional filtering can be done to retrieve DHCP statistics per SAP, SDP or interface.

To clear the DHCP statistics:

```
*A:BSA-1# clear service id 1 dhcp statistics ?
- statistics [sap <sap-id> | sdp <sdp-id:vc-id> |
               interface <ip-int-name|ip-address>]

---snip---
```

```
*A:BSA-1#

*A:BSR-1# clear router 1 dhcp statistics ?
- statistics [<ip-int-name|ip-address>]
<ip-int-name|ip-ad*> : ip-int-name    - 32 chars max
                      ip-address      - a.b.c.d

---snip---

*A:BSR-1#
```

DHCP Hosts Advanced Topics

High Availability

The DHCP lease state supports High Availability (HA): the lease state table is synchronized to the standby CPM. When the active CPM fails, all DHCP hosts stay active without service interruption.

DHCP Lease State Persistency

A DHCP session does not have a keep-alive mechanism to detect unavailability. A new DHCP session set up is only attempted after expiration of the DHCP lease time. A node reboot causing the loss of DHCP lease state and the corresponding anti-spoof filters could therefore result in unacceptable long service outages.

The DHCP lease state can be made persistent across node reboots: DHCP lease state is restored from a persistency file stored on the compact flash file system. As a result, DHCP sessions will only loose connectivity during the time of reboot without being completely disconnected.

To activate the DHCP lease state persistency:

```
configure
  system
    persistence
      subscriber-mgmt
        description "DHCP lease state persistency"
        location cf3:
      exit
    exit
  exit
exit
```

A dedicated persistency file will be created on the specified compact flash file system. The file is initialized to store the maximum number of allowed hosts; its size is constant to avoid file system space problems during operations.

```
*A:BSA-1# file dir cf3:\sub*
Volume in drive cf3 on slot A is SROS VM.
Volume in drive cf3 on slot A is formatted as FAT32
Directory of cf3:
04/10/2019  08:25a                536871424 submgmt.012
04/10/2019  08:25a                12583424 submgmt.i12
                2 File(s)                549454848 bytes.
                0 Dir(s)                 330903552 bytes free.
```

```
*A:BSA-1#
```

Each time a DHCP ack is received from the DHCP server, the persistency file is updated together with the lease state. If the file update fails, an event is generated to indicate that persistency can not be guaranteed.

The content of the persistency file may vary between different SR OS software releases. When upgrading, the persistency file is automatically upgraded to the new format. To downgrade the persistency file to a lower SR OS release version, use the following command:

```
*A:BSA-1# tools perform persistence downgrade ?
- downgrade target-version <target> [reboot]
<target>          : the version you want to downgrade to
                   submgt
                   14.0 (current) - cf3:\submgt.012
                   13.0           - cf3:\submgt.011
                   12.0           - cf3:\submgt.010
                   11.0           - cf3:\submgt.009
                   10.0           - cf3:\submgt.008
                   9.0            - cf3:\submgt.007
                   8.0            - cf3:\submgt.006
                   7.0            - cf3:\submgt.005
                   6.0            - cf3:\submgt.004
                   5.0            - cf3:\submgt.003
                   4.0            - cf3:\submgt.pst
<reboot>          : reboot system after successful conversion
```

The content of the persistency file can be looked at using the following command:

```
*A:BSA-1# show service id 1 dhcp lease-state mac 00:00:00:11:11:11 detail
=====
DHCP lease states for service 1
=====
Service ID       : 1
IP Address       : 10.1.0.100
Client HW Address : 00:00:00:11:11:11
SAP              : 1/1/3:1
Termination Type : local
Up Time          : 0d 00:00:15
Remaining Lease Time : 0d 11:59:45
Remaining SessionTime: N/A
Persistence Key   : 0x00000000

---snip---

Relay Agent Information
  Circuit Id      : 11
  Radius User-Name : "00:00:00:11:11:11"
-----
Number of lease states : 1
=====
*A:BSA-1#
```

```
*A:BSA-1# tools dump persistence submgt record 0x00000000
-----
Persistence Record
-----
Client       : submgt
Persist-Key  : 0x00000000
Filename     : cf3:\submgt.011
Entries      : Index  FedHandle  Last Update          Action Valid
```

```

000064 0x00000000 2019/02/24 08:21:13 (UTC) ADD      Yes
Data      : 300 bytes
Host Type  : DHCP lease state
Service ID : 1
SAP ID     : 1/1/3:1
NH MAC     : 00:00:00:11:11:11
Srvr Lse Start : 2019/04/09 08:21:13 (UTC)
IP         : 10.1.0.100
CHADDR     : 00:00:00:11:11:11
Srvr Last Renew: 2019/04/09 08:21:13 (UTC)
Srvr Lse End   : 2019/04/09 20:21:13 (UTC)
Srvr Addr     : 172.16.0.1
Option82      : 4 bytes
RADIUS Fallback: NO
Acct-Sess-Id   : 02D9FF0000000856CD67F9
Multi-Sess-Id  : 02D9FF0000000956CD67F9
Class Attr     : 0 bytes
User-Name      : "00:00:00:11:11:11"
Address Origin : DHCP
host is authenticated by radius: true
Subscriber-Id  : "sub-11"
Sub-Profile-Str: "sub-prof-1"
SLA-Profile-Str: "sla-prof-1"
Framed IP Netmask: 255.255.0.0
Broadcast IP Address: 10.1.255.255
Default Router : 10.1.0.254
Lease-Time     : 43200
*A:BSA-1#

```

Limiting the Number of DHCP Hosts

Lease populate limit

The maximum number of DHCP lease state entries for a VPLS SAP, for an IES/VP RN interface or for each SAP on an IES/VP RN group-interface is defined when enabling the lease-populate. When omitted, a single DHCP host is allowed:

```

configure
  service
    vpls 1
      ---snip---
      sap 1/1/3:2 split-horizon-group "rshg-1" create
        description "sub-1"
        dhcp
          snoop
          lease-populate 1
          no shutdown
        exit
      exit
    exit
  exit
exit

```

When trying to instantiate a new DHCP host while the configured number of leases is reached, the DHCP ack is dropped (DHCP debug log output):

```

110 2019/04/09 16:25:03.030 CEST MINOR: DEBUG #2001 Base SVC MGR

```

```
"SVCMMGR: Dropped DHCP Packet  
VPLS 1, spoke-sdp 1:1
```

```
Problem: lease-populate limit (1) exceeded on SAP 1/1/3:2
```

The following event is generated (log-id 99):

```
95 2019/04/09 16:26:33.871 CEST WARNING: DHCP #2002 Base Maximum number of lease states  
"Lease state for (CiAddr = 10.1.0.102, ChAddr = 00:00:00:22:22:22, leaseTime = 43200)  
was not stored because the number of DHCP lease states on SAP 1/1/3:2 in service 1 has  
reached its upper limit"
```

With ESM enabled, the following additional limits apply:

- sla-profile host-limits
- multi-sub-sap limit

SLA-profile host limits

The SLA-profile contains host limits defining the maximum number of dynamic subscriber hosts per subscriber for this sla-profile. Static hosts are not counted in the host-limits.

```
*A:BSA-1>config>subscr-mgmt>sla-prof# host-limits ?  
- host-limits  
- no host-limits  
  
[no] ipv4-arp      - Maximum number of IPv4 ARP hosts  
[no] ipv4-dhcp     - Maximum number of IPv4 DHCP hosts  
[no] ipv4-overall  - Maximum number of IPv4 hosts  
[no] ipv4-ppp      - Maximum number of IPv4 PPP hosts  
[no] ipv6-overall  - Maximum number of IPv6 hosts  
[no] ipv6-pd-ipoe-d* - Maximum number of IPv6-PD IPOE DHCP hosts  
[no] ipv6-pd-overall - Maximum number of IPv6-PD hosts  
[no] ipv6-pd-ppp-dh* - Maximum number of IPv6-PD PPP DHCP hosts  
[no] ipv6-wan-ipoe-* - Maximum number of IPv6-Wan IPOE DHCP hosts  
[no] ipv6-wan-ipoe-* - Maximum number of IPv6-Wan IPOE SLAAC hosts  
[no] ipv6-wan-overa* - Maximum number of IPv6-Wan hosts  
[no] ipv6-wan-ppp-d* - Maximum number of IPv6-Wan PPP DHCP hosts  
[no] ipv6-wan-ppp-s* - Maximum number of IPv6-Wan PPP SLAAC hosts  
[no] lac-overall   - Maximum number of L2TP LAC hosts  
[no] overall       - Maximum number of hosts  
[no] remove-oldest - Remove oldest
```

Optionally the remove-oldest command can be used. In that case, the new host is accepted and the DHCP lease state for the oldest host (with the least remaining lease time) is cleared. A DHCP release message is sent to the DHCP server.

The following example limits the amount of ipv4-dhcp hosts.

```
configure  
  subscriber-mgmt  
    sla-profile "sla-profile-2" create  
      host-limits  
        ipv4-dhcp 1  
      exit  
    exit  
  exit  
exit
```

If the configured host-limit is reached for a subscriber, access is denied for a new host, an event is generated (log-id 99) and the corresponding DHCP ack message is dropped:

```
113 2019/04/09 16:34:05.002 CEST WARNING: DHCP #2005 Base Lease State Population Error
"Lease state table population error on SAP 1/1/3:2 in service 1 - subscriber sub-21,
sla-profile sla-profile-2 : host-limit ipv4-dhcp (1) exceeded "
```

Multi-sub-sap limit

The multi-sub-sap parameter defines the maximum number of subscribers (dynamic and static) that can be simultaneously active on this SAP. By default only a single subscriber is allowed (no multi-sub-sap).

```
# Bridged C0 @ BSA-1
configure
service
  vpls 1
    sap 1/1/3:2
      sub-sla-mgmt
        multi-sub-sap 2
      exit
    exit
  exit
exit
```

```
# Routed C0 @ BSR-1
configure
service
  vprn 1
    subscriber-interface "sub-int-1"
    group-interface "group-int-1"
      sap 1/1/3:2
        sub-sla-mgmt
          multi-sub-sap 2
        exit
      exit
    exit
  exit
exit
exit
exit
```

If the limit is reached, a new subscriber will be denied access, an event is generated (log-id 99) and the corresponding DHCP ack message is dropped:

```
112 2019/04/09 16:38:22.027 CET WARNING: DHCP #2005 vprn1 Lease State Population Error
"Lease state table population error on SAP 1/1/3:2 in service 1 - Number of
subscribers exceeds the configured multi-sub-sap limit (2)"
```

DHCP Host Connectivity Verification

Because the DHCP protocol does not have a keep-alive mechanism and IP address renewal is not frequent enough, alternative mechanisms are needed to track reachability of DHCP hosts.

The first alternative is called Subscriber Host Connectivity Verification (SHCV). An ARP unicast message is periodically sent to the DHCP host. The connectivity test fails:

- If for X consecutive unicast ARP requests no ARP reply is received within the specified retry-timeout ([10 — 60] seconds, default 10). The number of retries (X-1) is specified by the retry-count ([2 — 29], default 2). Hence, at minimum 3 unicast ARP requests are sent before connectivity is lost.
- If the ARP reply contains an inconsistent IP/MAC compared with the local DHCP lease state

For a failed connectivity test, an event is raised and optionally the DHCP lease state is removed from the system by cleaning up all related resources (e.g. anti-spoof table) and sending a DHCP release to the DHCP server. When ESM is enabled, the DHCP host also is removed.

The interval for the periodic checks can be configured between 1 and 6000 minutes. If not specified, the default value of 10 minutes will be used.

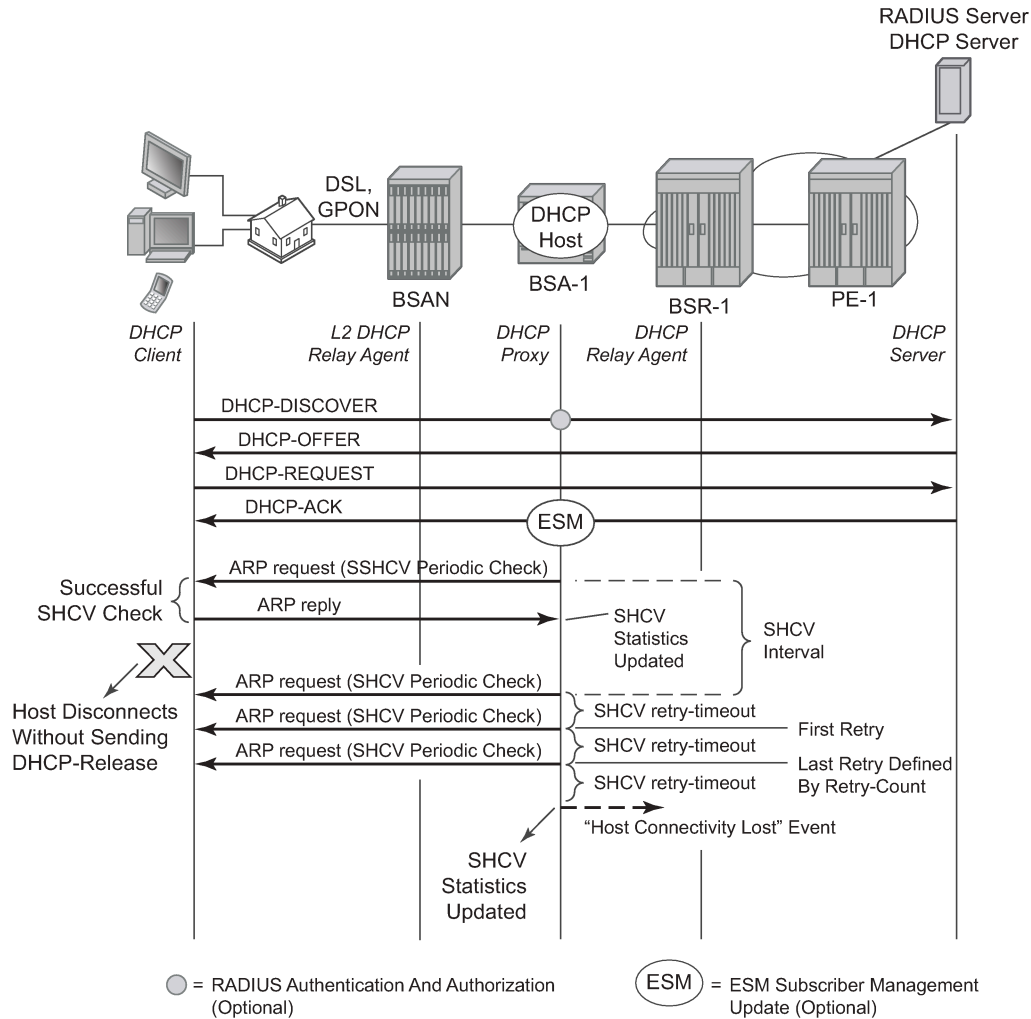
The maximum time for DHCP host connectivity loss detection in this case is:

$$((\text{host-connectivity-verify interval}) + ((\text{retry-count}) * (\text{retry-timeout})))$$

The parameters for the **host-connectivity-verify** command are:

```
*A:BSA-1>config>service>vpls>sap# host-connectivity-verify ?
- host-connectivity-verify source-ip <ip-address> [source-mac <ieee-address>] [interval
<interval>] [action {remove|alarm}] [timeout <retry-timeout>]
[retry-count <count>]
<ip-address>          : a.b.c.d
<ieee-address>        : xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx
<interval>            : [1..6000] minutes
<{remove|alarm}>      : keywords
<retry-timeout>       : [10..60] seconds
<count>               : [2..29]
```


Figure 103: Subscriber Host Connectivity Verification



The configuration for Bridged CO is as follows, where the configured source IP should be an unused unique IP address in the DHCP client subnet or alternatively use source-ip 0.0.0.0. As the host-connectivity-verify application is sending a unicast ARP to the DHCP host, its ARP table is updated with the configured source-ip and source-mac (chassis MAC if not configured). If an existing IP address is used, the DHCP host ARP table gets poisoned, breaking the connectivity to that host.

```
# Bridged CO @ BSA-1
configure
service
vpls 1
sap 1/1/3:2
host-connectivity-verify source-ip 0.0.0.0 interval 1 action alarm
exit
```

The configuration for Routed CO is as follows, where the source IP is not configurable. The source-ip used in the unicast ARP is set to the local subscriber interface address in the subnet of the DHCP hosts that is checked for connectivity.

```
# Routed CO @ BSR-1
configure
service
  vprn 1
    subscriber-interface "sub-int-1"
    group-interface "group-int-1"
    host-connectivity-verify interval 1 action remove
  exit
exit
```

To verify the result of the connectivity check:

```
*A:BSA-1# show service id 1 host-connectivity-verify statistics

=====
Host connectivity check statistics
=====
SvcId  SapId/SdpId      HostIp      MAC
DestIp  Oper      Last-reply/Conn-lost
-----
1      1/1/3:2      10.1.0.101  00:00:00:22:22:21
10.1.0.101
Up      04/09/2019 16:53:20 (elapsed: 0d 00:01:05)
-----
1 host-connectivity states : 1 Up / 0 Down / 0 Retry pending
=====
*A:BSA-1#
```

With action alarm, the lease-state is not removed in case the connectivity with the host is lost. An event is generated (log-id 99) and the statistics show:

```
*A:BSA-1# show service id 1 host-connectivity-verify statistics

=====
Host connectivity check statistics
=====
SvcId  SapId/SdpId      HostIp      MAC
DestIp  Oper      Last-reply/Conn-lost
-----
1      1/1/3:2      10.1.0.101  00:00:00:22:22:21
10.1.0.101
Down    04/10/2019 09:19:18 (elapsed: 0d 00:01:47)
-----
1 host-connectivity states : 0 Up / 1 Down / 0 Retry pending
=====
*A:BSA-1#
```

In case the connectivity with the host is lost, following event is generated:

```
229 2019/04/10 09:20:57.933 CEST WARNING: SVCMGR #2206 Base Host connectivity lost
"host connectivity lost on 1/1/3:2 in service 1 for inetAddr = 10.1.0.101,
chAddr=00:00:00:22:22:21, verify-addr=10.1.0.101."
```

When connectivity is restored, following event (log-id 99) is generated:

```
231 2019/04/10 09:27:22.933 CEST WARNING: SVCNMR #2207 Base Host connectivity restored
"host connectivity restored on 1/1/3:2 in service 1, for inetAddr = 10.1.0.101,
chAddr=00:00:00:22:22:21, verify-addr=10.1.0.101."
```

Connectivity to a DHCP host can also be checked using an OAM command:

```
*A:BSA-1# oam host-connectivity-verify service 1 sap 1/1/3:2
=====
Triggering host connectivity verify for service 1 sap 1/1/3:2 ...
Waiting 3 seconds ...

=====
Host connectivity check statistics
=====
SvcId  SapId/SdpId      HostIp      MAC
DestIp
Oper    Last-reply/Conn-lost
-----
1       1/1/3:2          10.1.0.101  00:00:00:22:22:21
10.1.0.101
Up      04/10/2019 09:33:09 (elapsed: 0d 00:00:02)
-----
1 host-connectivity states : 1 Up / 0 Down / 0 Retry pending
=====
*A:BSA-1#
```

Note that in this case, no action is triggered. If the connectivity test is successful, the host-connectivity-verify statistics are updated with the new timestamp last-reply. If the connectivity test fails, the host-connectivity state becomes Retry Pending (oper state unknown) until an automatic test is scheduled again in the next interval.

To troubleshoot host-connectivity-verify, enable following debug log (additional filtering is possible on ip address, mac address and/or SAP):

```
debug
  service
    id 1
      host-connectivity-verify
      exit
    exit
  exit
exit
```

DHCP Lease Split

The second alternative to the DHCP protocol not having a keep-alive mechanism to verify connectivity is to use a DHCP proxy server with the lease-split option.

A finer granularity of DHCP lease time is used between the DHCP client and the DHCP proxy server than between the DHCP proxy server and the DHCP server.

The maximum time for DHCP host connectivity loss detection in this case is the configured DHCP lease-split lease time.

DHCP communication between the DHCP client and DHCP server is snooped. In the DHCP ack message, the offered lease-time from the DHCP server is replaced with the configured DHCP proxy server lease-split

lease time. Note that the lease time is only updated if the configured lease-split lease time is less than half of the original lease time value. The minimum value for the proxy server lease-split lease time is 5 minutes. When the DHCP client renews the DHCP session, the DHCP proxy server sends a DHCP ack on behalf of the DHCP server as long as the next renew time is earlier than half of the DHCP server expiry time for this session. With ESM enabled, RADIUS re-authentication will occur only when the DHCP request must be sent to the DHCP server. In other words, configuring a DHCP proxy with lease-split does not put extra load on the RADIUS server.

In the example in [Figure 104: DHCP Proxy Server: Lease Split Operation](#), the DHCP server offers a lease time of 960 seconds. The lease time in the offer sent to DHCP client will be updated with the lease time of 300 seconds as configured in the DHCP proxy server lease-split on BSA-1.

```
# Bridged C0 @ BSA-1
configure
  service
    vpls 1
      sap 1/1/3:2
        dhcp
          proxy-server
            lease-time min 5
            no shutdown
          exit
        exit
      exit
    exit
```

```
# Routed C0 @ BSR-1
configure
  service
    vprn 1
      subscriber-interface "sub-int-1"
        group-interface "group-int-1"
          dhcp
            proxy-server
              lease-time min 5
              no shutdown
            exit
          exit
        exit
      exit
    exit
```



Note:

The emulated server address in the DHCP proxy-server configuration does not have to be configured for lease-split operation. This parameter is needed for an alternative use of the DHCP proxy server: RADIUS based IP configuration of a subscriber host. This is out of the scope of this configuration note.

If DHCP lease split is operational for a DHCP host, it will be shown in the Remaining Lifetime field of the detailed lease-state output. Note that the Session Timeout field is the original offered lease time from the DHCP server.

```
*A:BSA-1# show service id 1 dhcp lease-state detail
```

```
=====
DHCP lease states for service 1
=====
Service ID       : 1
IP Address       : 10.1.0.101
Client HW Address : 00:00:00:22:22:21
SAP              : 1/1/3:2
Termination Type : local
Up Time         : 0d 00:00:10
```

```
Remaining Lease Time : 0d 00:04:49 (Lease Split)
Remaining SessionTime: N/A
Persistence Key      : N/A
```

```
Sub-Ident           : "sub-21"
Sub-Profile-String  : "sub-profile-2"
SLA-Profile-String  : "sla-profile-2"
App-Profile-String  : ""
Lease ANCP-String   : ""
Lease Int Dest Id   : ""
Category-Map-Name   : ""
```

```
Lease Info origin   : DHCP
```

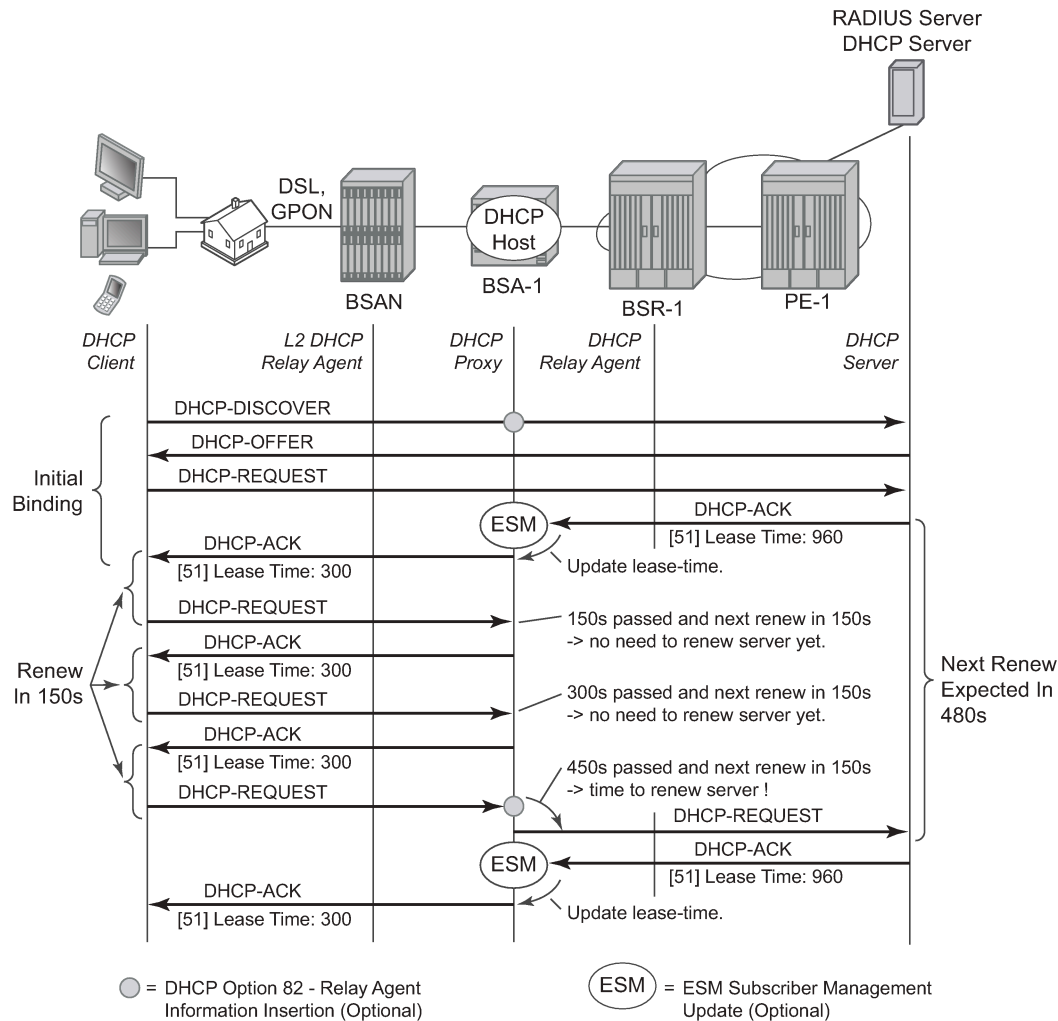
```
Ip-Netmask          : 255.255.0.0
Broadcast-Ip-Addr    : 10.1.255.255
Default-Router       : 10.1.0.254
Primary-Dns          : N/A
Secondary-Dns         : N/A
Primary-Nbns         : N/A
Secondary-Nbns       : N/A
```

```
ServerLeaseStart    : 04/10/2019 09:44:06
ServerLastRenew     : 04/10/2019 09:44:06
ServerLeaseEnd       : 04/10/2019 21:44:06
Session-Timeout      : N/A
IPoE|PPP session    : No
Lease-Time           : 0d 12:00:00
DHCP Server Addr     : 172.16.0.1
Radius User-Name     : "00:00:00:22:22:21"
```

```
-----
Number of lease states : 1
```

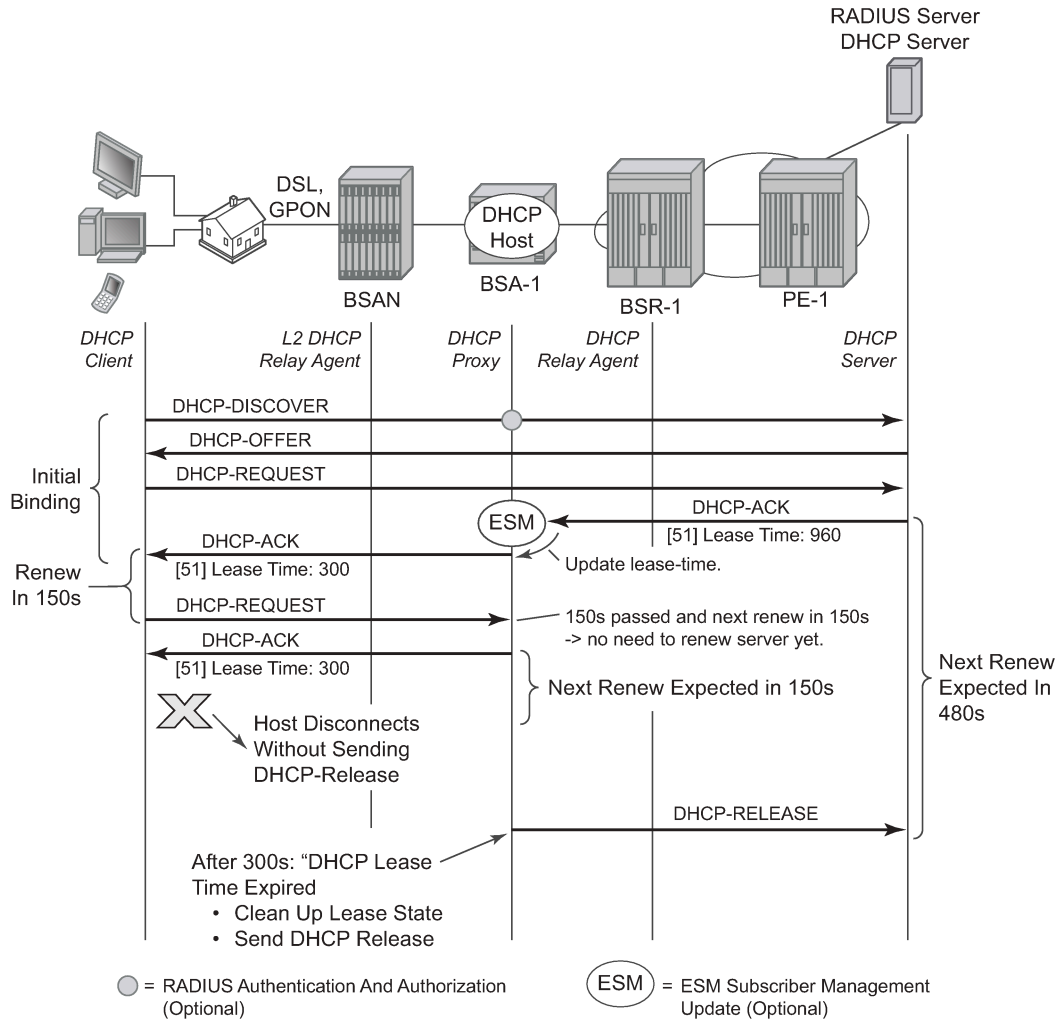
```
=====
*A:BSA-1#
```

Figure 104: DHCP Proxy Server: Lease Split Operation



When the DHCP client disconnects without sending a DHCP release, the DHCP lease state in the BSA/BSR will be removed only when the DHCP lease time expires. With DHCP proxy server lease-split, the DHCP client disconnection can be sped up. In the example below, the DHCP client disconnection is detected in less than 5 minutes (lease-split lease time) while it would have taken up to 16 minutes without the lease-split. Note that the values are illustrative; in reality the DHCP lease times will be higher.

Figure 105: DHCP Proxy Server: Lease Split Operation, DHCP Client Disconnected

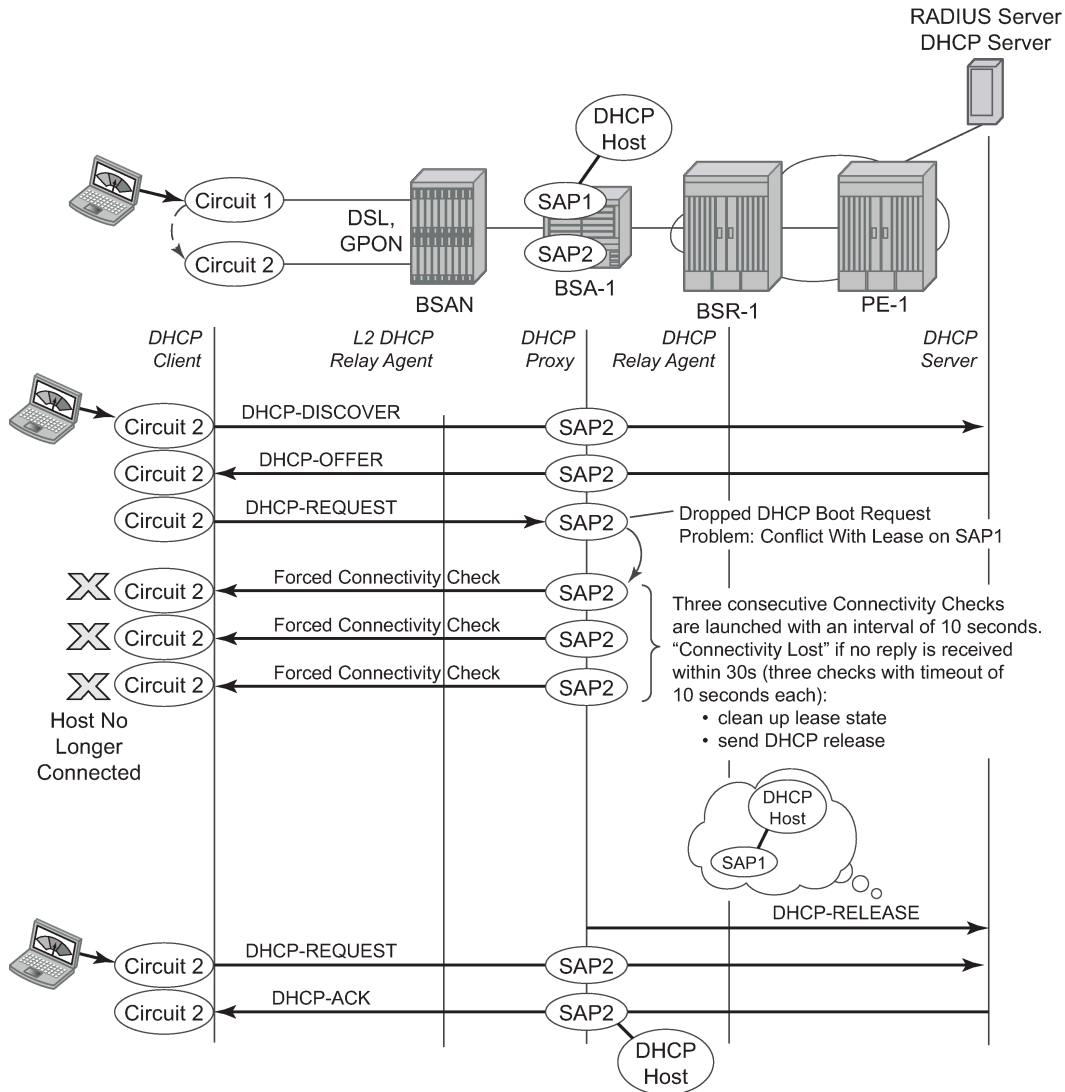


DHCP Host Mobility

A field technician verifying DSLAM operation often connects and disconnects from different ports rapidly. This will require the node to clear its own DHCP host state, the DHCP server state as well as flush MAC addresses learned within the VPLS network or clear ARP entries from the routing instance.

A DHCP request enters on SAP2. On SAP1 there exists a lease state with the same Client Hardware address. The packet is dropped and a forced SHCV check verifies the existing lease state on SAP1. Three consecutive checks are launched with a timeout of 10 seconds. If the host indeed moved from SAP1 to SAP2, the connectivity check will fail on SAP1. The existing lease state is deleted and a DHCP release message is sent to the DHCP server. Any subsequent DHCP session setup will proceed as normal.

Figure 106: DHCP Host Mobility



Note that for host mobility to function, host-connectivity-verification must be enabled. Next to periodic connectivity checks, it also enables forced checks triggered by moving hosts.

For Bridged CO, host-connectivity-verify must be enabled on the SAPs. When no interval is specified, it will default to 10 minutes for the periodic connectivity checks.

```
# Bridged CO
configure
service
vpls 1
sap 1/1/3:1
host-connectivity-verify source-ip 10.1.0.253
exit
sap 1/1/3:2
host-connectivity-verify source-ip 10.1.0.253
exit
exit
```



```
    exit
exit
```

The configured source-ip should be an unused unique ip address in the DHCP client subnet or alternatively use source-ip 0.0.0.0. As the host-connectivity-verify application is sending a unicast ARP to the DHCP host, its ARP table is updated with the configured source-ip and source-mac (chassis MAC if not configured). If you would use an existing IP address, the DHCP host ARP table gets poisoned, breaking the connectivity to that host.

For Routed CO, host-connectivity-verify must be enabled on the group-interface. When no interval is specified, it will default to 10 minutes for the periodic connectivity checks.

```
# Routed CO
configure
  service
    vprn 1 customer 1 create
      subscriber-interface "sub-int-1" create
        group-interface "group-int-1" create
          host-connectivity-verify
        exit
      exit
    exit
  exit
exit
```

The source IP address is not configurable. The source-ip used in the unicast ARP is fixed to the local subscriber interface address in the subnet of the DHCP hosts that is checked for connectivity.

Conclusion

This chapter provides configuration and troubleshooting commands for dynamic DHCP hosts. DHCP hosts can be instantiated in a Layer 2 bridged CO (VPLS) environment as well as in a Layer 3 Routed CO (IES/ VPRN subscriber interface) context.

L2TP for Subscriber Access — LAC

This chapter provides information about L2TP for subscriber access.

Topics in this chapter include:

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

Applicability

This chapter describes L2TP Access Concentrator (LAC) support for the L2TP Access Aggregation (LAA) architecture model and was initially written for SR OS Release 11.0.R4. The CLI in the current edition is based on Release 16.0.R7. PPP hosts are supported in a Routed CO model (with IES or VPRN services) using ATM, Ethernet or Ethernet over Pseudowire SAPs. A description of the L2TP Tunnel Switch (LTS) and L2TP Network Server (LNS) functions are out of the scope of this chapter.

Overview

PPP access architectures

The Broadband Forum proposes two architectures for point-to-point protocol (PPP) access.

- The PPP terminated aggregation architecture (PTA)
- The L2TP access aggregation architecture (LAA)


The PTA architecture (local-access model) uses the Broadband Network Gateway (BNG) to terminate user PPP sessions (see scenario PPPoE-5 in [Figure 107: PPP access architectures](#)).

The LAA architecture (which is a tunneled access model) uses a LAC and an LNS to transport PPP sessions from the LAC to the LNS which performs tunnel termination (see scenario PPPoE-1 and PPPoE-2 in [Figure 107: PPP access architectures](#)).

Optionally, an LTS can be used in the transport network to perform the grooming of traffic between tunnels (see scenarios PPPoE-3 and PPPoE-4 in [Figure 107: PPP access architectures](#)).

The LNS is the logical termination point of the PPP sessions originated by the remote clients and tunneled by the LAC/LTS.

Figure 107: PPP access architectures




al_0521

L2TP tunnels - LAC and LNS reachability options

The router instance where the L2TP tunnel starts and where ESM is handled can be one and the same, but does not need to be the same. The LNS peer address can be reachable via IP, BGP/IGP shortcuts, over a spoke SDP (GRE/MPLS), RFC 4364 VPRNs, but cannot be an address belonging to a directly connected interface. See [Figure 108: Supported L2TP reachability options](#).

Figure 108: Supported LT2P reachability options



al_0519

Recapitulation of the L2TPv2 protocol

L2TPv2 is a client-server protocol relying on UDP and encapsulates Layer 2 packets such as PPP for transmission across a network. L2TPv2 passes control and data messages over separate control and data channels, thus defines following message types:

- Control messages—The in-band control channel passes sequenced control messages, supporting connection management, call management, error reporting, and session control. Optionally, a shared-secret challenge authentication method can be used between the tunnel endpoints.

The following messages are used for L2TP tunnel management:

- Tunnel setup (control connection management)
 - Start-control-connection-request (SCCRQ)
 - Start-control-connection-reply (SCCRP)
 - Start-control-connection-connected (SCCCN)

- Stop-control-connection-notification (StopCCN)
- Tunnel keepalive
 - Hello

The following messages are used for L2TP session (call) management:

- Session setup over an existing tunnel
 - Incoming-call-request (ICRQ)
 - Incoming-call-reply (ICRP)
 - Incoming-call-connected (ICCN)
 - Call-disconnect-notify (CDN)

Zero-length body (ZLB) messages are control packets with an L2TP header only and are used to explicitly acknowledge packets, making the control channel reliable.

L2TP message encoding is done through attribute value pairs (AVP).

- Data messages — Data messages encapsulate the PPP frames that are sent into the L2TP tunnel.

L2TPv2 sessions run over an L2TP tunnel and are referenced by an L2TP session-id. An L2TP tunnel can carry none, one, or multiple L2TP sessions. An L2TP session corresponds to a PPPoE session. L2TPv3 for LAC-LNS dynamic tunnel setup is not supported.

L2TP header and AVP layout

The L2TPv2 header consists of following fields (RFC 2611):

0	8	16	31
T	L	-	S
-	O	P	-
Version		Length	
Tunnel-ID		Session-ID	
Ns		Nr	
Offset Size		Offset Pad	

al_0513A

Table 26: L2TPv2 header fields and descriptions

Field	Description
T	Type of L2TP message (1 bit): 0—data message 1—control message
L	Indicates if the optional Length field is present in the message (1 bit): 0—the field is left out of the message entirely 1—the field is included (must be included in control messages)
-	Reserved for future use, must be set to zero.
S	Indicates if the Ns and Nr fields are present (1 bit): 0 — the fields are left out of the message; entirely 1 — the fields are included (must be included in control messages)
O	Indicates if the Offset field is present (1 bit): 0 — the field is left out of the message entirely (must be left out of control messages); 1 — the field is included

Field	Description
P	Used with data messages only. Indicates priority of the data message (1 bit): 0 — no (this value is used for all control messages); 1 — yes
Version	The version of the message (4 bits): 2 — this is the latest version of the L2TP data message header; 1 — indicates an L2F packet as described in RFC 2341. Packets with an unknown version number are discarded.
Length	The total length (in bytes) of the L2TP message (16 bits).
Tunnel-ID	Identifies the L2TP tunnel (that is, the control connection). This number has local significance — each end gives the same tunnel different tunnel IDs. The ID refers to the receiver, not the sender, and is assigned during tunnel creation (16 bits).
Session-ID	Identifies the PPP session within a tunnel. This number has local significance — each end gives the same session different session IDs. The ID refers to the receiver, not the sender, and is assigned during session creation (16 bits).
Ns	The sequence number of the message. This is mandatory for control messages (to enable re-transmission of lost messages) but optional for data messages (to re-order data messages that were mis-sequenced during forwarding). The number, which starts at 0 and increments by 1, is assigned by an L2TP peer for each session in a tunnel (16 bits).
Nr	The sequence number of the next control message expected to be received. This is equal to the sequence number of last received control message plus 1. Used by the receiving peer to ensure that control messages are sent in order without duplication. In data messages, the field (if present as indicated by the S bit) is ignored (16 bits).
Offset Size	The location of the L2TP payload, expressed as the number of octets from the start of the message header (16 bits).
Offset pad	User-defined bytes used to pad the message header so that the payload starts at the location indicated by the Offset Size field (16 bits).

The AVP header consists of following fields (RFC 2611):

0	6	16	31
M	H	—	Length
Attribute Type			Vendor-ID
Attribute Type			Attribute Value

al_0513B

Table 27: AVP header fields and descriptions

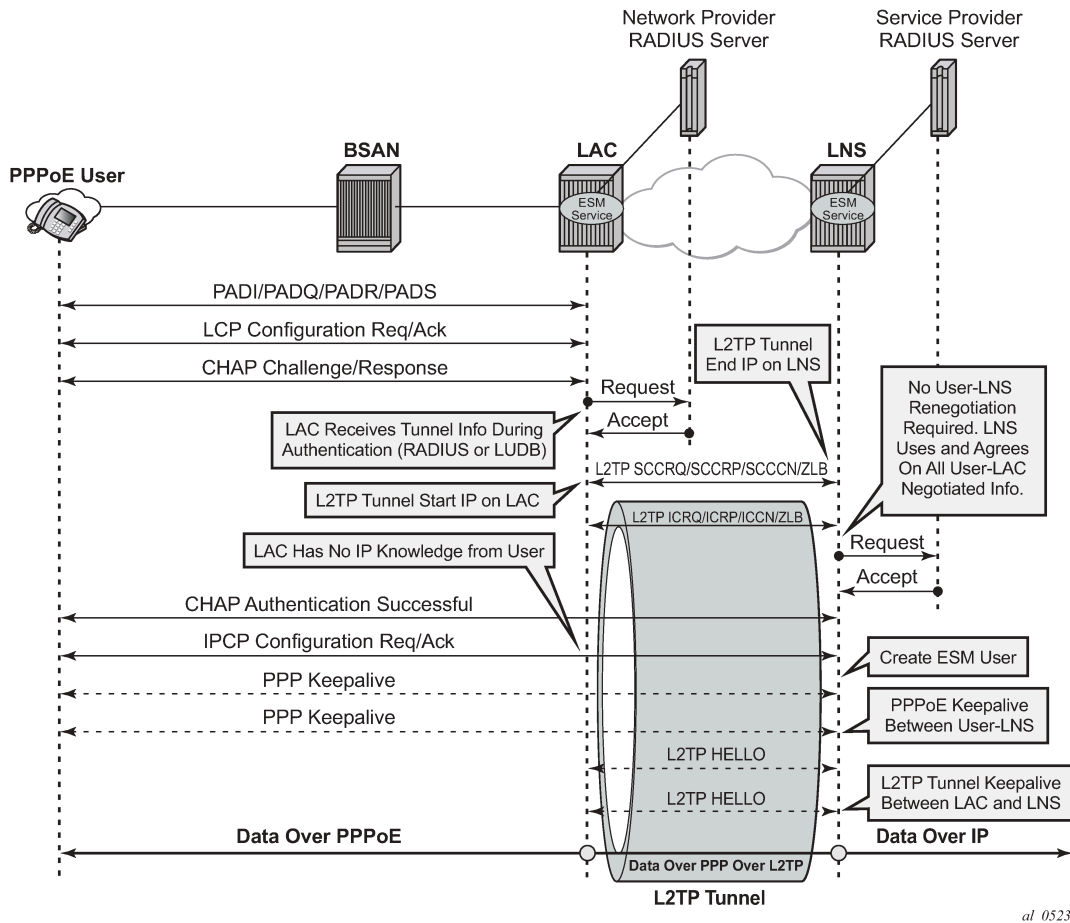
Field	Description
M	Mandatory bit — If the M bit is set on an unrecognized AVP within a message associated with a particular session, the session associated with this message MUST be terminated (1 bit).
H	Hidden bit — Identifies the hiding of data in the Attribute-Value field of an AVP. This capability can be used to avoid the passing of sensitive data, such as user passwords, as clear text in an AVP. The H-bit MUST only be set if a shared secret exists between the LAC and LNS. The shared secret is the same secret that is used for tunnel authentication. If the H-bit is set in any

Field	Description
	AVP(s) in a given control message, a Random Vector AVP must also be present in the message and MUST precede the first AVP having an H bit of 1 (1 bit).
-	Reserved for future use, must be set to zero (4 bits).
Length	Indicates the total number of bytes (including the overall length and bitmask fields) contained in this AVP (10 bits).
Vendor-id	Any vendor wishing to implement their own L2TP extensions can use their own Vendor ID along with private Attribute values. Vendor-ID=0 means that the standard AVPs are used (2 bytes).
Attribute type	A value with a unique interpretation across all AVPs defined under a given Vendor (2 bytes).
Attribute value	This is the actual value as indicated by the Vendor ID and Attribute Type (2 bytes).

RADIUS-triggered tunnel/session setup without LNS renegotiation

[Figure 109: RADIUS-triggered tunnel/session setup without LNS renegotiation](#) depicts the complete PPP session setup, using RADIUS authentication on both LAC and LNS. After the discovery phase (PADI/PADO/PADR/PADS) and LCP negotiation phase (LCP config_request/ack), the LAC initiates the L2TP tunnel setup based on RADIUS authentication information (RADIUS request/accept) and includes the negotiated PPP user-LAC information (called LCP proxy information). The LNS replies directly with a successful CHAP authentication if it agrees with the received proxy information. IP negotiation (IPCP config_request/ack) is handled between the user and the LNS, and the LAC has no IP knowledge of this PPP session.

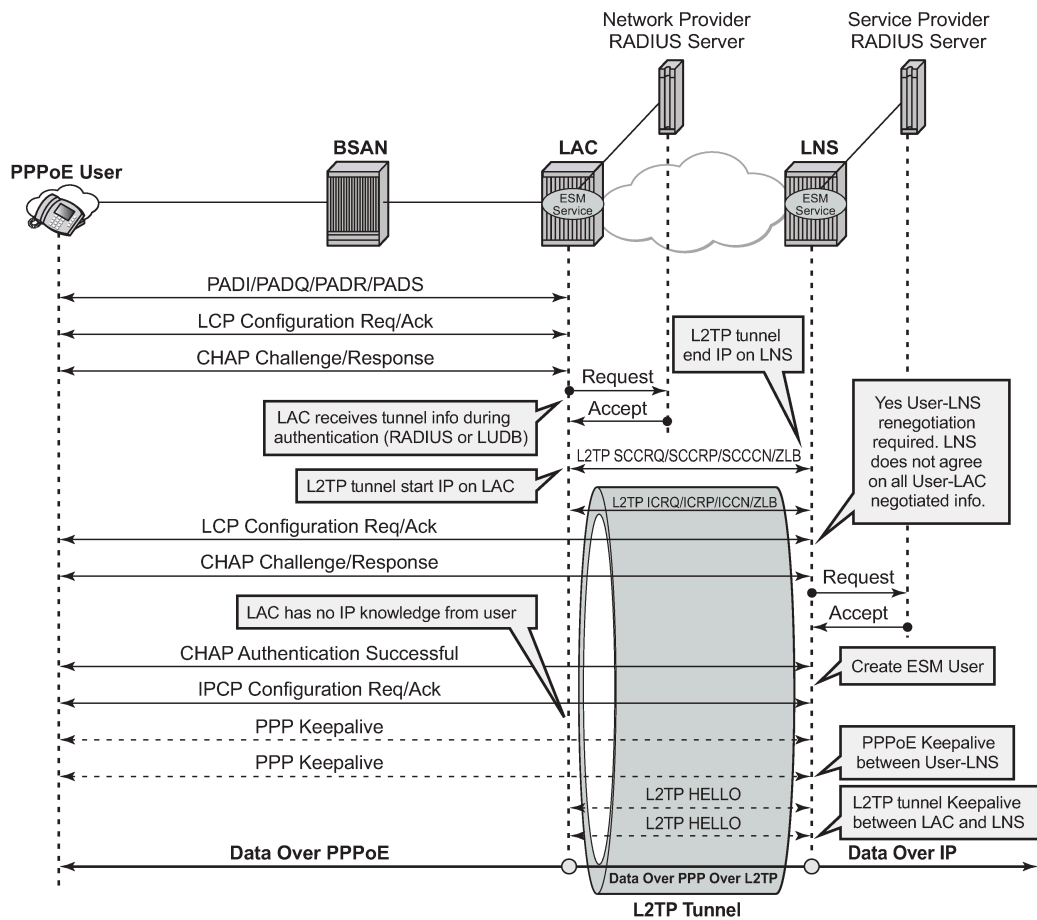
Figure 109: RADIUS-triggered tunnel/session setup without LNS renegotiation



RADIUS-triggered tunnel/session setup with LNS renegotiation

Figure 110: RADIUS-triggered tunnel/session setup with LNS renegotiation shows the scenario where the LNS does not agree with the received LCP proxy information and (re)starts the LCP phase (LCP config_request/ack) directly with the PPP user. The rest of this scenario is the same as shown in Figure 109: RADIUS-triggered tunnel/session setup without LNS renegotiation.

Figure 110: RADIUS-triggered tunnel/session setup with LNS renegotiation



al_0524

Running multiple PPP sessions over a single L2TP tunnel

Figure 111: Running multiple PPP sessions over a single L2TP tunnel shows multiple PPP sessions tunneled over a single L2TP Tunnel. The LAC encapsulates each PPP session with a different L2TP session-id (SID) but with the same L2TP Tunnel-id (TID).

The diagram illustrates the L2TP architecture and data flow. It shows three users (User 1, User 2, User 3) connected to a BSAN (Base Station Access Network). The BSAN is connected to a LAC (L2TP Access Concentrator), which is connected to a cloud representing the network. The cloud is connected to a LNS (L2TP Network Server). The LNS is connected to a L2TP Tunnel (represented by a large cylinder). The L2TP Tunnel is divided into three sections, each representing a different user's session. The sections are labeled: User 1 | SID-1 | TID-1, User 2 | SID-2 | TID-1, and User 3 | SID-3 | TID-1. The L2TP Tunnel is also labeled L2TP Tunnel-id 1. The diagram shows the flow of data from the users through the BSAN, LAC, and LNS to the L2TP Tunnel, and then to the destination. The flow is as follows: User 1 sends PPPoE User 1 Setup to the LAC, which then sends L2TP SCCRQ/SCCRP/SCCN to the LNS. The LNS responds with Tunnel-id 1 : TID-1. The LAC then sends L2TP Tunnel Setup and L2TP Session Setup to the LNS, which responds with Session-id 1 : SID-1. The LAC then sends Data Over PPPoE User 1 to the LNS, which sends L2TP ICRQ/ICRP/ICCN to the L2TP Tunnel. The L2TP Tunnel then sends Data Over PPPoL2TP to the LNS, which sends Data to the destination. Similar flows are shown for User 2 and User 3. A callout box indicates that 1 LT2P tunnel has many L2TP sessions.

Figure 112: PPP user-initiated release/terminate shows the user initiated terminate_request tunneled by the LAC followed by the user initiated PADT terminated on the LAC. The LAC informs the LNS about the termination of the session via the L2TP CDN message. The L2TP tunnel can be optionally (idle-timeout) terminated via the L2TP StopCCN message.

The diagram illustrates the termination of an L2TP tunnel between a BSAN (Base Station Access Network) and an LNS (Local Network Server) via an LAC (Local Access Concentrator). The process involves the following steps:

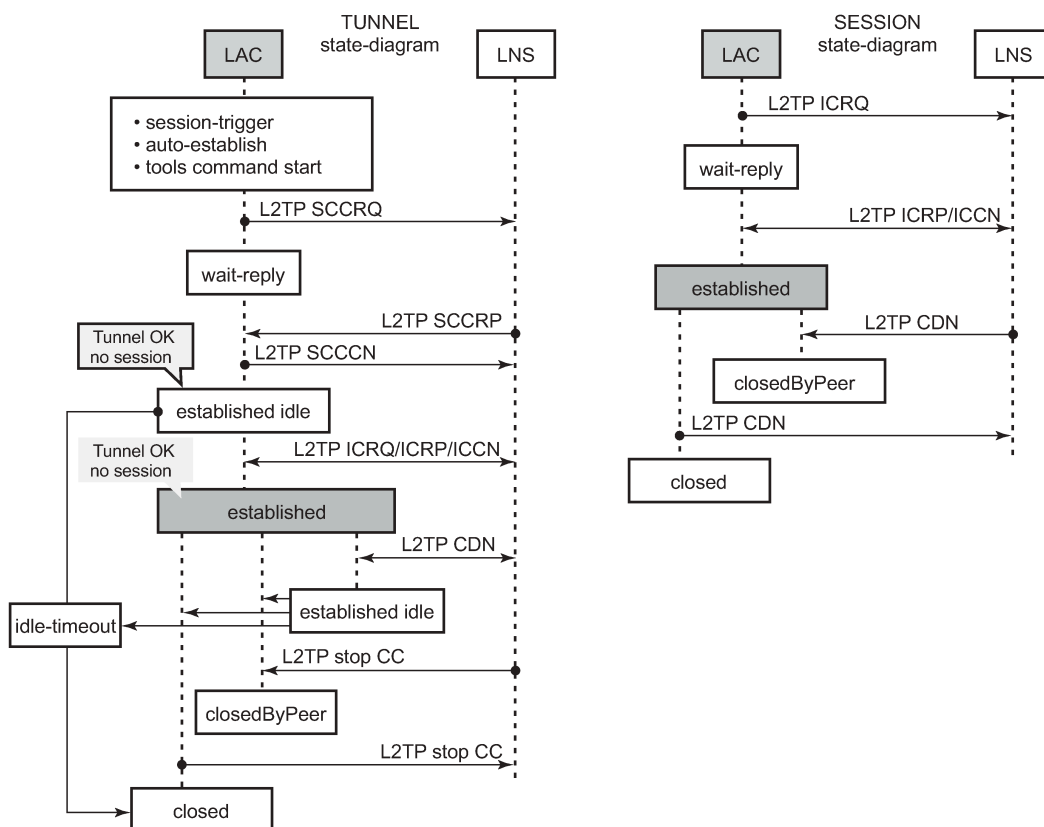
- PPPoE User** (represented by a mobile phone icon) initiates the process.
- BSAN** (Base Station Access Network) sends a **PPPoE LCP Terminate Request** to the **LAC**.
- LAC** (Local Access Concentrator) sends a **PPPoE LCP Terminate Request** to the **LNS**.
- LNS** (Local Network Server) sends a **PADT** (PPP over ATM Data Transfer) packet to the **LAC**.
- LAC** sends a **PADT** packet to the **PPPoE User**.
- Idle-timeout (See Further)** is indicated by a bracket on the L2TP Tunnel.
- The **L2TP Tunnel** is shown as a cylinder containing **L2TP CDN** (L2TP Control Network Data) and **L2TP StopCCN** (L2TP Stop Control Network Data).

L2TP tunnel/session state diagram

Figure 113: L2TP tunnel and session state diagram gives an overview of the main L2TP tunnel and session states. An L2TP tunnel in the establishedIdle state is a tunnel without sessions. A **tools** command (see [Advanced topics](#)) can put an L2TP tunnel in a draining state (this prevents adding new sessions on tunnel but leaves the current sessions intact) or in a drained state (moved from draining to drained when all sessions terminated). The draining and drained state are not shown in the state diagram.

The L2TP tunnel setup occurs first with the triggers being: session activation, auto-establish, and a **tools start** command (see the [Advanced topics](#) section). An L2TP session setup trigger is always session based.

Figure 113: L2TP tunnel and session state diagram



al_0515

Configuration

The following scenarios are configured:

- [Scenario 1: RADIUS-derived L2TP parameters](#)
- [Scenario 2: node-derived L2TP parameters](#)

Scenario 1: RADIUS-derived L2TP parameters

In the first scenario, the LAC receives an incoming connection and contacts the LAC RADIUS server. The RADIUS server retrieves the attributes for the user's domain (for example @wholesale.com) and passes the tunnel attributes to the LAC. Based on these RADIUS provided tunnel attributes, the LAC selects or initiates a new tunnel to the LTS or directly to the LNS. Once the tunnel is established, the LNS authenticates the end user using its own RADIUS server. Configuring the LNS and the LTS are out of the scope of this example.

In a RADIUS driven L2TP setup, either all or some of the required L2TP attributes are returned via RADIUS. If the RADIUS server only returns the L2TP [67] Tunnel-Server-Endpoint attributes, then the L2TP tunnel/session is established using the 'l2tp node parameter values' for the other required L2TP parameters. The 'l2tp node parameters' are defined under the configure router/service l2tp hierarchy. If the RADIUS server does not return all of the L2TP attributes and the node values are not configured, then the system falls back to default settings for these L2TP parameters.

The standard and vendor specific [26-6572] L2TP RADIUS attributes are listed in the tables below, together with the corresponding l2tp node parameters and defaults.

Table 28: Generic L2TP RADIUS attributes

Attribute ID	Attribute Name	Mandatory	CLI Node Parameter	Corresponding Defaults	
64	Tunnel-Type	Y	-	-	-
65	Tunnel-Medium-Type	Y	-	-	-
66	Tunnel-Client-Endpoint: [0-31]	N	local-address	no local-address	system-ip
67	Tunnel-Server-Endpoint	N	-	-	-
69	Tunnel-Password	N	password	no password	-
82	Tunnel-Assignment-ID:0	N	-	-	default_radius_group
82	Tunnel-Assignment-ID: [1..31]	N	-	-	Unnamed
83	Tunnel-Preference	N	preference	no preference	50
90	Tunnel-Client-Auth-ID	N	local-name	no local-name	system-name
91	Tunnel-Server-Auth-ID	N	-	-	-

Table 29: Nokia Specific L2TP RADIUS attributes


26-6527	Attribute Name	Mandatory	CLI Node Parameter	Corresponding Defaults	
-46	Alc-Tunnel-Group	N	-	-	-

26-6527	Attribute Name	Mandatory	CLI Node Parameter	Corresponding Defaults	
-47	Alc-Tunnel-Algorithm	N	session-assign-method	no session-assign-method	existingFirst
-48	Alc-Tunnel-Max-Sessions:0	N	-	group-session-limit	131071
-48	Alc-Tunnel-Max-Sessions:[1..31]	N	-	tunnel-session-limit	32767
-49	Alc-Tunnel-Idle-Timeout	N	idle-timeout	no idle-timeout	Infinite
-50	Alc-Tunnel-Hello-Interval	N	hello-interval	no hello-interval	300 sec
-51	Alc-Tunnel-Destruct-Timeout	N	destruct-timeout	no destruct-timeout	60 sec
-52	Alc-Tunnel-Max-Retries-Estab	N	max-retries-estab	no max-retries-estab	5
-53	Alc-Tunnel-Max-Retries-Not-Estab	N	max-retries-not-estab	no max-retries-not-estab	5
-54	Alc-Tunnel-AVP-Hiding	N	avp-hiding	no avp-hiding	Never
-97	Alc-Tunnel-Challenge	N	challenge	no challenge	Never
-104	Alc-Tunnel-Serv-Id	N	-	-	Base
-120	Alc-Tunnel-Rx-Window-Size	N	receive-window-size	no receive-window-size	64
-144	Alc-Tunnel-Acct-Policy	N	radius-accounting-policy	no radius-accounting-policy	-

Base router hosted LAC with single endpoint/single tunnel

Using the mandatory L2TP RADIUS attributes (see the following RADIUS user file) the LAC initiates an L2TP tunnel, as shown in [Figure 114: Base router hosted LAC with single endpoint/single tunnel](#). The source address for the tunnel is the IPv4 address of a loopback interface in the base router system (LAC tunnel endpoint). The destination for the tunnel is defined by the Tunnel-Server-Endpoint RADIUS attribute [67], and is also known as the peer tunnel LNS endpoint address.

Figure 114: Base router hosted LAC with single endpoint/single tunnel



The PPPoE user terminates on IES service 1, sap 1/1/3:100, and is authenticated via RADIUS **authentication-policy authentication-1** which provides wholesale/retail (L2TP) information.

```
configure
service
  ies 1 customer 1 create
    subscriber-interface "sub-l2tp" create
    unnumbered "system"
    group-interface "grp-l2tp" create
    authentication-policy "radius-1"
    sap 1/1/3:100 create
      sub-sla-mgmt
      sub-ident-policy "all-subscribers"
      multi-sub-sap 1000
      no shutdown
    exit
  exit
  pppoe
    sap-session-limit 10
    no shutdown
  exit
exit
exit
no shutdown
exit
exit
exit
exit
```

The excerpt from the FreeRADIUS users file below shows the attributes to be returned.

```
user1@wholesale.com    Cleartext-Password := "letmein", NAS-Identifier == "LAC"
                        Alc-Subsc-ID-Str = "%{User-name}",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Tunnel-Type:1 += L2TP,
                        Tunnel-Medium-Type:1 += IP,
                        Tunnel-Server-Endpoint:1 += 192.168.22.2,
```

L2TP is enabled (**no shutdown**) in the related service instance.

The L2TP tunnel is set up in the base instance and not in a VRF because the attribute Alc-Tunnel-Serv-Id is not returned from RADIUS.

Missing L2TP parameters are taken from defaults defined in the router l2tp context.

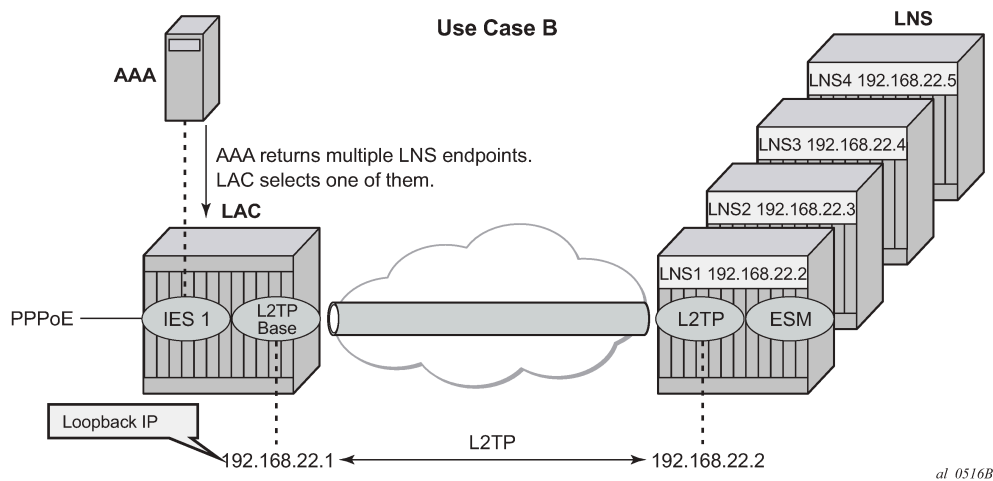
```
configure router l2tp
  calling-number-format "%S %s" # L2TP AVP 22 format
                                # Default format 'system-name sap-id'
  ---snip---
  no local-name                 # default name equals system-name
  no max-retries-estab          # default value equals 5
  ---snip---
  no shutdown                   # enable L2TP
```

This scenario shows the PPPoE session termination (base IES service 1) and the L2TP tunnel setup in the base router instance.

Base router hosted LAC with multiple endpoints

Figure 115: Base router hosted LAC with multiple endpoints shows a scenario with the PPPoE session termination (base IES service 1) and the L2TP tunnel setup in the base router instance.

Figure 115: Base router hosted LAC with multiple endpoints



The following excerpt from the FreeRADIUS users file shows that user *user1@wholesale.com* has four possible endpoints (LNS), each with its own tunnel preference. The LAC selects one L2TP endpoint out of these four tunnel specifications according to the configured L2TP selection process. This use case uses weighted load balancing between RADIUS-tunnel-1 and RADIUS-tunnel-2. The L2TP tunnel selection process is out of the scope of this chapter.

```
user1@wholesale.com    Cleartext-Password := "letmein", NAS-Identifier == "LAC"
                      Alc-Subsc-ID-Str = "%{User-name}",
                      Alc-Subsc-Prof-Str = "sub-profile-1",
                      Alc-SLA-Prof-Str = "sla-profile-1",
# group related info
                      Tunnel-Client-Endpoint:0 = 192.168.22.1,
                      Alc-Tunnel-Algorithm:0 = weighted-access,
```

```
Tunnel-Client-Auth-Id:0 = "lac-pe1",
Tunnel-Assignment-Id:0 = "RADIUS-group",
Alc-Tunnel-Max-Retries-Estab:0 = 2,

# tunnel-1 related info
Tunnel-Type:1 += L2TP,
Tunnel-Medium-Type:1 += IP,
Tunnel-Server-Endpoint:1 += 192.168.22.2,
Tunnel-Assignment-Id:1 += "RADIUS-tunnel-1",
Tunnel-Preference:1 += 100,

# tunnel-2 related info
Tunnel-Type:2 += L2TP,
Tunnel-Medium-Type:2 += IP,
Tunnel-Server-Endpoint:2 += 192.168.22.3,
Tunnel-Assignment-Id:2 += "RADIUS-tunnel-2",
Tunnel-Preference:2 += 80,

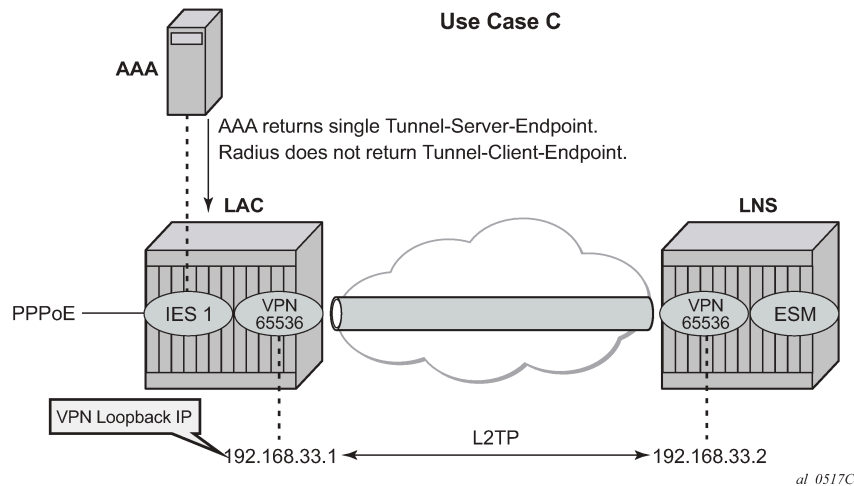
# tunnel-3 related info
Tunnel-Type:3 += L2TP,
Tunnel-Medium-Type:3 += IP,
Tunnel-Server-Endpoint:3 += 192.168.22.4,
Tunnel-Assignment-Id:3 += "RADIUS-tunnel-3",

---snip---
```

VRF hosted LAC

Figure 116: VRF hosted LAC shows the PPPoE session termination (base IES service 1) and the L2TP tunnel setup in a different router instance (VPRN 65536).

Figure 116: VRF hosted LAC



Using the following L2TP RADIUS attributes, the LAC initiates an L2TP tunnel in VPRN 65536. The PPPoE session is still handled by IES service 1, which proves that both router instances can be different. (See use-case A for configuration details of IES service 1).

```
user1@wholesale.com    Cleartext-Password := "letmein", NAS-Identifier == "LAC"
                        Alc-Subsc-ID-Str = "%{User-name}",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-Tunnel-Serv-Id = 65536,
```



```
Tunnel-Client-Auth-Id:0 = "lac-pe1",  
Tunnel-Assignment-Id:0 = "RADIUS-returned-TG",  
Tunnel-Type:1 += L2TP,  
Tunnel-Medium-Type:1 +=IP,  
Tunnel-Server-Endpoint:1 += 192.168.33.2,  
Tunnel-Assignment-Id:1 += "RADIUS-returned-TN",
```

If RADIUS does not return the L2TP source IP address (Tunnel-Client-Endpoint), then the IP address from the VPRN 65536 interface named 'system' is used as the L2TP source address. The tunnel setup fails if this system interface does not exist.

```
configure  
  service  
    vprn 65536  
      interface "system" create  
        address 192.168.33.1/32  
        loopback  
      exit  
    l2tp  
      no shutdown  
    exit
```

Scenario 2: node-derived L2TP parameters


In the second scenario, the LAC receives the incoming connection and an L2TP tunnel-group-name is assigned during LUDB or RADIUS authentication. This tunnel-group-name refers to the CLI preconfigured tunnel-group name context (**configure router <router-name> l2tp group <tunnel-group-name>**), which provides the context for all relevant tunnel attributes.

Based on these attributes, the LAC selects and initiates a tunnel to the LTS or directly to the LNS as in [Scenario 1: RADIUS-derived L2TP parameters](#).

RADIUS returns L2TP tunnel group

[Figure 117: RADIUS returns L2TP tunnel group](#) shows use case D, where the L2TP tunnel-group-name is assigned during RADIUS authentication.

Figure 117: RADIUS returns L2TP tunnel group



```
user1@wholesale.com    Cleartext-Password := "letmein", NAS-Identifier == "LAC"
                        Alc-Subsc-ID-Str = "%{User-name}",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-Tunnel-Serv-Id = 65536,
                        Alc-Tunnel-Group = "wholesale.com",
```

The L2TP tunnel is initiated from VPRN 65536 (Alc-Tunnel-Serv-Id) and all L2TP tunnel information is taken from the l2tp group wholesale.com hierarchy (Alc-Tunnel-Group) as defined on the node.


```
configure
  service
    vprn 65536
    ---snip---
    interface "system" create
      address 192.168.33.1/32
      loopback
    exit
    l2tp
      group "wholesale.com" create
        tunnel "wholesale.com" create
          local-address 192.168.33.1
          local-name "lac-pe1"
          peer 192.168.33.2
          no auto-establish
          no shutdown
        exit
        no shutdown
      exit
      no shutdown
    exit
    no shutdown
  exit
exit
```

An L2TP tunnel is set up by either a PPP session-trigger, a **tools** command or by the l2tp group tunnel auto-establish parameter configuration. See the [Advanced topics](#) section for the non-session-triggered tunnel setup.

LUDB returns L2TP tunnel group

[Figure 118: LUDB returns L2TP tunnel group](#) shows use case E, where the L2TP tunnel-group-name is assigned during LUDB authentication, so this essentially is a RADIUS-less scenario.

Figure 118: LUDB returns L2TP tunnel group



al_0518

The PPPoE user enters on an IES service 1, sap 1/1/3:100, and is authenticated via the LUDB which provides L2TP wholesale/retail and ESM information. The PPPoE context refers to a local-user database *l2tp* to provide the subscriber authentication and the tunnel setup parameters, so no RADIUS is needed.

```
configure
service
  ies 1 customer 1 create
    subscriber-interface "sub-l2tp" create
      unnumbered "system"
    group-interface "grp-l2tp" create
      sap 1/1/3:100 create
        sub-sla-mgmt
        sub-ident-policy "all-subscribers"
        multi-sub-sap 1000
        no shutdown
      exit
    exit
  exit
  pppoe
    sap-session-limit 10
    user-db "l2tp"
    no shutdown
  exit
exit
```

```
        exit
      no shutdown
    exit
  exit
exit
```

The referenced local user database *l2tp* configuration provides all of the required L2TP and ESM information.

```
configure
  subscriber-mgmt
    local-user-db "l2tp" create
    ppp
      match-list username
      host "wholesale.com" create
      host-identification
        username "wholesale.com" domain-only
      exit
      password ignore
      identification-strings 254 create
        subscriber-id "user@wholesale.com"
        sla-profile-string "sla-profile-1"
        sub-profile-string "sub-profile-1"
      exit
      l2tp
        group "wholesale.com" service-id 65536
      exit
      no shutdown
    exit
  exit
  no shutdown
exit
exit
exit
```

Operation and troubleshooting

The subsequent sections explain how the use cases A to E described in the configuration section are verified using `show`, `debug`, and `tools` commands.

The standard router debugging tools can be used to monitor and troubleshoot the L2TP tunnel and session setup.

Useful **show** commands are:

```
show service id <service-id> ppp session [detail]
show router l2tp tunnel [detail]
show router l2tp session [detail]
show router l2tp peer [ip-address]
```

To debug and show PPPoE packets:

```
debug service id <service-id> ppp packet mode egr-ingr-and-dropped
debug service id <service-id> ppp packet detail-level medium
```

To debug and show RADIUS authentication:

```
debug router radius packet-type authentication
```

To debug and show LUDB authentication:

```
debug subscriber-mgmt local-user-db <local-user-db-name> detail all
```

To debug and show the LAC tunnel selection process and L2TP state machine:

```
debug router l2tp event lac-session-setup  
debug router l2tp event finite-state-machine
```

To debug and show the L2TP tunnel and session setup:

```
debug router l2tp packet direction both  
debug router l2tp packet detail-level high
```

Understanding the L2TP debug output

The following L2TP ICRQ message (**debug router l2tp packet**) is used to explain how the displayed debug output should be interpreted. See [Recapitulation of the L2TPv2 protocol -L2TP header and AVP layout](#) for more details.

```
19 2019/05/21 14:02:10.811 CEST MINOR: DEBUG #2001 Base L2TP(v2, ctrl, egress)  
"L2TP(v2, ctrl, egress): UDP 192.0.2.1:1701 -> 192.168.22.2:1701  
tunnel 13008 session 0, ns 2 nr 1, flags:, reserved=0  
  AVP MessageType(0,0), flags: mandatory, reserved=0  
    IncomingCallRequest(10)  
  AVP AssignedSessionId(0,14), flags: mandatory, reserved=0  
    26540  
  AVP CallSerialNumber(0,15), flags: mandatory, reserved=0  
    25216  
  AVP CallingNumber(0,22), flags: mandatory, reserved=0  
    "LAC 1/1/3:100"  
  AVP AgentCircuitId(3561,1), flags:, reserved=0  
    "circuit0"  
  AVP AgentRemoteId(3561,2), flags:, reserved=0  
    "remote0"  
  AVP ActDataRateUp(3561,129), flags:, reserved=0  
    2000000  
  AVP ActDataRateDown(3561,130), flags:, reserved=0  
    4000000"
```

- L2TP(v2, ctrl, egress): UDP 192.0.2.1:1701 -> 192.168.22.2:1701
 - version: v2
 - type field (T-bit): control message (ctrl)
 - 192.0.2.1:1701 -> 192.168.22.2:1701
 - 192.0.2.1:1701 - source tunnel-end-point:source udp port
 - 192.168.22.2:1701 - destination tunnel-end-point:destination udp port
- tunnel 13008 session 0, ns 2 nr 1, flags:, reserved=0
 - tunnel-id: 13008
 - session-id: 0
 - ns:2

- nr:1
- flags: 0 (refers to T/L/S/O/P bits L2TP header)
- reserved field:0
- AVP CallingNumber(0,22), flags: mandatory, reserved=0
 - AVP MessageType(0,22): "LAC 1/1/3:100"
 - Vendor-id: 0 - Standard Attribute
 - Attribute Type: 22 – Calling Number AVP
 - Attribute Value: "LAC 1/1/3:100"

Scenario 1: RADIUS-derived L2TP parameters

Base router hosted LAC with single endpoint/single tunnel

The **debug service id <service-id> ppp packet mode egr-ingr-and-dropped** command shows the PPPoE packet exchange. The following PADI packet shows the service, SAP, and received PPPoE tags. The received PPPoE DSL forum tags are by default copied during the LAC L2TP tunnel setup into the Incoming Call Request (ICRQ) DSL Forum AVPs (RFC 5515).

```
1 2019/05/21 14:02:10.779 CEST MINOR: DEBUG #2001 Base PPPoE
"PPPoE: RX Packet
  IES 1, SAP 1/1/3:100

  DMAC: ff:ff:ff:ff:ff:ff
  SMAC: 00:00:00:00:01:01
  Ether Type: 0x8863 (Discovery)

  PPPoE Header:
  Version: 1                      Type      : 1
  Code   : 0x09 (PADI)           Session-Id: 0x0000 (0)
  Length : 48

  PPPoE Tags:
  [0x0101] Service-Name: ""
  [0x0103] Host-Uniq: len = 1, value = 31
  [0x0105] Vendor-Specific: vendor-id = 0x0de9 (ADSL Forum)
    [0x01] Agent-Circuit-Id: "circuit0"
    [0x02] Agent-Remote-Id: "remote0"
    [0x81] Actual-Upstream: 2000
    [0x82] Actual-Downstream: 4000
"
```

The **debug router radius packet-type authentication** command shows the actual authentication parameters returned by RADIUS. This example returns the minimum set of L2TP related RADIUS attributes.

```
12 2019/05/21 14:02:10.806 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Access-Accept(2) id 16 len 89 from 172.16.1.11:1812 vrid 1 pol rsp-radius-1
    VSA [26] 15 Nokia(6527)
      SUBSC PROF STR [12] 13 sub-profile-1
    VSA [26] 15 Nokia(6527)
```

```
SLA PROF STR [13] 13 sla-profile-1
TUNNEL TYPE [64] 4 1 L2TP(3)
TUNNEL MEDIUM TYPE [65] 4 1 IPv4(1)
TUNNEL SERVER ENDPOINT [67] 13 1 192.168.22.2
"
```

The **debug router l2tp event lac-session-setup** command shows the LAC tunnel selection for this example. An L2TP group-name *default_radius_group* with tunnel-name *unnamed* is created in this case, because RADIUS did not return an explicit group and tunnel name.

```
13 2019/05/21 14:02:10.808 CEST MINOR: DEBUG #2001 Base PPPoE 13->L2TP
"PPPoE 13->L2TP: UDP 192.0.2.1:1701 -> 192.168.22.2:1701
preference 50 tunnel default_radius_group:unnamed
request to open new tunnel 1640"
```

```
14 2019/05/21 14:02:10.808 CEST MINOR: DEBUG #2001 Base PPPoE 13->L2TP
"PPPoE 13->L2TP: UDP 192.0.2.1:1701 -> 192.168.22.2:1701
preference 50 tunnel default_radius_group:unnamed
create session 107505580"
```

The **debug router l2tp packet detail-level** command shows the L2TP tunnel and session setup for this example.

For the tunnel setup, the LAC sends a start-control-connection-request (SCCRQ) containing the assigned tunnel ID (no tunnel authentication in the example). The tunnel is now in a wait-reply state.

```
15 2019/05/21 14:02:10.808 CEST MINOR: DEBUG #2001 Base L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.0.2.1:1701 -> 192.168.22.2:1701
tunnel 0 session 0, ns 0 nr 0, flags:, reserved=0
  AVP MessageType(0,0), flags: mandatory, reserved=0
    StartControlConnectionRequest(1)
  AVP ProtocolVersion(0,2), flags: mandatory, reserved=0
    version=1, revision=0
  AVP HostName(0,7), flags: mandatory, reserved=0
    "lac-pe1"
  AVP WindowSize(0,10), flags: mandatory, reserved=0
    64
  AVP FramingCapabilities(0,3), flags: mandatory, reserved=0
    sync=no, async=no
  AVP BearerCapabilities(0,4), flags: mandatory, reserved=0
    digital=yes, analogue=no
  AVP FirmwareRevision(0,6), flags:, reserved=0
    4869
  AVP VendorName(0,8), flags:, reserved=0
    "Nokia"
  AVP AssignedTunnelId(0,9), flags: mandatory, reserved=0
    1640"
```

The LNS can bring up the tunnel, so the LNS replies with a start-control-connection-reply (SCCRP) including the assigned tunnel-id.

```
16 2019/05/21 14:02:10.809 CEST MINOR: DEBUG #2001 Base L2TP(v2, ctrl, ingress)
"L2TP(v2, ctrl, ingress): UDP 192.168.22.2:1701 -> 192.0.2.1:1701
tunnel 1640 session 0, ns 0 nr 1, flags:, reserved=0
  AVP MessageType(0,0), flags: mandatory, reserved=0
    StartControlConnectionReply(2)
  AVP ProtocolVersion(0,2), flags: mandatory, reserved=0
    version=1, revision=0
  AVP HostName(0,7), flags: mandatory, reserved=0
    "lns-pe2"
```

```

AVP WindowSize(0,10), flags: mandatory, reserved=0
64
AVP FramingCapabilities(0,3), flags: mandatory, reserved=0
sync=no, async=no
AVP BearerCapabilities(0,4), flags: mandatory, reserved=0
digital=yes, analogue=no
AVP FirmwareRevision(0,6), flags:, reserved=0
4869
AVP VendorName(0,8), flags:, reserved=0
"Nokia"
AVP AssignedTunnelId(0,9), flags: mandatory, reserved=0
13008"

```

As the last step in the tunnel setup phase, the LAC responds with a start-control-connection-connected (SCCCN) message. After an LNS ZLB acknowledgment, the tunnel is in the establishedIdle state.

```

17 2019/05/21 14:02:10.810 CEST MINOR: DEBUG #2001 Base L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.0.2.1:1701 -> 192.168.22.2:1701
tunnel 13008 session 0, ns 1 nr 1, flags:, reserved=0
  AVP MessageType(0,0), flags: mandatory, reserved=0
    StartControlConnectionConnected(3)"

```

Once the tunnel exists the session setup starts, a three-way exchange for session establishment within the tunnel is performed. The LAC sends an incoming-call-request (ICRQ) with the parameter information for the session. The session is now in the wait-reply state.

```

19 2019/05/21 14:02:10.811 CEST MINOR: DEBUG #2001 Base L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.0.2.1:1701 -> 192.168.22.2:1701
tunnel 13008 session 0, ns 2 nr 1, flags:, reserved=0
  AVP MessageType(0,0), flags: mandatory, reserved=0
    IncomingCallRequest(10)
  AVP AssignedSessionId(0,14), flags: mandatory, reserved=0
    26540
  AVP CallSerialNumber(0,15), flags: mandatory, reserved=0
    25216
  AVP CallingNumber(0,22), flags: mandatory, reserved=0
    "LAC 1/1/3:100"
  AVP AgentCircuitId(3561,1), flags:, reserved=0
    "circuit0"
  AVP AgentRemoteId(3561,2), flags:, reserved=0
    "remote0"
  AVP ActDataRateUp(3561,129), flags:, reserved=0
    2000000
  AVP ActDataRateDown(3561,130), flags:, reserved=0
    4000000"

```

The LNS then sends an incoming-call-reply (ICRP) that contains the assigned session ID. The session is now in the connect state.

```

21 2019/05/21 14:02:10.813 CEST MINOR: DEBUG #2001 Base L2TP(v2, ctrl, ingress)
"L2TP(v2, ctrl, ingress): UDP 192.168.22.2:1701 -> 192.0.2.1:1701
tunnel 1640 session 26540, ns 1 nr 3, flags:, reserved=0
  AVP MessageType(0,0), flags: mandatory, reserved=0
    IncomingCallReply(11)
  AVP AssignedSessionId(0,14), flags: mandatory, reserved=0
    2466"

```

Finally, the LAC sends an incoming call connected (ICCN) and provides the LNS with additional information from the user initiated session. This information includes the LCP information from the negotiation that the LAC and remote user performed. This information is used by the LNS to decide

whether to start LCP re-negotiation and/or authentication re-negotiation with the PPP user or not. After an LNS ZLB acknowledgment, the session is in the established state.

```
24 2019/05/21 14:02:10.814 CEST MINOR: DEBUG #2001 Base L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.0.2.1:1701 -> 192.168.22.2:1701
tunnel 13008 session 2466, ns 3 nr 2, flags:, reserved=0
  AVP MessageType(0,0), flags: mandatory, reserved=0
    IncomingCallConnected(12)
  AVP FramingType(0,19), flags: mandatory, reserved=0
    sync=no, async=no
  AVP TxConnectSpeed(0,24), flags: mandatory, reserved=0
    4294967295
  AVP InitialRxLcpConfReq(0,26), flags:, reserved=0
    01 04 05 d4
    [1] MRU: 1492
  AVP LastTxLcpConfReq(0,27), flags:, reserved=0
    01 04 05 d4 03 05 c2 23 05 05 06 75 25 ad d3
    [1] MRU: 1492
    [3] Authentication-Protocol: 0xc223 (CHAP), Algorithm = 5 (MD5)
    [5] Magic-Number: 0x7525add3
  AVP LastRxLcpConfReq(0,28), flags:, reserved=0
    01 04 05 d4
    [1] MRU: 1492
  AVP ProxyAuthenType(0,29), flags:, reserved=0
    chap(2)
  AVP ProxyAuthenName(0,30), flags:, reserved=0
    "user1@wholesale.com"
  AVP ProxyAuthenChallenge(0,31), flags:, reserved=0
    13 ba fc db 18 15 b5 21 03 c9 61 8d 8a 1b 43 00
    c3 4a 80 51 df 52 f4 06 26 c8 16 db ce 2b 7d 62
    e5 7a bd 7d 0f
  AVP ProxyAuthenId(0,32), flags:, reserved=0
    id=1, reserved=0
  AVP ProxyAuthenResponse(0,33), flags:, reserved=0
    da c4 40 35 e4 4b 3f 72 3f eb 84 7b 09 99 5d f7
  AVP RxConnectSpeed(0,38), flags:, reserved=0
    4294967295"
```

The operational PPPoE session information for the IES 1 (base router) instance is as follows.

```
*A:LAC# show service id 1 ppp session

=====
PPP sessions for service 1
=====
User-Name
  Descr.
      Up Time      Type  Termination      IP/L2TP-Id/Interface-Id MC-Stdby
-----
user1@wholesale.com
  svc:1 sap:1/1/3:100 mac:00:00:00:00:01:01 sid:1
      0d 00:10:01   oE    lac          107505580
-----
No. of PPP sessions: 1
=====
*A:LAC#
```

The operational tunnel information in the base instance shows that the tunnel is established.

```
*A:LAC# show router l2tp tunnel

=====
Conn ID   Loc-Tu-ID Rem-Tu-ID State                Blacklist-state   Ses Active
-----
```

Group Assignment	Ses Total
-----	-----
107479040 1640 13008 established not-blacklisted	1
default_radius_group	1
unnamed	
-----	-----
No. of tunnels: 1	
=====	=====
*A:LAC#	

Detailed operational tunnel information is obtained using following command.

```
*A:LAC# show router l2tp tunnel tunnel-id 1640 detail

=====
L2TP Tunnel 107479040
=====

Connection ID: 107479040
Protocol      : v2
State        : established
IP           : 192.0.2.1
UDP          : 1701
Peer IP      : 192.168.22.2
Peer UDP     : 1701
Tx dst-IP    : 192.168.22.2
Tx dst-UDP   : 1701
Rx src-IP    : 192.168.22.2
Rx src-UDP   : 1701
Name         : lac-pe1
Remote Name  : lns-pe2
Assignment ID: unnamed
Group Name   : default_radius_group
Acct. Policy : N/A
Error Message: N/A

Tunnel ID      : 1640
Preference     : 50
Hello Interval (s): 300
Idle T0 (s)    : infinite
Max Retr Estab : 5
Cfg'd Sess Limit : unlimited
Transport Type : udpIp
Time Started   : 05/21/2019 14:02:11
Time Established : 05/21/2019 14:02:11
Stop CCN Result : noError
Blacklist-state : not-blacklisted
Set Dont Fragment : true

Remote Conn ID : 852492288
Remote Tunnel ID : 13008
Receive Window  : 64
AVP Hiding      : never
Destruct T0 (s) : 60
Max Retr Not Estab: 5
Oper Session Limit: 32767
Challenge       : never
Time Idle       : N/A
Time Closed     : N/A
General Error   : noError

Failover
State          : not-recoverable
Recovery Conn ID : N/A
Recovery state  : not-applicable
Recovered Conn ID : N/A
Recovery method : mcs
Track SRRP     : (Not specified)
Ctrl msg behavior : handle
Recovery time (ms) :
Requested      : N/A
Peer           : N/A
-----
```

```
-----
No. of tunnels: 1
=====
*A:LAC#
```

The operational L2TP session information shows the L2TP session is established.

```
*A:LAC# show router l2tp session

=====
L2TP Session Summary
=====
ID                Control Conn ID    Tunnel-ID    Session-ID    State
-----
107505580         107479040          1640         26540         established
-----
No. of sessions: 1
=====
*A:LAC#
```

For detailed operational L2TP session information use the following command.

```
*A:LAC# show router l2tp session session-id 26540 detail

=====
L2TP Session 107505580
=====

Connection ID: 107505580
State          : established
Tunnel Group   : default_radius_group
Assignment ID: unnamed
Error Message: N/A

Control Conn ID : 107479040          Rem Cntrl Conn ID : 852492288
Tunnel ID       : 1640              Remote Tunnel ID  : 13008
Session ID      : 26540             Remote Session ID : 2466
PW Type         : ppp               Remote Conn ID    : 852494754
Time Started    : 05/21/2019 14:02:11
Time Established : 05/21/2019 14:02:11 Time Closed       : N/A
CDN Result      : noError           General Error      : noError
-----
No. of sessions: 1
=====
*A:LAC#
```

Base router hosted LAC with multiple endpoints

The **debug router radius packet-type authentication** command shows the actual RADIUS authentication parameters returned. This example returns multiple tunnel endpoints from which the LAC selects one. This example uses weighted load balancing. (The L2TP tunnel selection process is out of the scope of this example).

```
12 2019/05/22 10:57:38.331 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
Access-Accept(2) id 83 len 225 from 172.16.1.11:1812 vrid 1 pol rsp-radius-1
VSA [26] 15 Nokia(6527)
SUBSC PROF STR [12] 13 sub-profile-1
```

```
VSA [26] 15 Nokia(6527)
  SLA PROF STR [13] 13 sla-profile-1
  TUNNEL CLIENT ENDPOINT [66] 12 192.168.22.1
VSA [26] 6 Nokia(6527)
  TUNNEL ALGORITHM [47] 4 weighted access(1)
  TUNNEL CLIENT AUTH ID [90] 7 lac-pe1
  TUNNEL ASSIGNMENT ID [82] 12 RADIUS-group
VSA [26] 6 Nokia(6527)
  TUNNEL MAX RETRIES ESTAB [52] 4 0 2
  TUNNEL TYPE [64] 4 1 L2TP(3)
  TUNNEL MEDIUM TYPE [65] 4 1 IPv4(1)
  TUNNEL SERVER ENDPOINT [67] 13 1 192.168.22.2
  TUNNEL ASSIGNMENT ID [82] 16 1 RADIUS-tunnel-1
  TUNNEL PREFERENCE [83] 4 1 100
  TUNNEL TYPE [64] 4 2 L2TP(3)
  TUNNEL MEDIUM TYPE [65] 4 2 IPv4(1)
  TUNNEL SERVER ENDPOINT [67] 13 2 192.168.22.3
  TUNNEL ASSIGNMENT ID [82] 16 2 RADIUS-tunnel-2
  TUNNEL PREFERENCE [83] 4 2 80
"
```

The **debug router l2tp event lac-session-setup** command shows the LAC tunnel LNS2-T2 is selected for this example.

```
13 2019/05/22 10:57:38.332 CEST MINOR: DEBUG #2001 Base PPPoE 38->L2TP
"PPPoE 38->L2TP: UDP 192.168.22.1:1701 -> 192.168.22.3:1701
preference 80 tunnel RADIUS-group:RADIUS-tunnel-2
request to open new tunnel 8880"
```

```
14 2019/05/22 10:57:38.332 CEST MINOR: DEBUG #2001 Base PPPoE 38->L2TP
"PPPoE 38->L2TP: UDP 192.168.22.1:1701 -> 192.168.22.3:1701
preference 80 tunnel RADIUS-group:RADIUS-tunnel-2
create session 581989428"
```

The operational PPPoE session information in IES 1/base instance is shown as follows.

```
*A:LAC# show service id 1 ppp session

=====
PPP sessions for service 1
=====
User-Name
  Descr.
      Up Time      Type  Termination      IP/L2TP-Id/Interface-Id MC-StdbY
-----
user1@wholesale.com
  svc:1 sap:1/1/3:100 mac:00:00:00:00:01:01 sid:1
      0d 00:01:19   oE   lac      1066935934
-----
No. of PPP sessions: 1
=====
*A:LAC#
```

The operational L2TP tunnel information (base instance) is shown below.

```
*A:LAC# show router l2tp tunnel

=====
Conn ID   Loc-Tu-ID Rem-Tu-ID State      Blacklist-state   Ses Active
Group                                           Ses Total
Assignment
-----
```

```
581959680 8880 11367 established not-blacklisted 1
RADIUS-group 1
RADIUS-tunnel-2
```

```
-----
No. of tunnels: 1
=====
```

```
*A:LAC#
```

Operational session information (base instance) shows the session is in the established state.

```
*A:LAC# show router l2tp session
```

```
=====
L2TP Session Summary
=====
```

ID	Control Conn ID	Tunnel-ID	Session-ID	State
581989428	581959680	8880	29748	established

```
-----
No. of sessions: 1
=====
```

```
*A:LAC#
```

The L2TP endpoint/peer information shows there are two tunnels for tunnel endpoint 192.168.22.2.

```
*A:LAC# show router l2tp peer
```

```
=====
L2TP Peers
=====
```

Peer IP	Drain	Reachability	Port	Tun Active	Ses Active
				Tun Total	Ses Total
192.168.22.2			1701	0	0
				0	0
192.168.22.3			1701	1	1
				1	1

```
-----
No. of peers: 2
=====
```

```
*A:LAC#
```

The following command gives a system overview of subscriber session related data. This system overview shows the current and peak values per session type (local PTA, LAC, LTS, LNS) and an overview of the number of originated or terminated L2TP tunnels. Peak values can be cleared via the **clear subscriber-mgmt peakvalue-stats** command.

```
*A:LAC# show subscriber-mgmt statistics session system
```

```
=====
Subscriber Management Statistics for System
=====
```

Type	Current	Peak	Peak Timestamp

PPP Session Statistics			

Local PPP Sessions - PPPoE	0	0	
Local PPP Sessions - PPPoEoA	0	0	
Local PPP Sessions - PPPoA	0	0	
Local PPP Sessions - L2TP (LNS)	0	0	

LAC	PPP Sessions	-	PPPoE	1	2	05/22/2019	10:13:30
	PPP Sessions	-	PPPoEoA	0	0		
	PPP Sessions	-	PPPoA	0	0		
	PPP Sessions	-	L2TP (LTS)	0	0		
Total	PPP Sessions	-	established	1	2	05/22/2019	10:13:30
	PPP Sessions	-	in setup	0	1	05/22/2019	10:57:38
	PPP Sessions	-	local	0	0		
	PPP Sessions	-	LAC	1	2	05/22/2019	10:13:30
L2TP	L2TP Tunnels	-	originator	1	3	05/21/2019	14:49:38
	L2TP Tunnels	-	receiver	0	0		
	Total L2TP Tunnels			1	3	05/21/2019	14:49:38

IPoE Session Statistics							

Total	IPoE Sessions	-	established	0	0		
	IPoE Sessions	-	in setup	0	0		

=====							
Peak values last reset at : n/a							
*A:LAC#							

VRF hosted LAC

This example returns VPRN 65536 as the L2TP service instance [26-6527-104 Alc-Tunnel-Serv-Id]. The VPRN 65536 interface system address is used as the L2TP source address because the attribute Tunnel-Client-Endpoint is not returned.

The IP address 192.168.33.1 (Tunnel-Server-Endpoint) needs to be routable in VRF 65536 over a SAP or to a remote PE. This example uses BGP/MPLS IP virtual private networks (VPNs) (RFC 4364) to access the remote PE.

```
*A:LAC# show router 65536 route-table
```

Route Table (Service: 65536)				
Dest Prefix[Flags]	Type	Proto	Age	Pref
Next Hop[Interface Name]			Metric	

---snip---				
192.168.33.1/32	Local	Local	01d21h26m	0
system			0	
192.168.33.2/32	Remote	BGP VPN	20h03m54s	170
192.0.2.2 (tunneled)			10	
---snip---				
=====				
*A:LAC#				

Operational PPPoE session information for IES 1 (base instance) is shown using following command.

```
*A:LAC# show service id 1 ppp session
```

PPP sessions for service 1	
=====	

```
User-Name
Descr.
Up Time      Type  Termination      IP/L2TP-Id/Interface-Id MC-Stdby
-----
user1@wholesale.com
svc:1 sap:1/1/3:100 mac:00:00:00:00:01:01 sid:1
      0d 00:07:40  oE   lac          893192647
-----
No. of PPP sessions: 1
=====
*A:LAC#
```

Operational tunnel information for VPRN 65536 is displayed as follows.

```
*A:LAC# show router 65536 l2tp tunnel
=====
Conn ID      Loc-Tu-ID  Rem-Tu-ID  State          Blacklist-state  Ses Active
Group                               Ses Total
Assignment
-----
893190144    13629     14251     established    not-blacklisted  1
RADIUS-returned-TG                               1
RADIUS-returned-TN
-----
No. of tunnels: 1
=====
*A:LAC#
```

Operational session information for VPRN 65536 is displayed using following command, and shows that the session is established.

```
*A:LAC# show router 65536 l2tp session
=====
L2TP Session Summary
=====
ID            Control Conn ID  Tunnel-ID  Session-ID  State
-----
893192647     893190144       13629     2503       established
-----
No. of sessions: 1
=====
*A:LAC#
```

Scenario 2: Node-derived L2TP parameters

RADIUS returns L2TP group

This example returns VPRN 65536 as the L2TP service instance [26-6527-104] Alc-Tunnel-Serv-Id and an L2TP group name wholesale.com [26-6527-46] Alc-Tunnel-Group.

```
12 2019/05/22 11:15:46.604 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
Access-Accept(2) id 85 len 95 from 172.16.1.11:1812 vrid 1 pol rsp-radius-1
VSA [26] 15 Nokia(6527)
SUBSC PROF STR [12] 13 sub-profile-1
VSA [26] 15 Nokia(6527)
```

```
SLA PROF STR [13] 13 sla-profile-1
VSA [26] 6 Nokia(6527)
TUNNEL SERVICE ID [104] 4 65536
VSA [26] 15 Nokia(6527)
TUNNEL GROUP [46] 13 wholesale.com
"
```

For operational PPPoE session information in IES 1/base instance, use following command.

```
A:LAC# show service id 1 ppp session

=====
PPP sessions for service 1
=====
User-Name
Descr.
Up Time      Type  Termination      IP/L2TP-Id/Interface-Id MC-Stdby
-----
user1@wholesale.com
svc:1 sap:1/1/3:100 mac:00:00:00:00:01:01 sid:1
      0d 00:02:35  oE    lac          217909462
-----
No. of PPP sessions: 1
=====
*A:LAC#
```

Operational tunnel information for VPRN 65536 shows the tunnel is in the established state.

```
*A:LAC# show router 65536 l2tp tunnel

=====
Conn ID   Loc-Tu-ID Rem-Tu-ID State           Blacklist-state  Ses Active
Group                                           Ses Total
Assignment
-----
217907200 3325      374       established      not-blacklisted  1
wholesale.com
wholesale.com
-----
No. of tunnels: 1
=====
*A:LAC#
```

The operational session information for VPRN 65536 shows the session is in the established state.

```
*A:LAC# show router 65536 l2tp session

=====
L2TP Session Summary
=====
ID          Control Conn ID  Tunnel-ID  Session-ID  State
-----
217909462   217907200       3325       2262        established
-----
No. of sessions: 1
=====
*A:LAC#
```


LUDB returns L2TP group

This example returns VPRN 65536 as the L2TP service instance and L2TP group name wholesale.com (LUDB L2TP group "wholesale.com" service-id 65536).

The **debug subscriber-mgmt local-user-db l2tp detail all** command shows the LUDB authentication access (The returned parameter details are not shown).

```
11 2019/05/22 11:26:21.277 CEST MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
  user-name:
    original:  user1@wholesale.com
    masked:    user1@wholesale.com

Host wholesale.com found in user data base l2tp"
```

To show the operational data from LUDB */l2tp*, use the following command.

```
*A:LAC# show subscriber-mgmt local-user-db "l2tp" ppp-host "wholesale.com"
| match N/A invert-match
| match none invert-match

=====
PPP Host "wholesale.com"
=====
Admin State           : Up
Last Mgmt Change      : 05/20/2019 13:45:52

Host Identification
  User Name           : wholesale.com (domain only)

Matched Objects       : userName

Password Type         : ignore
PAD0 Delay            : 0msec
Diameter app policy   : (Not Specified)
Diameter auth policy  : (Not Specified)
Force IPv6CP          : Disabled
Ignore DF Bit         : Disabled

DHCPv6 lease times
  Renew timer         : > 9999 days
  Rebind timer        : > 9999 days
  Preferred lifetime  : 0d 00:00:00
  Valid lifetime      : 0d 00:00:00

Identification Strings (option 254)
  Subscriber Id       : user@wholesale.com
  SLA Profile String   : sla-profile-1
  Sub Profile String   : sub-profile-1

L2TP
  Service             : 65536
  Tunnel Group        : wholesale.com

MSAP defaults

Filter Overrides

Access loop info
=====
```

```
*A:LAC#
```

The **debug router l2tp event lac-session-setup** command shows the LAC tunnel selected for this example.

```
12 2019/05/22 11:26:21.278 CEST MINOR: DEBUG #2001 vprn65536 PPPoE 41->L2TP
"PPPoE 41->L2TP: UDP 192.168.33.1:1701 -> 192.168.33.2:1701
preference 50 tunnel wholesale.com:wholesale.com
request to open new tunnel 11734"
```

```
13 2019/05/22 11:26:21.278 CEST MINOR: DEBUG #2001 vprn65536 PPPoE 41->L2TP
"PPPoE 41->L2TP: UDP 192.168.33.1:1701 -> 192.168.33.2:1701
preference 50 tunnel wholesale.com:wholesale.com
create session 769017741"
```

For the operational PPPoE session information in IES 1/base instance, use the following command.

```
*A:LAC# show service id 1 ppp session

=====
PPP sessions for service 1
=====
User-Name
  Descr.
      Up Time      Type  Termination      IP/L2TP-Id/Interface-Id MC-Stdb
-----
user1@wholesale.com
  svc:1 sap:1/1/3:100 mac:00:00:00:00:01:01 sid:1
      0d 00:03:12   oE    lac          769017741
-----
No. of PPP sessions: 1
=====
*A:LAC#
```

Operational tunnel information for VPRN 65536 can be obtained using following command.

```
*A:LAC# show router 65536 l2tp tunnel

=====
Conn ID   Loc-Tu-ID Rem-Tu-ID State           Blacklist-state  Ses Active
Group                                           Ses Total
Assignment
-----
768999424 11734     4184      established    not-blacklisted  1
wholesale.com
wholesale.com
-----
No. of tunnels: 1
=====
*A:LAC#
```

The operational session information for VPRN 65536 shows the session is in the established state.

```
*A:LAC# show router 65536 l2tp session

=====
L2TP Session Summary
=====
ID          Control Conn ID  Tunnel-ID  Session-ID  State
-----
769017741   768999424       11734     18317       established
```

```
-----  
No. of sessions: 1  
=====
```

*A:LAC#

Advanced topics

Non-session-triggered L2TP tunnel setup

In addition to the PPP-session-triggered setup, an L2TP tunnel can also be set up via a **tools** command or an **auto-establish** command.

These non-session triggers are useful, for example, during the initial configuration phase where the LAC-LNS tunnel setup can be tested without the need for a user to attempt and establish a PPPoE connection.

The PPPoE user still triggers the L2TP session-setup over this L2TP tunnel and RADIUS needs to return an L2TP group name with the relevant name during authentication.

Auto-establish

Every minute, a check is performed to determine if tunnels need to be established (a process referred to as scan auto-establish). The tunnel state is establishedIdle when the tunnel is setup, and becomes established when user triggered sessions are set up over this tunnel.

```
configure  
  service  
    vprn 65536 customer 1 create  
      l2tp  
        group "wholesale.com" create  
          tunnel "wholesale.com" create  
            local-address 192.168.33.1  
            local-name "lac-pe1"  
            peer 192.168.33.2  
            auto-establish  
            no shutdown  
          exit  
        no shutdown  
      exit  
    no shutdown  
  exit  
exit  
exit  
exit  
exit
```

There is no difference in operational behavior for a tunnel set up via a session-trigger or an auto-establish command. Removing the auto-establish parameter has no impact on active tunnels (establishedIdle or established).

```
*A:LAC# show router 65536 l2tp tunnel  
=====
```

Conn ID	Loc-Tu-ID	Rem-Tu-ID	State	Blacklist-state	Ses Active
Group					Ses Total
Assignment					

```
-----
```

```
475136000 7250      6770      establishedIdle  not-blacklisted  0
wholesale.com      0
wholesale.com
-----
No. of tunnels: 1
=====
*A:LAC#
```

Tools tunnel start

First revert to the original situation, without auto-establish, as follows:

```
*A:LAC# configure service vprn 65536 l2tp group "wholesale.com"
                                tunnel "wholesale.com" no auto-establish
*A:LAC#
```

Verify the tunnel does not exist anymore, as follows:

```
*A:LAC# show router 65536 l2tp tunnel
No entries found.
*A:LAC#
```

Issue the tools command to manually establish the L2TP tunnel, as follows:

```
*A:LAC# tools perform router 65536 l2tp group "wholesale.com"
                                tunnel "wholesale.com" start
*A:LAC#
```

Verify a new tunnel has been created, as follows:

```
*A:LAC# show router 65536 l2tp tunnel
=====
Conn ID   Loc-Tu-ID Rem-Tu-ID State           Blacklist-state  Ses Active
Group                                           Ses Total
Assignment
-----
282394624 4309      4437      establishedIdle  not-blacklisted  0
wholesale.com      0
wholesale.com
-----
No. of tunnels: 1
=====
*A:LAC#
```

How long remains a tunnel idle before being torn down?

An L2TP tunnel can be torn down automatically, after the expiration of an idle timer, or manually through a tools command.

Idle timeout

A persistent tunnel is a tunnel that remains available after the last session over that tunnel is closed. To create a persistent tunnel, the **idle-timeout** parameter must be set to *infinite*.

A non-persistent tunnel is torn down immediately (idle-timeout zero) after the last session over that tunnel is closed or after a configurable delay. The idle-timeout parameter is set via the RADIUS [26-6527-49] Alc-Tunnel-Idle-Timeout attribute or the corresponding node parameter. The default value for this parameter is infinite (persistent).

```
configure router l2tp | configure service vprn l2tp
  idle-timeout [0..3600] s
  ---snip---
  group <tunnel-group-name>
    idle-timeout [0..3600] s | infinite
    ---snip---
    tunnel <tunnel-name>
      idle-timeout [0..3600] s | infinite
      ---snip---
```

The following shows an example of a non-persistent tunnel (idle-timeout 30 seconds). The tunnel changes state from established to establishedIdle when the last session is terminated. Idle-timeout seconds later, the session changes to the closed state. For the purpose of troubleshooting, the operational data stays available for destruct-timeout seconds (see later).

```
*A:LAC# show router l2tp tunnel detail

=====
L2TP Tunnel 921436160
=====

Connection ID: 921436160
Protocol      : v2
State       : closed
IP           : 192.0.2.1

---snip---

Name          : lac-pe1
Remote Name   : lns-pe2
Assignment ID : unnamed
Group Name    : default_radius_group
Acct. Policy  : N/A
Error Message: idle timeout (30 seconds) expired

Tunnel ID      : 14060
Preference     : 50
Hello Interval (s): 300
Idle T0 (s)   : 30
Max Retr Estab : 5
Cfg'd Sess Limit : unlimited
Transport Type : udpIp

Remote Conn ID : 280231936
Remote Tunnel ID : 4276
Receive Window  : 64
AVP Hiding      : never
Destruct T0 (s) : 60
Max Retr Not Estab: 5
Oper Session Limit: 32767
Challenge       : never

---snip---

No. of tunnels: 1
=====
*A:LAC#
```

The following shows an example of a persistent tunnel (idle-timeout infinite).

```
*A:LAC# show router l2tp tunnel detail

=====
L2TP Tunnel 235405312
=====
```

```

=====
Connection ID: 235405312
Protocol      : v2
State        : establishedIdle
IP           : 192.0.2.1

---snip---

Name          : lac-pe1
Remote Name   : lns-pe2
Assignment ID : unnamed
Group Name    : default_radius_group
Acct. Policy  : N/A
Error Message : N/A

Tunnel ID      : 3592
Preference     : 50
Hello Interval (s): 300
Idle T0 (s)    : infinite
Max Retr Estab : 5
Cfg'd Sess Limit : unlimited
Transport Type : udpIp

Remote Conn ID : 154861568
Remote Tunnel ID : 2363
Receive Window : 64
AVP Hiding     : never
Destruct T0 (s) : 60
Max Retr Not Estab: 5
Oper Session Limit: 32767
Challenge      : never

---snip---

No. of tunnels: 1
=====
*A:LAC#

```

Tools tunnel stop

In addition to the idle timeout used for tunnel termination, a tools stop command is also available that can be used to terminate persistent and non-persistent tunnels at any moment in time. Be aware that this command is very destructive and destroys all sessions carried over the closed tunnel.

Following command shows the tunnel is in the establishedIdle state.

```

*A:LAC# show router 65536 l2tp tunnel
=====
Conn ID   Loc-Tu-ID Rem-Tu-ID State           Blacklist-state  Ses Active
Group                                           Ses Total
Assignment
-----
60293120  920         7499     establishedIdle not-blacklisted  0
wholesale.com
wholesale.com
-----
No. of tunnels: 1
=====
*A:LAC#

```

The following command terminates the L2TP tunnel. The tunnel is aborted (the LAC sends StopCCN) using the <connection-id> or <tunnel-group-name>+<tunnel-name> as input. This StopCCN indicates "operator request" as the error reason.

```

*A:LAC# tools perform router 65536 l2tp group "wholesale.com"
          tunnel "wholesale.com" stop
INFO: CLI stopped 1 tunnels, destructed 0 tunnels.

```

```
*A:LAC#
```

The following debug output shows the tunnel being aborted.

```
13 2019/05/22 12:12:28.377 CEST MINOR: DEBUG #2001 vprn65536 L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.168.33.1:1701 -> 192.168.33.2:1701
tunnel 12828 session 0, ns 2 nr 1, flags:, reserved=0
  AVP MessageType(0,0), flags: mandatory, reserved=0
    StopControlConnectionNotification(4)
  AVP ResultCode(0,1), flags: mandatory, reserved=0
    result-code: "generalRequestToClearControlConnection"(1),
      error-code: "noGeneralError"(0)
    error-msg: "operator request"
  AVP AssignedTunnelId(0,9), flags: mandatory, reserved=0
    3695"
```

Alternatively, the tunnel can also be stopped with the following command. The effect would be the same.

```
*A:LAC# tools perform router 65536 l2tp tunnel 964886528 stop
```

Keepalive - L2TP hello

A keepalive mechanism is employed by L2TP in order to differentiate between tunnel outages and no control or data activity on a tunnel for an extended period. This is accomplished by injecting hello control messages after a specified period of time has elapsed since the last data or control message (ZLB not included) was received on a tunnel. As for any other L2TP control message, if the hello message is not reliably delivered, then the tunnel is declared down and reset, as defined in RFC 2661. This means that SR OS does not send hello packets if session control traffic is handled over this tunnel. The hello timer is reset if the system transmits any control packet over this tunnel (ZLB packets and data traffic are not taken into account).

The keepalive function is disabled (not recommended) using RADIUS [26-6527-50] Alc-Tunnel-Hello-Interval -1 or hello-interval infinite (default 300). The number of retries for unsuccessful hello packet delivery equals RADIUS [26-6527-52] Alc-Tunnel-Max-Retries-Estab or node parameter max-retries-estab (default 5). The retry interval is initially set to 1 second and doubles on each retry with a maximum interval of 8 seconds. Using a max-retries-estab 7 results in a retry of [1,2,4,8,8,8,8 seconds].

```
configure router l2tp | configure service vprn l2tp]
  hello-interval [60..3600] s | infinite # default 300 s
  max-retries-estab [2..7]                # default 5
  ---snip---
  group <tunnel-group-name>
    hello-interval [60..3600] s | infinite
    max-retries-estab [2..7]
    ---snip---
  tunnel <tunnel-name>
    hello-interval [60..3600] s | infinite
    max-retries-estab [2..7]
    ---snip---
```

For example, the LAC can be configured with an hello timer of 1 minute and the LNS with an hello timer of 2 minutes. The hello timer interval for LAC and LNS do not have to be same because the keepalive mechanism works asynchronous. See [Figure 119: L2TP keepalive mechanism](#).

```
*A:LAC# show router l2tp tunnel
```

Conn ID	Loc-Tu-ID	Rem-Tu-ID	State	Blacklist-state	Ses Active
Group	Assignment				Ses Total
794361856	11364	5391	established	not-blacklisted	1
default_radius_group					1
unnamed					
No. of tunnels: 1					

Figure 119: L2TP keepalive mechanism

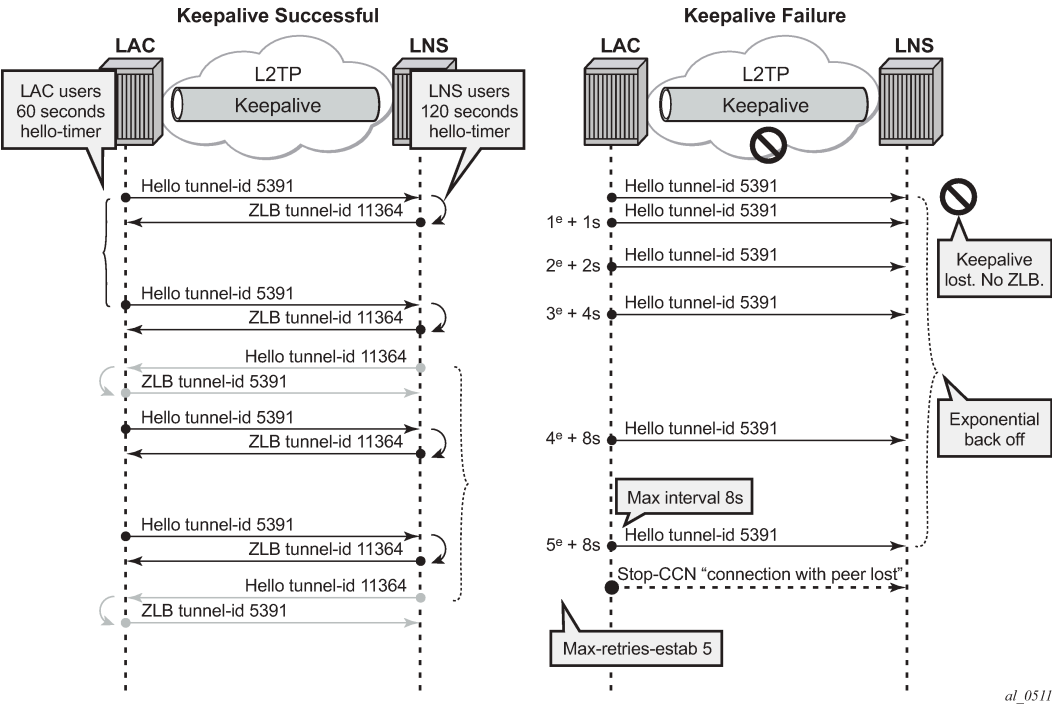


Figure 119: L2TP keepalive mechanism shows the tunnel being closed after 5 unsuccessful Hello deliveries with error-message *connection with peer lost*.

```
*A:LAC# show router 65536 l2tp tunnel detail

=====
L2TP Tunnel 479985664
=====

Connection ID: 479985664
---snip---

Acct. Policy : N/A
Error Message: connection with peer lost

Tunnel ID      : 7324
Preference     : 50
Hello Interval (s): 60
Remote Conn ID : 201195520
Remote Tunnel ID : 3070
Receive Window : 64
AVP Hiding     : never
```



```
---snip---
```

```
No. of tunnels: 1
```

```
=====
```

```
*A:LAC#
```

Keeping closed tunnel and session information

The destruct-timeout parameter (expressed in seconds) controls the period of time that the tunnel, or session data related to a closed (disconnected) tunnel, or session persists before being removed. The destruct timeout is a debugging aid by keeping underlying memory structures after the tunnel or session is terminated. It is configured via the RADIUS [26-6527-51] Alc-Tunnel-Destruct-Timeout attribute or the corresponding node parameter. Default value for this parameter is 60 seconds.

```
configure router l2tp | configure service vprn l2tp
  destruct-timeout [60..86400]
  ---snip---
  group <tunnel-group-name>
    destruct-timeout [60..86400]
    ---snip---
  tunnel <tunnel-name>
    destruct-timeout [60..86400]
```

The following output shows a session that is closed and the reason for it being terminated.

```
*A:LAC# show router l2tp session detail
```

```
=====
```

```
L2TP Session 900466242
```

```
=====
```

```
Connection ID: 900466242
```

```
State : closed
```

```
Tunnel Group : default_radius_group
```

```
Assignment ID: unnamed
```

```
Error Message: Terminated by PPPoE: Received PPPoE PADT
```

```
Control Conn ID : 900464640
```

```
Rem Cntrl Conn ID : 899416064
```

```
Tunnel ID : 13740
```

```
Remote Tunnel ID : 13724
```

```
Session ID : 1602
```

```
Remote Session ID : 23489
```

```
PW Type : ppp
```

```
Remote Conn ID : 899439553
```

```
Time Started : 05/22/2019 13:54:44
```

```
Time Established : 05/22/2019 13:54:44
```

```
Time Closed : 05/22/2019 13:54:50
```

```
CDN Result : generalError
```

```
General Error : vendorSpecific
```

```
-----
```

```
No. of sessions: 1
```

```
=====
```

```
*A:LAC#
```

The following output shows a tunnel that is closed and the reason for it being closed.

```
*A:LAC# show router l2tp tunnel detail
```

```
=====
```

```
L2TP Tunnel 900464640
```

```
=====
```

```
Connection ID: 900464640
```

```

Protocol      : v2
State         : closed
IP            : 192.0.2.1
UDP           : 1701
Peer IP       : 192.168.22.2
Peer UDP      : 1701
Tx dst-IP     : 192.168.22.2
Tx dst-UDP    : 1701
Rx src-IP     : 192.168.22.2
Rx src-UDP    : 1701
Name          : lac-pe1
Remote Name   : lns-pe2
Assignment ID : unnamed
Group Name    : default_radius_group
Acct. Policy  : N/A
Error Message: idle timeout (60 seconds) expired

---snip---

No. of tunnels: 1
=====
*A:LAC#

```

When the destruct timeout expires the tunnel and session is deleted, as follows:

```

*A:LAC# show router l2tp session detail
No entries found.
*A:LAC# show router l2tp tunnel detail
No entries found.
*A:LAC#

```

Floating peers

A floating peer exists if the peer LNS address indicated in the source address of the SCCRPs is different from the peer address known on the LAC. Floating peer allowance is configuration driven and is rejected by default.

The parameter `peer-address-change-policy` specifies whether the LAC accepts, ignores or rejects requests from a peer to change the destination IP address or UDP port.

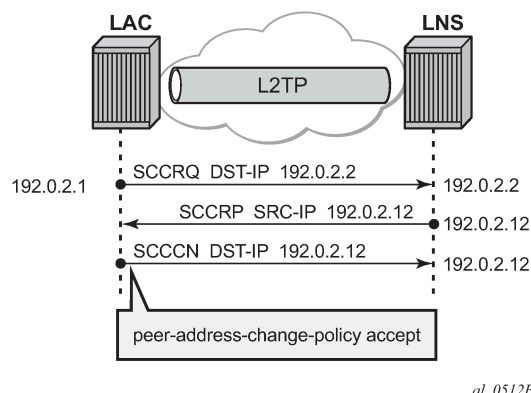
```

configure router l2tp | configure service vprn l2tp
  peer-address-change-policy accept | ignore | reject

```


- `accept` — Specifies that this system accepts any source IP address change for received L2TP control messages related to a locally originated tunnel in the state `wait-reply` and rejects any peer address change for other tunnels. In case the new peer IP address is accepted, it is learned and used as destination address in subsequent L2TP messages, as shown in [Figure 120: Floating peers accept](#).

Figure 120: Floating peers accept




- ignore — Specifies that this system ignores any source IP address change for received L2TP control messages, does not learn any new peer IP address and does not change the destination address in subsequent L2TP messages, as shown in [Figure 121: Floating peers ignore](#).

Figure 121: Floating peers ignore



- reject — Specifies that this system rejects any source IP address change for received L2TP control messages and drops those messages, as shown in [Figure 122: Floating peers reject](#).

Figure 122: Floating peers reject



The values Peer IP, Tx dst-IP and Rx src-IP in the **show router l2tp tunnel detail** command indicates if floating peers are used or not.

An example of a floating peer (peer-address-change-policy accept) is as follows.

```
*A:LAC# show router l2tp tunnel detail
=====
L2TP Tunnel Status
=====
Connection ID: 897122304
State       : established
IP          : 192.0.2.1
UDP         : 1701
Peer IP     : 192.0.2.2 # (1) peer address used in SCCRP
Peer UDP    : 1701
Tx dst-IP   : 192.0.2.12 # (3) peer address used in SCCRP
Tx dst-UDP  : 1701
Rx src-IP   : 192.0.2.12 # (2) SCCRP different IP received
Rx src-UDP  : 1701
---snip---
```

Tx/Rx connect speed - AVP 24/38

The connect speed (TX AVP 24 and RX AVP 38) is passed in the ICCN messages sent from the LAC to the LNS. The L2TP AVP 24 defines the (Tx) connect speed in bps from the perspective of traffic flowing from the LAC towards the subscriber (BNG downstream rate). The L2TP AVP 38 defines the (Rx) connect speed in bps from the perspective of traffic flowing from the subscriber toward the LAC (BNG upstream rate).

The report-rate configuration option indicates what rate is reported to the LNS when creating an L2TP session.

```
configure subscriber-mgmt sla-profile <sla-profile-name> ingress | egress
  report-rate agg-rate-limit|scheduler|pppoe-actual-rate|policer|rfc5515-actual-rate
```

- **agg-rate-limit** — Take the aggregate rate as received from the RADIUS Access-Accept message in VSA Alc-Subscriber-QoS-Override. When this RADIUS VSA is not present in the Access-Accept, or when RADIUS is not used, then take the configured aggregate rate limit. In the case where this is not configured, then take the port rate.

- **scheduler <scheduler-name>** — Take the rate of the specified scheduler. In case the scheduler is not linked with the scheduler-policy from the subscriber-profile, then take the port rate.
- **pppoe-actual-rate** — Take the rate from the DSL-Forum Vendor-Specific PPPoE Tag when available, otherwise take the port rate.
- **rfc5515-actual-rate** — Put the same value as the transmitted Actual-Data-Rate-Upstream AVP in the Rx-Connect-Speed AVP, and the same value as the transmitted Actual-Data-Rate-Downstream AVP in the Tx-Connect-Speed AVP.

Calling number AVP 22 format

The format of AVP 22 calling number in the ICRQ message is configurable via the parameter **calling-number-format**. The default format is "%S<space>%s" and corresponds to the concatenation of system-name<space>sap-id. Available parameters are %S (system-name), %c (Agent Circuit Id), %r Agent Remote Id, %s (sap-id), %l (Logical Line ID) and fixed strings. A combination can be configured from any of these parameters, but the total configured format cannot exceed 255 characters.

Example 1: Default configuration.

```
configure router l2tp calling-number-format "%S %s"
```

```
19 2019/05/22 14:01:04.885 CEST MINOR: DEBUG #2001 Base L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.0.2.1:1701 -> 192.168.22.2:1701
tunnel 5593 session 0, ns 2 nr 1, flags:, reserved=0
  AVP MessageType(0,0), flags: mandatory, reserved=0
    IncomingCallRequest(10)
  AVP AssignedSessionId(0,14), flags: mandatory, reserved=0
    15342
  AVP CallSerialNumber(0,15), flags: mandatory, reserved=0
    84781
  AVP CallingNumber(0,22), flags: mandatory, reserved=0
    "LAC 1/1/3:100"
  AVP AgentCircuitId(3561,1), flags:, reserved=0
    "circuit0"
  AVP AgentRemoteId(3561,2), flags:, reserved=0
    "remote0"
  AVP ActDataRateUp(3561,129), flags:, reserved=0
    2000000
  AVP ActDataRateDown(3561,130), flags:, reserved=0
    4000000"
```

Example 2: Customized configuration and all parameters (%S %s %c) are available to construct the requested AVP 22.

```
configure router l2tp calling-number-format "start-%S###%s###%c-end"
```

```
19 2019/05/22 14:05:28.116 CEST MINOR: DEBUG #2001 Base L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.0.2.1:1701 -> 192.168.22.2:1701
tunnel 832 session 0, ns 2 nr 1, flags:, reserved=0
  AVP MessageType(0,0), flags: mandatory, reserved=0
    IncomingCallRequest(10)
  AVP AssignedSessionId(0,14), flags: mandatory, reserved=0
    24748
  AVP CallSerialNumber(0,15), flags: mandatory, reserved=0
    84782
  AVP CallingNumber(0,22), flags: mandatory, reserved=0
```

```
"start-LAC###1/1/3:100###circuit0-end"
AVP AgentCircuitId(3561,1), flags:, reserved=0
"circuit0"
AVP AgentRemoteId(3561,2), flags:, reserved=0
"remote0"
AVP ActDataRateUp(3561,129), flags:, reserved=0
2000000
AVP ActDataRateDown(3561,130), flags:, reserved=0
4000000"
```

Example 3: Customized configuration and not all parameters are available to construct the requested AVP 22. Option-82 circuit-id (%c),remote-id (%r), and LLID (%l) information are lacking and therefore missing (skipped) in the formatted attribute.

```
configure router l2tp calling-number-format "%S#%c#%r#%l#%s"
```

```
19 2019/05/22 14:07:26.302 CEST MINOR: DEBUG #2001 Base L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.0.2.1:1701 -> 192.168.22.2:1701
tunnel 255 session 0, ns 2 nr 1, flags:, reserved=0
  AVP MessageType(0,0), flags: mandatory, reserved=0
    IncomingCallRequest(10)
  AVP AssignedSessionId(0,14), flags: mandatory, reserved=0
    20814
  AVP CallSerialNumber(0,15), flags: mandatory, reserved=0
    84783
  AVP CallingNumber(0,22), flags: mandatory, reserved=0
    "LAC#circuit0#remote0##1/1/3:100"
  AVP AgentCircuitId(3561,1), flags:, reserved=0
    "circuit0"
  AVP AgentRemoteId(3561,2), flags:, reserved=0
    "remote0"
  AVP ActDataRateUp(3561,129), flags:, reserved=0
    2000000
```

Prevent LAC from transmitting calling number AVP 22 to LNS

By default, the LAC includes the calling number AVP 22 in the L2TP incoming-call-request (ICRQ) packets transmitted to LNS. This AVP identifies the interface that is connected to the customer in the access network. Network access interface information can be hidden by configuring the LAC not to send the calling number AVP to the LNS.

Use the following command to disable the sending of L2TP calling number AVP 22.

```
configure router l2tp
  exclude-avps calling-number
```

AVP 100 - Cisco-Nas-Port

Interoperation with a Cisco LNS requires that the LAC communicates a NAS port type to the LNS via the L2TP ICRQ 'Cisco Nas Port Info AVP (100)'. This AVP (100) includes information that identifies the NAS port and indicates whether the port type is Ethernet or ATM and is configured via the cisco-nas-port parameter.

The Cisco AVP 100 format is as follows:

- First 5 bytes are NAS-Port-Type:
 - 0f10090203 (Ethernet)
 - 0f10090201 (ATM)
- Remaining four bytes correspond with the configured cisco-nas-port value

Example:

- Ethernet 12b s-vlan-id; 10b c-vlan-id; 3b slot number; 2b MDA nbr; 5b port
- ATM 12b VPI; 10b VCI; 3b slot number; 2b MDA nbr; 5b port

```
configure router l2tp
    cisco-nas-port ethernet "*12o*10i*3s*2m*5p" atm "*12v*10c*3s*2m*5p"
```

nas-port 1/1/3:100 corresponds to 102563 (000000000000 0001100100 001 01 00011).

```
19 2019/05/22 14:11:49.431 CEST MINOR: DEBUG #2001 Base L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.0.2.1:1701 -> 192.168.22.2:1701
tunnel 13002 session 0, ns 2 nr 1, flags:, reserved=0
  AVP MessageType(0,0), flags: mandatory, reserved=0
    IncomingCallRequest(10)
  AVP CiscoNasPort(9,100), flags:, reserved=0
    102563 type=ethernet(0f:10:09:02:03)
  AVP AssignedSessionId(0,14), flags: mandatory, reserved=0
    2461
  AVP CallSerialNumber(0,15), flags: mandatory, reserved=0
    84784
  AVP CallingNumber(0,22), flags: mandatory, reserved=0
    "LAC 1/1/3:100"
  AVP AgentCircuitId(3561,1), flags:, reserved=0
    "circuit0"
  AVP AgentRemoteId(3561,2), flags:, reserved=0
    "remote0"
  AVP ActDataRateUp(3561,129), flags:, reserved=0
    2000000
  AVP ActDataRateDown(3561,130), flags:, reserved=0
    4000000"
```

L2TP group/peer/tunnel draining

When the LAC has established sessions, the LAC can avoid the creation of new sessions for a specific group, peer, or tunnel, via the **drain** command.

No new sessions are created for a group, peer or tunnel that is being drained (draining state) but the current sessions are left intact.

After the **drain** command is issued, the group, peer, or tunnel moves from a draining to drained state when the last session is closed. A drained group, peer, or tunnel can then be managed (reconfigured, deleted) without any user impact.

Be aware that a group, peer, or tunnel in a draining or drained state is skipped in the tunnel selection process. The next example shows a tunnel draining; group and peer draining works according in the same way.

A tunnel has 1 session and is in established state.

```
*A:LAC# show router 65536 l2tp tunnel
=====
```

```

Conn ID   Loc-Tu-ID Rem-Tu-ID State           Blacklist-state  Ses Active
Group                                           Ses Total
Assignment
-----
1023868928 15623      9957      established      not-blacklisted  1
wholesale.com                                     1
wholesale.com
-----
No. of tunnels: 1
=====
*A:LAC#

```

The following tools **drain** command puts the tunnel in a draining state and leaves the sessions intact.

```

*A:LAC# tools perform router 65536 l2tp tunnel 1023868928 drain
*A:LAC#

```

Initially the tunnel is in the draining state.

```

*A:LAC# show router 65536 l2tp tunnel
=====
Conn ID   Loc-Tu-ID Rem-Tu-ID State           Blacklist-state  Ses Active
Group                                           Ses Total
Assignment
-----
1023868928 15623      9957      draining         not-blacklisted  1
wholesale.com                                     1
wholesale.com
-----
No. of tunnels: 1
=====
*A:LAC#

```

The tunnel moves to the drained state at the moment the last session is closed. Debugging shows that a drained tunnel is also not used as last resort and is skipped during the tunnel selection process.

```

*A:LAC# show router 65536 l2tp tunnel
=====
Conn ID   Loc-Tu-ID Rem-Tu-ID State           Blacklist-state  Ses Active
Group                                           Ses Total
Assignment
-----
1023868928 15623      9957      drained          not-blacklisted  0
wholesale.com                                     1
wholesale.com
-----
No. of tunnels: 1
=====
*A:LAC#

```

The following output shows new sessions cannot select a drained tunnel.

```

82289 2019/05/22 14:43:56.549 CEST MINOR: DEBUG #2001 vprn65536 PPPoE 19644->L2TP
"PPPoE 19644->L2TP: UDP 192.168.33.1:1701 -> 192.168.33.2:1701
preference 50 tunnel wholesale.com:wholesale.com
no additional session can be created in tunnel 15623"

```

```

82290 2019/05/22 14:43:56.549 CEST MINOR: DEBUG #2001 vprn65536 PPPoE 19644->L2TP
"PPPoE 19644->L2TP:

```



```
stop: no more tunnels can be tried"
```

The drained tunnel can then be closed without user impact.

```
*A:LAC# tools perform router 65536 l2tp tunnel 1023868928 stop
*A:LAC#
```

```
207376 2019/05/22 14:45:01.981 CEST MINOR: DEBUG #2001 vprn65536 L2TP(v2, ctrl,
                                                    ingress)
"L2TP(v2, ctrl, ingress): UDP 192.168.33.2:1701 -> 192.168.33.1:1701
tunnel 15623 session 0, ns 2 nr 8, flags:, reserved=0"

207377 2019/05/22 14:46:01.853 CEST MINOR: DEBUG #2001 vprn65536 L2TP(v2, ctrl,
                                                    ingress)
"L2TP(v2, ctrl, ingress): UDP 192.168.33.2:1701 -> 192.168.33.1:1701
tunnel 15623 session 0, ns 2 nr 8, flags:, reserved=0
  AVP MessageType(0,0), flags: mandatory, reserved=0
    StopControlConnectionNotification(4)
  AVP ResultCode(0,1), flags: mandatory, reserved=0
    result-code: "generalRequestToClearControlConnection"(1),
                error-code: "noGeneralError"(0)
    error-msg: "idle timeout (60 seconds) expired"
  AVP AssignedTunnelId(0,9), flags: mandatory, reserved=0
    9957"
```

For draining and undraining for example a group, following commands can be used.

```
tools perform router 65536 l2tp group "wholesale.com" drain
tools perform router 65536 l2tp group "wholesale.com" no drain
```

Conclusion

This example provides the LAC L2TP access server configuration and troubleshooting commands for the LAA architecture (tunneled-access) model.

Local User Database Basics

This chapter provides information about Local User Database (LUBD) Basics.

Topics in this chapter include:

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

Applicability

This chapter was initially written based on SR OS Release 13.0.R1, but the information and configuration in the current edition are based on SR OS Release 16.0.R4.


Overview

A local user database (LUBD) is a data source containing a set of host entries, providing full or partial enhanced subscriber management (ESM) data so that subscribers and subscriber hosts can be instantiated when end-users connect their devices.

An LUBD can be accessed for the following applications; see [Figure 123: LUBD applications](#).

1. To support ESM for retrieval of data to instantiate hosts and subscribers. This applies to the routed central office [CO] model only.
2. To support a local DHCPv4 server; for example for assigning fixed IP addresses to dedicated end-user devices.
3. To allow the system to provide the ESM data in case the RADIUS server referenced from the authentication policy is not available. The LUBD serves as a fallback for RADIUS authentication.

Figure 123: LUDB applications



al_0806

The LUDB lookup process is common to the applications shown in [Figure 123: LUDB applications](#), and performs the following steps:

- Applying match criteria, to select the input parameters that will be used for the lookup.
- Optionally, applying a mask to one or more of the remaining input parameters.
- Performing the lookup.

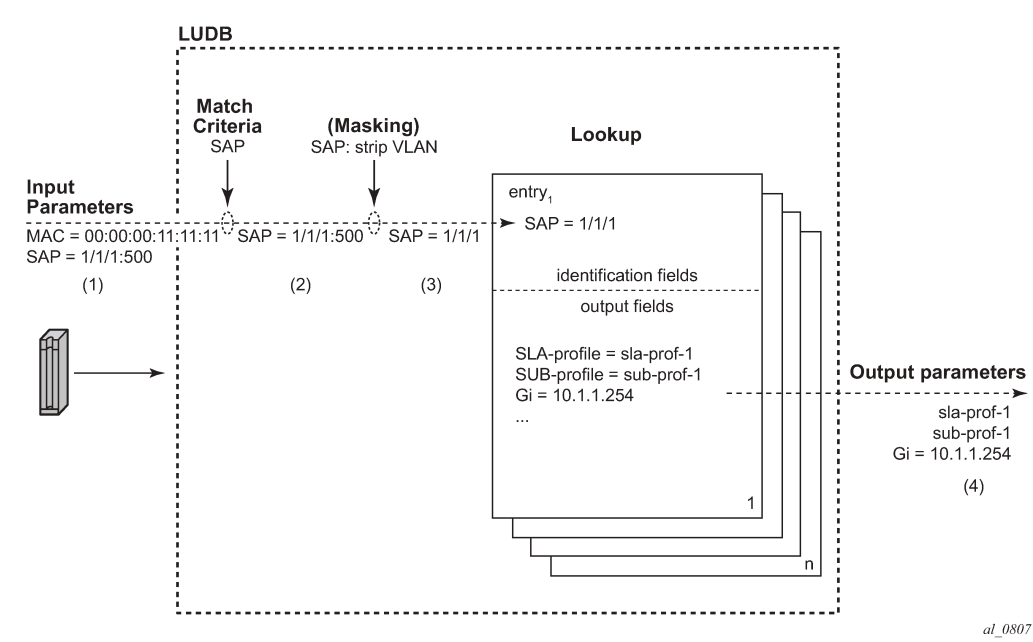
The LUDB lookup process translates a set of input parameters (the host identification fields) into a set of output parameters; see [Figure 124: Processing an LUDB lookup request](#) for the following example:

- An LUDB lookup is requested for a client with MAC address and SAP as input parameters (1).
- The match criteria indicate to consider the SAP only, so the MAC address is ignored (2).
- The masking defines the stripping of the VLAN-tag from the SAP (3).
- The lookup then uses SAP 1/1/1 and finds entry1 to be the matching entry, so the LUDB returns the SLA-profile string and the SUB-profile string together with the Gi address.

Optionally, an LUDB defines a **default** host entry, which is used in case none of the other entries matches the lookup request, so it serves as a wildcard (*).

Not finding any matching host entry in an LUDB results in a setup failure.

Figure 124: Processing an LUDB lookup request



Configuration

Creating LUDBs

An LUDB is identified by a name of 32 characters maximum ([Figure 125: Creating LUDBs and LUDB entries](#)).

```
*A:BNG-1>config>subscr-mgmt# local-user-db
- local-user-db <local-user-db-name> [create]
```

Multiple LUDBs can be defined, and their respective names must be unique.

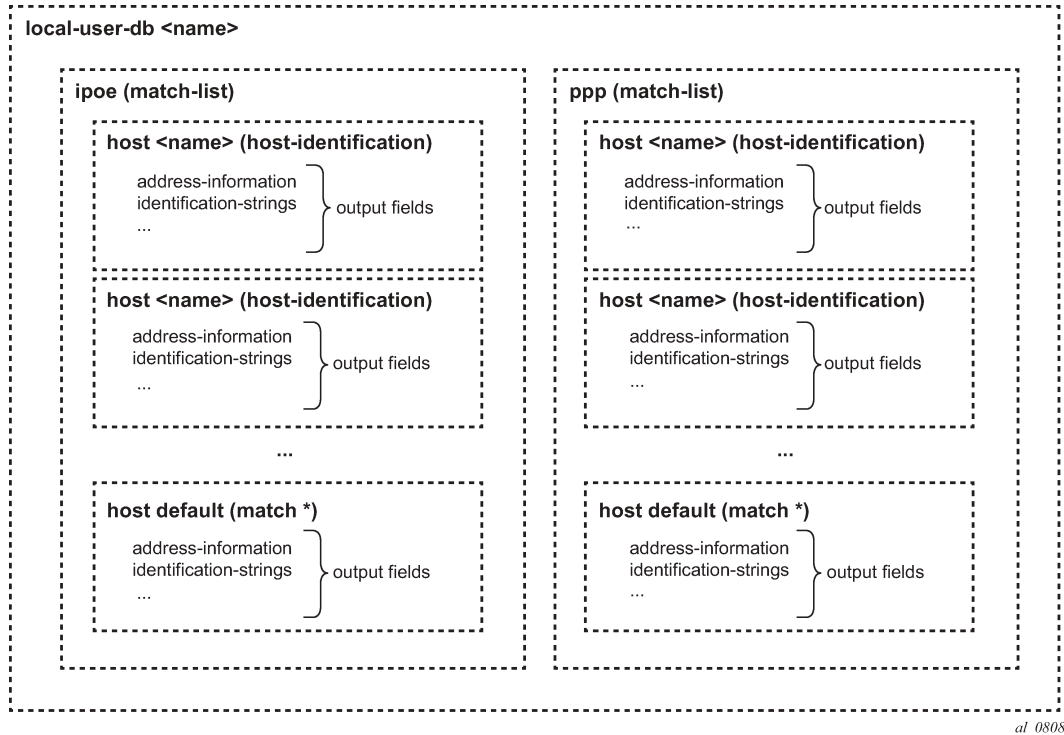
An LUDB can provide the data for IPoE (DHCPv4, DHCPv6, and SLAAC) as well as for PPP users.

```
*A:BNG-1>config>subscr-mgmt>loc-user-db$
[no] description      - Description for this local user database
    ipoe              + Configure IPoE hosts
    ppp               + Configure PPP hosts
[no] shutdown         - Administratively enable/disable this local user database
```

For an LUDB to be active, the LUDB must be in the **no shutdown** state.

Individual host entries in an LUDB can match single or multiple hosts.

Figure 125: Creating LUDBs and LUDB entries



Creating host entries

A host entry is identified by name of 32 characters maximum ([Figure 125: Creating LUDBs and LUDB entries](#)).

```
*A:BNB-1>config>subscr-mgmt>loc-user-db>ipoe$ host
- host <host-name> [create]
```

The name **default** can optionally be used as a wildcard for situations where otherwise the lookup fails to find a matching entry. A host entry belongs to either the IPoE or the PPP section of an LUDB. The name of a host entry must be unique within the section. A host entry contains two sets of fields. The first set of fields are the host-identification fields and are used for the lookup, the second set of fields are output to the lookup process.

The host-identification fields available for IPoE are, in alphabetical order:

- circuit-id
- derived-id, which must be defined using a Python script, which derives the value from DHCP messages
- encap-tag-range
- ip-prefix
- mac-address
- option 60
- remote-id

- sap-id
- service-id
- string
- system-id

The host-identification fields available for PPP are, in alphabetical order:

- circuit-id: taken from the PPPoE tags
- derived-id, which must be defined using a Python script, which derives the value from PPP messages
- encap-tag-range
- mac
- remote-id: taken from the PPPoE tags
- sap-id
- service-name
- username

The output fields of the lookup process include the identification strings, DHCP options, IP address information, MSAP information, and so on.

Entry validation

For a host entry to be active, it must be put in the **no shutdown** state.

Before adding the host entry to the lookup database, the system validates the host entry:

- A **default** host entry can be added, preferably without host identification fields.

```
*A:BNG-1>config>subscr-mgmt>loc-user-db>ipoe# host default create
*A:BNG-1>config>subscr-mgmt>loc-user-db>ipoe>host$ no shutdown
INFO: DHCP# 1138 This host will be considered as the default host
*A:BNG-1>config>subscr-mgmt>loc-user-db>ipoe>host$
```

Defining a **default** host entry with identification fields is not recommended, because it would turn the default host entry into a regular entry, instead of a match all entry, when the match-list is changed.

- A **non default** host entry without host identification fields cannot be added to the lookup database.

```
MINOR: DHCP# 1126 Host-identification must have at least 1 item defined
```

- A **non default** host entry with none of its host identification fields in common with the match-list is added to the unmatched host list.

```
INFO: DHCP# 1107 Host could not be inserted in lookup database - no
match values
```

- A **non default host** entry is added to the lookup database when at least one of the defined host identification fields is in common with the match-list, even when some of the host identification fields are not on the match-list.

- Two or more **non default** host entries with the same host-identification definitions are considered as duplicates. The second entry with the same host-identification definitions is not added to the lookup database; it is considered as a configuration mistake.

```
INFO: DHCPD #1107 Host could not be inserted in lookup database – duplicate
```

LUDB informational and error messages appear to be originating from the DHCPD application (DHCPD #nnn in the preceding outputs), even though the LUDB is not associated with a DHCPv4 server.

Creating a match list

Retrieving data from an LUDB requires one or more criteria to be put on a **match-list**. A match-list is a sequential list of parameters considered for the lookup; other parameters provided on LUDB access are ignored.

For IPoE, up to four criteria can be defined; for PPP, the maximum is three. The criteria on a match-list are processed in the order specified.

For IPoE users, the following match criteria are allowed, in alphabetical order:

- circuit-id
- derived-id (defined by a Python script)
- dual-stack-remote-id (IPv4 and IPv6, with IPv6 enterprise-id stripped off)
- encap-tag-range
- ip
- mac-address
- option 60
- remote-id (IPv4 and IPv6, including the IPv6 enterprise-id)
- sap-id
- service-id
- string
- system-id

```
*A:BNG-1>config>subscr-mgmt>loc-user-db>ipoe# match-list
- no match-list
- match-list <ipoe-match-type-1> [<ipoe-match-type-2>...(up to 4 max)]
<ipoe-match-type>      : circuit-id|derived-id|dual-stack-remote-id|
                        encap-tag-range|ip|mac|option60|remote-id|sap-id|
                        service-id|string|system-id

*A:BNG-1>config>subscr-mgmt>loc-user-db>ipoe#
```

For PPP users, the following match criteria are allowed, in alphabetical order:

- circuit-id
- derived-id (defined by a Python script)
- encap-tag-range
- mac-address

- remote-id (IPv4 and IPv6, including the IPv6 enterprise-id)
- sap-id
- service-name
- username (complete username, domain part only, host part only)

```
*A:BNG-1>config>subscr-mgmt>loc-user-db>ppp# match-list
- no match-list
- match-list <ppp-match-type-1> [<ppp-match-type-2>...(up to 3 max)]

<ppp-match-type>      : circuit-id|derived-id|mac|remote-id|sap-id|
                        encap-tag-range|service-name|username

*A:BNG-1>config>subscr-mgmt>loc-user-db>ppp#
```

Masking

Optionally, the parameters considered for the lookup can be masked.

Masking is prefix- or suffix- based, or a combination of both. A prefix or suffix string, or a prefix or suffix length, can be specified.

For PPP users, masks can be applied to the circuit-id, remote-id, sap-id, service-name, and username. For IPoE users, masks can be applied to the circuit-id, option 60, remote-id, sap-id, string, and system-id.

```
*A:BNG-1>config>subscr-mgmt>loc-user-db>ppp# mask ?
- mask type <ppp-match-type>
  {[prefix-string <prefix-string> | prefix-length <prefix-length>]
  [suffix-string <suffix-string> | suffix-length <suffix-length>]}
- no mask type <ppp-match-type>

<ppp-match-type>      : circuit-id|remote-id|sap-id|service-name|username
<prefix-string>       : [127 chars max] ('*' is wildcard)
<prefix-length>       : [1..127]
<suffix-string>       : [127 chars max] ('*' is wildcard)
<suffix-length>       : [1..127]

*A:BNG-1>config>subscr-mgmt>loc-user-db>ipoe# mask ?
- mask type <ipoe-match-type>
  {[prefix-string <prefix-string> | prefix-length <prefix-length>]
  [suffix-string <suffix-string> | suffix-length <suffix-length>]}
- no mask type <ipoe-match-type>

<ipoe-match-type>     : circuit-id|option60|remote-id|sap-id|string|system-id
<prefix-string>       : [127 chars max] ('*' is wildcard)
<prefix-length>       : [1..127]
<suffix-string>       : [127 chars max] ('*' is wildcard)
<suffix-length>       : [1..127]
```

The lookup occurs after applying the optional masks.

The examples in [Table 30: Masking examples](#) illustrate masking. For the third example, a combination of both prefix and suffix matching is used.

Table 30: Masking examples

Mask Type			prefix-length	suffix-length	prefix-string	suffix-string	result
username	circuit-id	remote-id					
-	-	87654321-BSAN-1	9	-	-	-	BSAN-1
-	BSAN-2 1 100 1/2/1	-	-	11	-	-	BSAN-2
all@domain-1.com	-	-	-	-	*@	.com	domain-1

Lookup

The following rules apply while scanning through an LUDB in search of a single matching entry:

1. Only criteria on the match-list are considered.

Assume a client with MAC-address, a circuit-id, and a remote-id. If the match-list only defines the MAC-address to be used as criterion, then the circuit-id and the remote-id are ignored. Only the MAC-address is used for selecting the proper host entry.

2. The order of the criteria on the match-list is important.

The match-list is a sequential list, and the criteria are processed left to right.

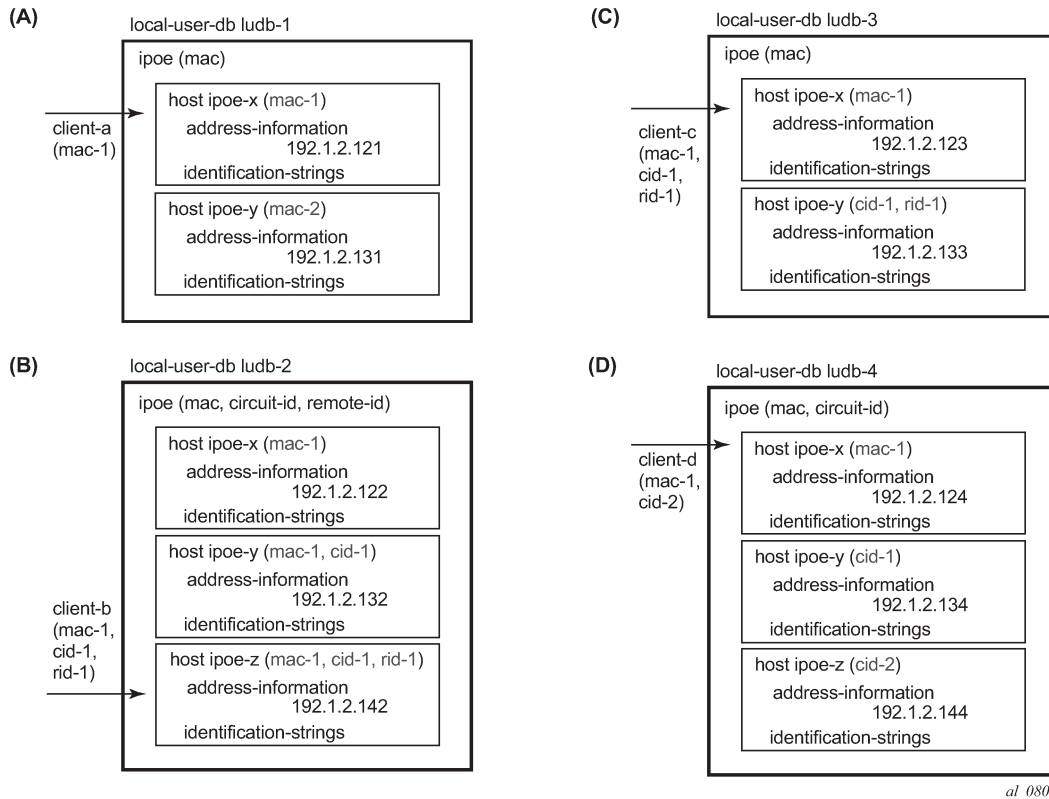
3. As many of the host-identification fields as possible must be matched, while still obeying rule 1.

Only the (optionally masked) parameters from the match list are verified.

4. A **default host** is excluded from the scan, if defined.

A **default host** is used as a fallback when scanning through an LUDB does not provide any result.

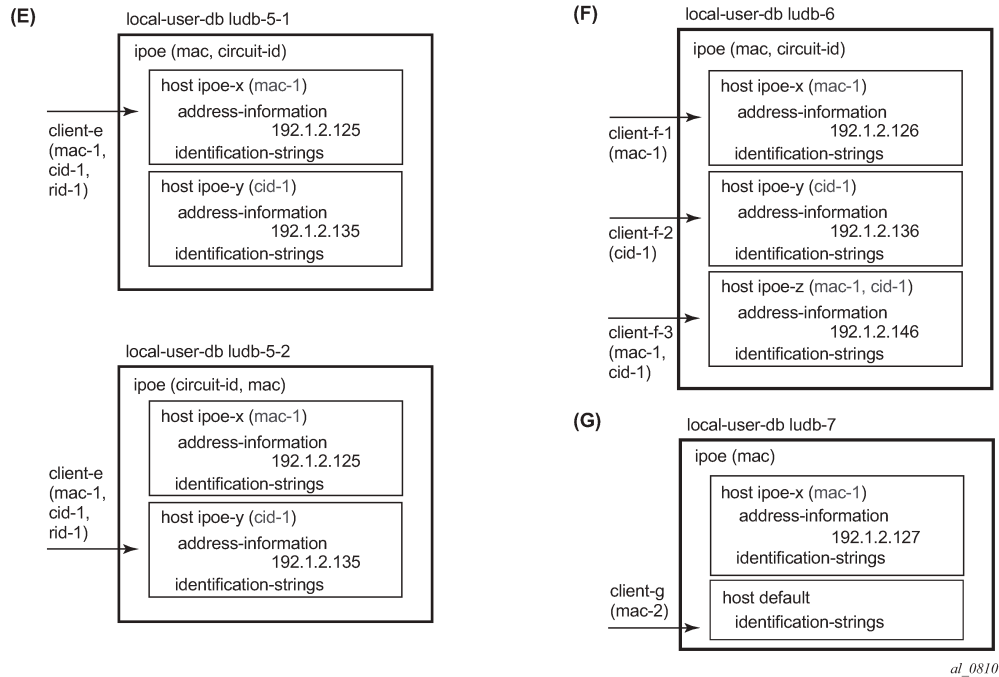
Figure 126: Host matching examples



The examples in [Figure 126: Host matching examples](#) and [Figure 127: Host matching examples \(continued\)](#) illustrate these rules:

1. Matching is based on the MAC-address only. When *client-a* with *mac-1* connects, host *ipoe-x* is matched.
2. Matching is based on the MAC-address, circuit-id, and remote-id, in this sequence. As client-b enters with *mac-1*, *cid-1* and *rid-1*, the match-list is scanned and matched left to right, so host *ipoe-z* is matched.
3. Matching is based on the MAC-address only. Even though client-c connects with *mac-1*, *cid-1*, and *rid-1*, the system ignores the circuit-id and the remote-id, so the matching host is *ipoe-x*. Note that host *ipoe-y* can never be matched using the match-list defined; it is on the unmatched host list.
4. Matching is based on the MAC-address and the circuit-id, in this sequence. Client-d connects with *mac-1* and *cid-2*, but because the system scans the match-list left through right, the MAC-address takes priority over the circuit-id. The matching host is *ipoe-x*.

Figure 127: Host matching examples (continued)



5. For the top part, matching is based on MAC-address and the circuit-id, in this sequence. When client-e connects (mac-1, cid-1, and rid-1), the system scans ludb-5-1 and matches host ipoe-x.

For the bottom part, matching is based on the circuit-id first, then the MAC-address. When client-e connects (mac-1, cid-1, and rid-1), the system scans ludb-5-2 and matches host ipoe-y.

6. Matching is based on MAC-address and the circuit-id, in this sequence. When client-f-1 (mac-1) connects, the matching host is ipoe-x because only the MAC-address is provided and checked. When client-f-2 (cid-1) connects, the matching host is ipoe-y because only the client-id is provided and checked. When client-f-3 (mac-1, cid-1) connects, the matching host is ipoe-z.
7. Matching is based on the MAC-address only. When client-g with mac-2 connects, host default is matched because there is no explicit entry matching mac-2.

As shown in these examples, the system only checks the parameters provided by the client in the sequence as defined by the match-list. Parameters not provided by a client will not be searched for.

Tools commands

The following **tools** command manually triggers the lookup of an IPoE host in an LUDB and is useful during commissioning, troubleshooting, and verification of the configured database, without the need for an external client.

```
*A:BNG-1# tools perform subscriber-mgmt local-user-db "ludb-1" ipoe host-lookup ?
- host-lookup [mac <ieee-address>] [remote-id <remote-id-ascii>]
  [sap-id <sap-id>] [service-id <service-id>] [string <vso-string>]
  [system-id <system-id>] [option60 <option-60-ascii>]
  [circuit-id <circuit-id-ascii>] [circuit-id-hex <circuit-id-hex>]
  [option60-hex <option60-hex>] [remote-id-hex <remote-id-hex>]
```

```
[derived-id <derived-id>] [ip-prefix <ip-prefix/ip-prefix-length>]

<ieee-address>      : xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx
<remote-id-ascii>    : [255 chars max]
<sap-id>             : [255 chars max]
<service-id>         : [1..2148278317]|<svc-name:64 char max>
<vso-string>         : [255 chars max]
<system-id>          : [255 chars max]
<option-60-ascii>    : [32 chars max]
<circuit-id-ascii>   : [127 chars max]
<circuit-id-hex>     : [0x0..0xFFFFFFFF...(max 254 hex nibbles)]
<option60-hex>       : [0x0..0xFFFFFFFF...(max 64 hex nibbles)]
<remote-id-hex>      : [0x0..0xFFFFFFFF...(max 510 hex nibbles)]
<derived-id>         : [255 chars max]
<ip-prefix/ip-pref*> : ipv4-prefix   - a.b.c.d (host bits must be 0)
                       ipv4-prefix-le - [0..32]
                       ipv6-prefix   - x:x:x:x:x:x:x      (eight 16-bit pieces)
                                   x:x:x:x:x:x:d.d.d.d
                                   x - [0..FFFF]H
                                   d - [0..255]D
                       ipv6-prefix-le - [0..128]
```

A similar command exists for the lookup of a PPP host in an LUDB.

```
*A:BNG-1# tools perform subscriber-mgmt local-user-db "ludb-1" ppp host-lookup ?
- host-lookup [circuit-id <circuit-id>] [circuit-id-hex <circuit-id-hex>]
  [derived-id <derived-id>] [mac <ieee-address>] [remote-id <remote-id>]
  [remote-id-hex <remote-id-hex>] [sap-id <sap-id>]
  [service-name <service-name>] [user-name <user-name>]

<circuit-id>          : [127 chars max]
<circuit-id-hex>      : [0x0..0xFFFFFFFF...(max 254 hex nibbles)]
<derived-id>          : [255 chars max]
<ieee-address>        : xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx
<remote-id>           : [255 chars max]
<remote-id-hex>       : [0x0..0xFFFFFFFF...(max 510 hex nibbles)]
<sap-id>              : [255 chars max]
<service-name>        : [255 chars max]
<user-name>           : [253 chars max]
```

Example 1: Single match criterion

The following shows an excerpt from ludb-1. Host entry-11 defines the parameters for an IPoE host, and host entry-55 defines the parameters for a PPPoE host. Host matching IPoE hosts is MAC-address based, whereas host matching PPP hosts is username based.

```
configure
  subscriber-mgmt
    local-user-db "ludb-1" create
      description "example user-db"
      ipoe
        match-list mac
        host "default" create
          address pool "pool4-1"
          no shutdown
        exit
        host "entry-11" create
          host-identification
            mac 00:00:00:11:11:11
          exit
          address 10.1.1.211
        ---snip---
```

```

        no shutdown
    exit
    ---snip---
exit
ppp
    match-list username
    host "entry-55" create
        host-identification
            username "sub55@domain1"
        exit
        password chap sub55
        address 10.1.2.252
    ---snip---
    no shutdown
    exit
    ---snip---
exit
no shutdown
exit

```

IPoE hosts

IPoE host entry lookup using a MAC-address only is triggered with following **tools** command.

```

*A:BNG-1# tools perform subscriber-mgmt local-user-db "ludb-1" ipoe
                                         host-lookup mac 00:00:00:11:11:11
=====
IPoE Host Lookup results
=====
Result                : Success
Matched Host Name     : entry-11
Admin State           : Up
Last Mgmt Change      : 11/20/2018 11:25:10

Host Identification
Circuit Id            : N/A
Mac Address           : 00:00:00:11:11:11
Remote Id             : N/A
Sap Id                : N/A
Service Id            : N/A
String                : N/A
Option 60             : N/A
System Id             : N/A
Encap Tag Range       : N/A
Derived Id            : N/A
IP prefix             : N/A

Matched Objects       : mac
---snip---
=====
*A:BNG-1#

```

The debug output confirms the successful lookup.

```

1 2018/11/20 11:40:41.132 CET MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
  mac: 00:00:00:11:11:11

  Host entry-11 found in user data base ludb-1"
*A:BNG-1#

```

The following command is using a MAC-address, a circuit-id, and a remote-id for the lookup. The output shows that only the MAC-address is used, the other input parameters are ignored (N/A) so again entry-11 is selected.

```
*A:BNG-1# tools perform subscriber-mgmt local-user-db "ludb-1" ipoe
                    host-lookup mac 00:00:00:11:11:11 circuit-id AA remote-id BB
=====
IPoE Host Lookup results
=====
Result                : Success
Matched Host Name     : entry-11
Admin State           : Up
Last Mgmt Change      : 11/20/2018 11:25:10

Host Identification
Circuit Id            : N/A
Mac Address           : 00:00:00:11:11:11
Remote Id             : N/A
Sap Id                : N/A
Service Id            : N/A
String                : N/A
Option 60             : N/A
System Id             : N/A
Encap Tag Range       : N/A
Derived Id            : N/A
IP prefix             : N/A

Matched Objects       : mac
---snip---
=====
*A:BNG-1#
```

The following command triggers the lookup of a non-existing MAC-address, leading to a host not found message.

```
*A:BNG-1# tools perform subscriber-mgmt local-user-db "ludb-1" ipoe
                    host-lookup mac 00:00:00:12:34:56
=====
IPoE Host Lookup results
=====
Result                : host not found
*A:BNG-1#
```

The host not found message is also confirmed by the debug output.

```
3 2018/11/20 11:43:30.351 CET MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host not found
 mac: 00:00:00:12:34:56

Host not found in user data base ludb-1"
```

To allow IPoE users with unknown MAC-addresses to successfully connect, a default host can be created, at which time an informational message is returned:

```
*A:BNG-1# configure subscriber-mgmt local-user-db "ludb-1" ipoe host "default" create
*A:BNG-1>config>subscr-mgmt>loc-user-db>ipoe>host# address pool pool4-1
*A:BNG-1>config>subscr-mgmt>loc-user-db>ipoe>host# no shutdown
INFO: DHCP #1138 This host will be considered as the default host
*A:BNG-1>config>subscr-mgmt>loc-user-db>ipoe>host# exit all
*A:BNG-1#
```

After the previous commands are executed, devices with MAC-addresses not explicitly listed in the LUDB can also connect.

```
*A:BNG-1# tools perform subscriber-mgmt local-user-db "ludb-1" ipoe
                                         host-lookup mac 00:00:00:12:34:56
=====
IPoE Host Lookup results
=====
Result           : Success
Matched Host Name : default
Admin State      : Up
Last Mgmt Change  : 11/20/2018 11:45:02

Host Identification
Circuit Id       : N/A
Mac Address      : N/A
Remote Id        : N/A
Sap Id           : N/A
Service Id       : N/A
String           : N/A
Option 60        : N/A
System Id        : N/A
Encap Tag Range  : N/A
Derived Id       : N/A
IP prefix        : N/A

Matched Objects   : N/A
---snip---
=====
*A:BNG-1#
```

PPP hosts

Manually authenticating a PPP host is done as follows.

```
*A:BNG-1# tools perform subscriber-mgmt local-user-db "ludb-1"
                                         ppp authentication user-name sub55@domain1 password sub55
=====
Authentication results
=====
Result           : Success
Matched Host Name : entry-55
Admin State      : Up
Last Mgmt Change  : 11/20/2018 11:25:10

Host Identification
Mac Address      : N/A
Circuit Id       : N/A
Remote Id        : N/A
Sap Id           : N/A
Service Name     : N/A
User Name        : sub55@domain1
Encap Tag Range  : N/A
Derived Id       : N/A

Matched Objects   : userName
---snip---
=====
*A:BNG-1#
```

When the wrong password is provided, the following message is returned:

```
*A:BNG-1# tools perform subscriber-mgmt local-user-db "ludb-1"
                        ppp authentication user-name sub55@domain1 password sub5x
=====
Authentication results
=====
Result                : invalid password
*A:BNG-1#
```

PPP host entry lookup is similar to the IPoE host lookup. The following example demonstrates a user-name based lookup.

```
*A:BNG-1# tools perform subscriber-mgmt local-user-db "ludb-1"
                        ppp host-lookup user-name sub55@domain1
=====
PPP host Lookup results
=====
Result                : Success
Matched Host Name     : entry-55
Admin State           : Up
Last Mgmt Change      : 11/20/2018 11:25:10

Host Identification
Mac Address           : N/A
Circuit Id            : N/A
Remote Id             : N/A
Sap Id                : N/A
Service Name          : N/A
User Name             : sub55@domain1
Encap Tag Range       : N/A
Derived Id            : N/A

Matched Objects       : userName
---snip---
=====
*A:BNG-1#
```

The following command is using a user-name and a MAC-address for the lookup.

```
*A:BNG-1# tools perform subscriber-mgmt local-user-db "ludb-1"
                        ppp host-lookup user-name sub55@domain1 mac 00:00:00:11:11:11
=====
PPP host Lookup results
=====
Result                : Success
Matched Host Name     : entry-55
Admin State           : Up
Last Mgmt Change      : 11/20/2018 11:25:10

Host Identification
Mac Address           : N/A
Circuit Id            : N/A
Remote Id             : N/A
Sap Id                : N/A
Service Name          : N/A
User Name             : sub55@domain1
Encap Tag Range       : N/A
Derived Id            : N/A

Matched Objects       : userName
---snip---
```



```
=====
*A:BNG-1#
```

Similar to the IPoE host lookup, the lookup of a non-existing user fails if no default entry is defined for PPP. In this case, a default host can be defined.

Example 2: Multiple match criteria

The following shows an excerpt from ludb-2, with multiple match criteria.

The match-list includes mac, circuit-id, and remote-id, in this sequence.

```
configure
  subscriber-mgmt
    local-user-db "ludb-2" create
      ipoe
        match-list mac circuit-id remote-id
        host "entry-11" create
          host-identification
            mac 00:00:00:11:11:11
          exit
          address 10.1.1.111
          ---snip---
          no shutdown
        exit
        host "entry-12" create
          host-identification
            circuit-id string "11"
            mac 00:00:00:11:11:11
          exit
          address 10.1.1.112
          ---snip---
          no shutdown
        exit
        host "entry-13" create
          host-identification
            circuit-id string "11"
            mac 00:00:00:11:11:11
            remote-id string "AA"
          exit
          address 10.1.1.113
          ---snip---
          no shutdown
        exit
        host "entry-14" create
          host-identification
            circuit-id string "11"
            remote-id string "AA"
          exit
          address 10.1.1.114
          ---snip---
          no shutdown
      exit
```

The following **tools** command uses a MAC-address only, with entry-11 being matched.

```
*A:BNG-1# tools perform subscriber-mgmt local-user-db "ludb-2"
                                     ipoe host-lookup mac 00:00:00:11:11:11
=====
IPoE Host Lookup results
=====
Result          : Success
Matched Host Name : entry-11
```

```
Admin State      : Up
Last Mgmt Change : 11/20/2018 11:25:10
```

```
Host Identification
Circuit Id      : N/A
Mac Address     : 00:00:00:11:11:11
Remote Id       : N/A
Sap Id          : N/A
Service Id      : N/A
String          : N/A
Option 60       : N/A
System Id       : N/A
Encap Tag Range : N/A
Derived Id      : N/A
IP prefix       : N/A
```

```
Matched Objects : mac
---snip---
```

```
=====
*A:BNG-1#
```

The corresponding debug output shows the parameters from the match-list and their values, in sequence. The values for the circuit-id and the remote-id are left empty as they were not provided for the lookup.

```
10 2018/11/20 11:52:30.777 CET MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
 mac: 00:00:00:11:11:11
 circuit-id:
 remote-id:
```

```
Host entry-11 found in user data base ludb-2"
```

The following **tools** command uses a circuit-id and a remote-id, with entry-14 being matched.

```
*A:BNG-1# tools perform subscriber-mgmt local-user-db "ludb-2"
                                     ipoe host-lookup circuit-id 11 remote-id AA
=====
IPoE Host Lookup results
=====
Result          : Success
Matched Host Name : entry-14
Admin State      : Up
Last Mgmt Change : 11/20/2018 11:25:10

Host Identification
Circuit Id      : 11
Mac Address     : N/A
Remote Id       : AA
Sap Id          : N/A
Service Id      : N/A
String          : N/A
Option 60       : N/A
System Id       : N/A
Encap Tag Range : N/A
Derived Id      : N/A
IP prefix       : N/A

Matched Objects : circ-id remote-id
---snip---
```

```
=====
*A:BNG-1#
```

The corresponding debug output shows that the original and the masked values of the circuit-id and the remote-id are the same, because no masks are applied.

```
11 2018/11/20 11:53:46.129 CET MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
mac:
circuit-id:
  original: 11
  masked:   11
remote-id:
  original: AA
  masked:   AA

Host entry-14 found in user data base ludb-2"
```

Accessing ludb-2 with a remote-id only returns a failure.

```
*A:BNG-1# tools perform subscriber-mgmt local-user-db "ludb-2"
                                                    ipoe host-lookup remote-id AA
=====
IPoE Host Lookup results
=====
Result          : host not found
*A:BNG-1#
```

```
12 2018/11/20 11:54:44.948 CET MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host not found
mac:
circuit-id:
remote-id:
  original: AA
  masked:   AA

Host not found in user data base ludb-2"
```

Example 3: Masking (1)

The following shows an excerpt from ludb-3, applying masks.

The match-list includes the circuit-id and the MAC-address, in this sequence. Masking applies to the circuit-id, which has the leading 8 characters and the trailing characters (behind the last vertical bar, and including the vertical bar) stripped.

```
configure
  subscriber-mgmt
    local-user-db "ludb-3" create
      description "masking example, ipoe"
      ipoe
        match-list circuit-id mac
        mask type circuit-id prefix-length 8 suffix-string "|"
        host "entry-111" create
          host-identification
            circuit-id string "grp-int-1-1"
            mac 00:00:00:11:11:11
          exit
        ---snip---
        no shutdown
      exit
    ---snip---
```

The following **tools** command uses circuit-id and mac-address, matching entry-111.

```
*A:BNG-1# tools perform subscriber-mgmt local-user-db "ludb-3" ipoe host-lookup
mac 00:00:00:11:11:11 circuit-id "BNG-1|1|grp-int-1-1|1/1/2/1:111"
=====
IPoE Host Lookup results
=====
Result                : Success
Matched Host Name     : entry-111
Admin State           : Up
Last Mgmt Change      : 11/20/2018 11:25:11

Host Identification
Circuit Id            : grp-int-1-1
Mac Address           : 00:00:00:11:11:11
Remote Id             : N/A
Sap Id                : N/A
Service Id            : N/A
String                : N/A
Option 60             : N/A
System Id             : N/A
Encap Tag Range       : N/A
Derived Id            : N/A
IP prefix             : N/A

Matched Objects       : circ-id mac
---snip---
=====
*A:BNG-1#
```

The debug output shows the values of the parameters before and after applying the mask.

```
13 2018/11/20 11:56:10.244 CET MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
circuit-id:
  original: BNG-1|1|grp-int-1-1|1/1/2/1:111
  masked:   grp-int-1-1
mac: 00:00:00:11:11:11

Host entry-111 found in user data base ludb-3"
```

Example 4: Masking (2)

The following shows an excerpt from ludb-4, applying masks.

The match-list includes the username, circuit-id, and remote-id, in this sequence. Masking applies to both the username and the circuit-id. The username has everything before the @-sign and the trailing .org stripped. The circuit-id has the trailing 11 characters stripped.

```
configure
  subscriber-mgmt
    local-user-db "ludb-4" create
      ppp
        match-list username circuit-id remote-id
        mask type circuit-id suffix-length 11
        mask type username prefix-string "*" suffix-string ".org"
        host "entry-11" create
          host-identification
            username domain1
            circuit-id string "BSAN-2"
          exit
        address pool "pool4-1"
```

```

        identification-strings 254 create
        sla-profile-string "sla-prof-1"
        sub-profile-string "sub-prof-2"
    exit
    no shutdown
exit
---snip---
exit
no shutdown
exit
---snip---
```

The following **tools** command does not result in a match, which is not the intention.

```

*A:BNG-1# tools perform subscriber-mgmt local-user-db "ludb-4" ppp host-lookup
        user-name sub11@domain1.org circuit-id "BSAN-2|100|1/2/1:111"
=====
PPP host Lookup results
=====
Result          : host not found
*A:BNG-1#
```

The debug output shows the original and the masked value of the user-name and the circuit-id; the remote-id was not provided.

```

14 2018/11/20 11:57:45.554 CET MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host not found
  user-name:
    original: sub11@domain1.org
    masked:   domain1
  circuit-id:
    original: BSAN-2|100|1/2/1:111
    masked:   BSAN-2|10
  remote-id:

Host not found in user data base ludb-4"
```

The preceding output shows that three more characters (the |10) must be stripped to have a successful lookup, and following configuration changes are needed.

```

configure
  subscriber-mgmt
    local-user-db "ludb-4" create
    ppp
      mask type circuit-id suffix-length 14
```

Modifying the mask results in host entry-11 being matched.

```

*A:BNG-1# tools perform subscriber-mgmt local-user-db "ludb-4" ppp host-lookup
        user-name sub11@domain1.org circuit-id "BSAN-2|100|1/2/1:111"
=====
PPP host Lookup results
=====
Result          : Success
Matched Host Name : entry-11
Admin State      : Up
Last Mgmt Change  : 11/20/2018 11:25:11

Host Identification
Mac Address      : N/A
```

```
Circuit Id       : BSAN-2
Remote Id       : N/A
Sap Id         : N/A
Service Name    : N/A
User Name      : domain1
Encap Tag Range : N/A
Derived Id     : N/A
```

```
Matched Objects : userName circ-id
```

```
---snip---
```

```
=====
*A:BNB-1#
```

The debug output shows the effect of the modified mask.

```
15 2018/11/20 11:58:29.339 CET MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
  user-name:
    original: sub11@domain1.org
    masked:   domain1
  circuit-id:
    original: BSAN-2|100|1/2/1:111
    masked:   BSAN-2
  remote-id:

Host entry-11 found in user data base ludb-4"
```

Example 5: VLAN range

LUDB matching also supports the use of encap-tag-range. Host-identification then needs a start and an end for the range, which both use the following format:

```
dot1q      - qtag1
  qinq      - (qtag1.qtag2 | qtag1.* | *.qtag2)
  atm       - (vpi/vci | vpi/* | */vci)
    qtag1    - [0..4094]
    qtag2    - [0..4094]
    vpi      - [0..4095] (NNI)
              [0..255] (UNI)
    vci      - [1..65535]
```

For VLAN tagging, the Ethernet frames could be single- or dual- tagged. For ATM, a virtual path identifier (VPI) and a virtual circuit identifier (VCI) can be defined. The asterisk (*) serves as a wildcard, meaning that the parameter is ignored.

The system validates, at configuration time, the values of the start-tag and the end-tag, applying following rules:

- The start-tag must be lower than the end-tag.
- When using the asterisk, it should be present in both the start-tag and the end-tag, as either the inner or the outer tag:
 - *.10 - *.100 — the outer tag is ignored
 - 201.* - 299.* — the inner tag is ignored
- The encapsulation type for start-tag and end-tag must be the same.
- Overlapping ranges (while on the same port) are not allowed.

The following shows an excerpt from ludb-5, using vlan-ranges.

```
configure
  subscriber-mgmt
    local-user-db "ludb-5" create
      description "example for vlan ranges"
      ipoe
        match-list encap-tag-range
        host "range-1" create
          host-identification
            encap-tag-range start-tag *.1 end-tag *.50
          exit
          address pool "pool4-3"
          ---snip---
          no shutdown
        exit
        host "range-2" create
          host-identification
            encap-tag-range start-tag *.51 end-tag *.100
          exit
          address pool "pool4-4"
          ---snip---
          no shutdown
        exit
      exit
    no shutdown
  exit
```

The following **tools** command specifies a sap-id including an outer and an inner tag, matching host range-1.

```
*A:BNG-1# tools perform subscriber-mgmt local-user-db "ludb-5"
                                                    ipoe host-lookup sap-id 1/1/1:50.4
=====
IPoE Host Lookup results
=====
Result                : Success
Matched Host Name     : range-1
Admin State           : Up
Last Mgmt Change      : 11/20/2018 11:25:11

Host Identification
Circuit Id            : N/A
Mac Address           : N/A
Remote Id             : N/A
Sap Id                : N/A
Service Id           : N/A
String               : N/A
Option 60            : N/A
System Id            : N/A
Encap Tag Range       : start-tag *.1 end-tag *.50
Derived Id           : N/A
IP prefix             : N/A

Matched Objects        : encap-tag-range

Address               : pool "pool4-3"
---snip---
=====
*A:BNG-1#
```

Operational considerations

Following operational considerations should be kept in mind:

- Names of LUDBs and LUDB entries cannot be changed.
- Modification of the host identification fields is possible only when the host-entry is put in the shutdown state. Modifying output fields does not require the host-entry to be in the shutdown state.

```
*A:BNG-1>config>subscr-mgmt>...>ppp>host>host-ident# circuit-id string x-y-z
MINOR: DHCPD #1133 Not allowed. Host is not shutdown
```

- Modifying a match-list requires the LUDB to be in the shutdown state.
- Modifying a match-list results in a re-evaluation of all host entries of the LUDB block, so that the lookup database and the unmatched host list are re-populated.

```
*A:BNG-1>config>subscr-mgmt>loc-user-db>ipoe# match-list circuit-id
INFO: DHCPD #1107 Host could not be inserted in lookup database
- lookup database constructed, 1 hosts not inserted: 1 no match, 0 duplicate
```

Modifying the match-list also imposes the risk of a default entry with host-identification fields suddenly not being the fallback (default) entry anymore, which is why defining a default entry with host-identification fields is not recommended.

- Modifying one or more mask types does not require the LUDB to be in the shutdown state.
- Deletion of an LUDB requires that the LUDB is not referenced and the LUDB is in the shutdown state. Use caution: the status of the individual entries is not taken into account when deleting an LUDB.

```
*A:BNG-1>config>subscr-mgmt# no local-user-db "ludb-1"
MINOR: DHCPD #1103 User data base still referenced

*A:BNG-1>config>subscr-mgmt# no local-user-db "ludb-11"
MINOR: DHCPD #1104 Not allowed when user db admin state is up
```

Troubleshooting and debugging LUDBs

The **tools** command can also be used for troubleshooting; the example is not repeated for brevity.

Show commands

The following command shows which LUDBs are available in the system, including the administrative state and the host count. The host count equals the total number of configured ipoe and ppp entries, regardless of their administrative state (shutdown/no shutdown).

```
*A:BNG-1# show subscriber-mgmt local-user-db

=====
Local User Databases
=====
Name                               Admin Host  Description
                                State Count
-----
```



```
ludb-1          Up    10    example user-db
ludb-2          Up     4
ludb-22        Up     5
ludb-3         Down   5
ludb-4          Up     1
ludb-5          Up     2    example for vlan ranges
-----
Number of Local User Databases : 6    Number of Hosts : 27
=====
*A:BNB-1#
```

For showing the host count and the IPoE and PPP match types for a single LUDB, following command is useful.

```
*A:BNB-1# show subscriber-mgmt local-user-db "ludb-1"

=====
Local User Database "ludb-1"
=====
Description      : example user-db
Admin State      : Up
Last Mgmt Change : 11/20/2018 11:25:10
Host Count       : 10
IPoE Match Types : mac
PPP Match Types  : userName
=====
*A:BNB-1#
```

Listing all IPoE hosts in a specific LUDB is performed with the following command.

```
*A:BNB-1# show subscriber-mgmt local-user-db "ludb-1" ipoe-all-hosts

=====
Local User Database "ludb-1" IPoE hosts
=====
Name                Admin    Matched objects
                   State
-----
default            Up      -
entry-11           Up      mac
---snip---
Number of IPoE Hosts : 5
=====
*A:BNB-1#
```

A similar command lists all PPP hosts.

```
*A:BNB-1# show subscriber-mgmt local-user-db "ludb-1" ppp-all-hosts

=====
Local User Database "ludb-1" PPP Hosts
=====
Name                Admin    Matched objects
                   State
-----
entry-55           Up      userName
---snip---
Number of PPP Hosts : 5
=====
*A:BNB-1#
```

```
*A:BNG-1#
```

To find the places where a specific LUDB is applied, use the following command.

```
*A:BNG-1# show subscriber-mgmt local-user-db "ludb-1" association
```

```
=====
DHCP Server associations for ludb-1
=====
```

Server-Name	Router-Name
dhcp4-srv	Base

dhcp4-srv	Base
-----------	------

```
No. of Server(s): 1
=====
```

```
=====
DHCP client interface associations for ludb-1
=====
```

Interface-Name	Svc-Id	Type
grp-int-1-1	1	IES
grp-int-1-2	1	IES
grp-int-2-1	1	IES
grp-int-2-2	1	IES

grp-int-1-1	1	IES
grp-int-1-2	1	IES
grp-int-2-1	1	IES
grp-int-2-2	1	IES

```
No. of Interface(s): 4
=====
```

```
=====
DHCP6 interface associations for ludb-1
=====
```

Interface-Name	Svc-Id	Type
grp-int-1-1	1	IES
grp-int-1-2	1	IES
grp-int-2-1	1	IES
grp-int-2-2	1	IES

grp-int-1-1	1	IES
grp-int-1-2	1	IES
grp-int-2-1	1	IES
grp-int-2-2	1	IES

```
No. of Interface(s): 4
=====
```

```
No Router solicit interface associations found.
=====
```

```
=====
PPP client interface associations for ludb-1
=====
```

Interface-Name	Svc-Id	Type
grp-int-1-1	1	IES
grp-int-1-2	1	IES
grp-int-2-1	1	IES
grp-int-2-2	1	IES

grp-int-1-1	1	IES
grp-int-1-2	1	IES
grp-int-2-1	1	IES
grp-int-2-2	1	IES

```
No. of Interface(s): 4
=====
```

```
No PPPoE client interface associations found.
```

```
No IPoE client interface associations found.
```

```
No capture SAP associations found.
```

```
No associated L2TP groups found.
No associated L2TP tunnels found.
No associated authentication policies found.
No WPP interface associations found.
No GTP APN policy associations found.
*A:BNG-1#
```

The following command is useful for displaying the details of a specific LUDB entry.

```
*A:BNG-1# show subscriber-mgmt local-user-db "ludb-1" ipoe-host "entry-33"

=====
IPoE Host "entry-33"
=====
Admin State           : Up
Last Mgmt Change      : 11/20/2018 11:25:10

Host Identification
Circuit Id            : N/A
Mac Address           : 00:00:00:33:33:33
---snip---

Matched Objects       : mac

Address                : use GI-address (scope subnet)
---snip---
IPv6 Address Pool     : pool6-3
IPv6 Del Pfx Pool     : pool6-3
IPv6 Slaac Pfx Pool   : N/A
IPv6 Del Pfx Length   : N/A

---snip---

Identification Strings (option 254)
Subscriber Id         : sub-33
SLA Profile String    : sla-prof-2
SPI Sharing Group Id  : N/A
Sub Profile String    : sub-prof-4
App Profile String    : N/A
ANCP String           : N/A
Inter Destination Id  : N/A
Category Map Name     : N/A

---snip---

Filter Overrides
Ing Ipv4 Fltr         : N/A
Egr Ipv4 Fltr         : N/A
Ing Ipv6 Fltr         : N/A
Egr Ipv6 Fltr         : N/A
=====
*A:BNG-1#
```

The following commands list the IPoE and the PPP host entries in a specific LUDB that are not matched. Duplicates are counted as unmatched hosts.

```
*A:BNG-1# show subscriber-mgmt local-user-db "ludb-22" ipoe-unmatched-hosts
```

```
=====
Local User Database "ludb-22" IPoE unmatched hosts
=====
Name                      Reason      Duplicate Host
-----
this-is-a-no-match        No match    N/A
this-is-a-duplicate        Duplicate    entry-12
-----
Number of IPoE Unmatched Hosts : 2
=====
*A:BNB-1#
```

```
*A:BNB-1# show subscriber-mgmt local-user-db "ludb-22" ppp-unmatched-hosts

=====
Local User Database "ludb-22" PPP unmatched hosts
=====
Name                      Reason      Duplicate Host
-----
No PPP Unmatched Hosts found
=====
*A:BNB-1#
```

Debugging commands

The following configuration enables debugging for both ludb-1 and for ludb-2.

```
debug
  subscriber-mgmt
    local-user-db "ludb-1"
    detail all
  exit
  local-user-db "ludb-2"
  detail all
exit
exit
exit
```

To ensure that the debug output is sent to the console, the following additional configuration is needed.

```
configure
  log
    log-id 1
    from debug-trace
    to session
    no shutdown
  exit
exit
```

After the preceding configuration, debug output appears as part of the session.

```
2 2018/11/20 11:42:27.965 CET MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
mac: 00:00:00:11:11:11
```

```
Host entry-11 found in user data base ludb-1"
```

```
12 2018/11/20 11:54:44.948 CET MINOR: DEBUG #2001 Base LUDB"LUDB: User lookup success - host  
not found  
mac:  
circuit-id:  
remote-id:  
  original:  AA  
  masked:    AA
```

```
Host not found in user data base ludb-2"
```

```
37 2018/11/20 12:18:59.141 CET MINOR: DEBUG #2001 Base LUDB  
"LUDB: User lookup failed  
  Problem: user db is shutdown"
```

Conclusion

In this chapter general LUDB concepts are explained. LUDBs are defined and host entries for both IPoE as well as for PPP are described. The different match criteria are explained and demonstrated by means of examples, including the use of single and multiple match criteria. Match criteria are handled left to right, in sequence, so that a natural priority is taken care of. Debugging aids are provided through **show**, **debug** and **tools** commands.

Local User Database for DHCPv4 Server

This chapter provides information about local user database (LUDB) for DHCPv4 server.

Topics in this chapter include:

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

Applicability

This chapter was initially written based on SR OS Release 13.0.R1, but the information and configuration in the current edition are based on SR OS Release 16.0.R4.

Basic LUDB knowledge is a prerequisite for understanding this chapter.

Overview

In SR OS, a local DHCPv4 server can be assigned a local user database (LUDB).

Assigning an LUDB to a DHCPv4 server allows the server to:

- control IP address assignment; for example, by assigning a fixed IP address based on the user's MAC address.
- control DHCPv4 options for native as well as for simulated DHCPv4 clients used by PPP. In the case of PPP, users are identified to the DHCPv4 server using the DHCPv4 Vendor-Specific Information Sub-option [82,9][6] in the DHCPv4 discover/request messages.
- provide ESM strings (referred to as identification-strings in CLI) using a user-defined unassigned DHCPv4 option (option 254 is provided as default).

Introduction

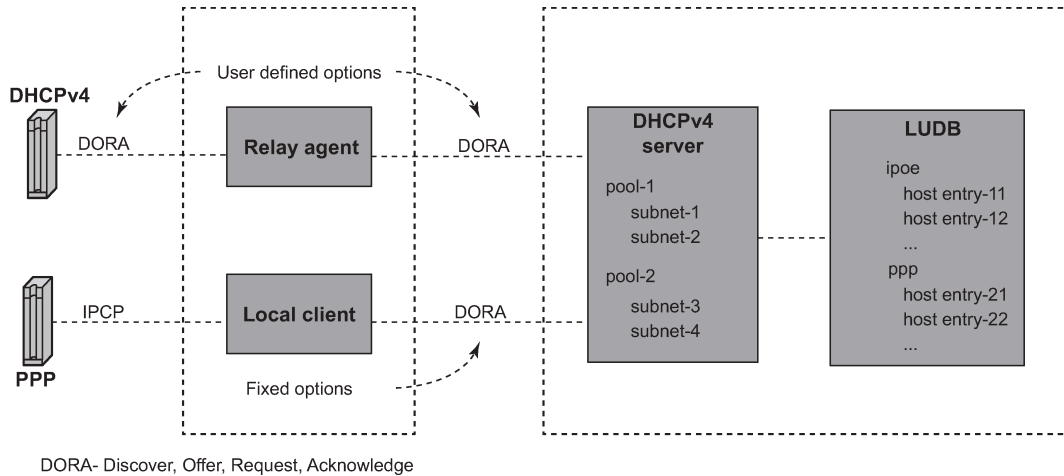
A local DHCPv4 server can be used for IPoE users as well as for PPP users (see [Figure 128: LUDB access via a DHCPv4 server](#)).

When a DHCPv4 user connects, the typical discover, offer, request, and acknowledge (DORA) message sequence running between the DHCPv4 client and the DHCPv4 server also passes through a relay agent.

When a PPP user connects through LCP/IPCP, an internal DHCPv4 client manages the communication toward the DHCPv4 server, on the condition that the relay agent also has relaying enabled for PPP applications. This internal DHCPv4 client is also referred to as a local (DHCPv4) client.

The DHCPv4 server can be located in the same node as the relay agent and the local client, but that is not required.

Figure 128: LUDB access via a DHCPv4 server



The relay agent must be configured correctly in order to forward the messages toward a DHCPv4 server. The DHCPv4 server IP address is defined. The relay agent can include following configurable options and sub-options to be used by the DHCPv4 server:

- [82,1] Agent circuit ID sub-option
- [82,2] Agent remote ID sub-option
- [82,9] Vendor-specific information sub-option (VSO)
 - [1] system ID
 - [2] client MAC address
 - [3] service ID (IES/VP RN service ID)
 - [4] SAP ID
 - [5] string
 - [13] pool name

The local client uses the same DHCPv4 server IP address as the preceding relay agent, and can include its own set of (unconfigurable) options and sub-options to be used by the DHCPv4 server:

- [60] vendor class (fixed string: ALU7XXXSBM)
- [82,1] Agent circuit ID sub-option
- [82,2] Agent remote ID sub-option
- [82,6] Subscriber ID sub-option (equals PPPoE username)
- [82,9] Vendor-specific information sub-option
 - [1] system ID (not included in a redundant node)
 - [2] client MAC address
 - [3] service ID (IES/VP RN service ID)
 - [4] SAP ID (not included for retail VP RN and redundant node)

- [6] client type (1=PPP)
- [13] pool name
- [14] service name (PPPoE tag service name)
- [17] session ID (PPPoE session ID)

When one or more of these options and sub-options are included, the DHCPv4 server can use them while accessing the LUDB, for selection of the section (client type 1 is the PPP section), and the host entry in that section. For example, not including the service ID VSO [82,9][3] to the DHCPv4 server, when LUDB host identification needs the MAC address and service ID, will result in an LUDB lookup failure on the DHCPv4 server and a silent drop of the DHCPv4 discover message.

LUDB input parameters

The following IPoE host identification fields are supported when accessing an LUDB from a DHCPv4 server:

- circuit-id
- encap-tag-range
- mac
- option60
- remote-id
- sap-id
- service-id
- string
- system-id

The LUDB lookup process can match up to four IPoE match-criteria, as defined by the IPoE match-list.

The following PPP host identification fields are supported when accessing an LUDB from a DHCPv4 server:

- circuit-id
- encap-tag-range
- mac
- remote-id
- sap-id
- service-name
- username

The LUDB lookup process can match up to three PPP match-criteria, as defined by the PPP match-list.

LUDB output parameters

Addressing information

The host entry address field has the following configuration options when the LUDB is associated with a DHCPv4 server:

- no address

Host access is not allowed. The clients mapping to this host entry will not get an IP address.

```
23 2018/11/21 13:20:09.958 CET MINOR: DEBUG #2001 Base DHCP server
"DHCP server:  dhcp4-srv
DISCOVER dropped: host=entry-66, host found but no valid address info
```

- address *<ip-address>*

A fixed IP address is offered to the client and should not overlap with the address ranges configured in the local DHCP server.

- address pool *<pool-name>* [secondary-pool *<sec-pool-name>*]

The DHCPv4 server allocates an address from one of the subnets in that pool on the condition that the DHCPv4 server is configured to use pool names for address selection. Optionally, a secondary pool can be defined, to be used in case the primary pool is exhausted.

Pool-name addressing is useful when the subscriber management node is not capable of inserting the pool-name VSO [82,9][13] or when a specific host requires a specific pool-name different from the pool-name included by the subscriber management node.

- address gi-address [scope *<subnet | pool>*]

When the scope is set to subnet, the DHCPv4 server allocates an address from the subnet that includes the Gi address. When the scope is set to pool, the DHCPv4 server allocates an address from the subnet that includes the Gi address, or from the other subnets belonging to the same pool. Gi-addressing is useful when the subnet that the Gi address belongs to is exhausted.

- address use-pool-from-client [delimiter *<delimiter>*]

The DHCPv4 server allocates an address from one of the subnets in the pool, as indicated by the pool-name VSO [82,9][13]. If two pools are available in this VSO, the configured delimiter distinguishes the first pool-name from the second pool-name.

Identification strings

An LUDB can optionally return identification strings (also known as ESM strings). The DHCPv4 server returns them in a user-defined DHCPv4 option (default: 254) to the requesting entity (the relay agent or the local client). The identification strings, in alphabetical order, are:

- ancp-string
- app-profile-string
- category-map-name
- inter-dest-id

- sla-profile-string
- sub-profile-string
- subscriber-id

Options

An LUDB can return options to be used by the relay agent, the internal client, or the end-user device.

The IPoE user options configurable by option-name are:

- default-router
- dns-server
- domain-name
- lease-rebind-timer
- lease-renew-timer
- lease-time
- netbios-name-server
- netbios-node-type
- subnet-mask

The PPP user options configurable by option-name are:

- dns-server
- netbios-name-server

Additional options, configurable by option-number, can be configured for both IPoE and PPP users by using the **custom-option** command.

```
*A:BNG-1>config>subscr-mgmt>loc-user-db>ipoe>host>options# custom-option ?
- custom-option <option-number> address [<ip-address>...(upto 4 max)]
- custom-option <option-number> hex <hex-string>
- custom-option <option-number> string <ascii-string>
- no custom-option <option-number>

<option-number>      : [1..254]
<ip-address>         : a.b.c.d
<ascii-string>       : [127 chars max]
<hex-string>         : [0x0..0xFFFFFFFF...(max 254 hex nibbles)]

*A:BNG-1>config>subscr-mgmt>loc-user-db>ipoe>host>options#
```

The encoding of these custom options is either in hexadecimal, ASCII, or IP address format. In debug output, these custom options are indicated as *Unknown options* and presented in Type-Length-Value (TLV) format.

Options and custom options can be configured at three different levels:

- LUDB host entry level
- DHCPv4 server pool level
- DHCPv4 server subnet level

Options and custom options defined at the host entry level overrule options defined at either of the server levels.

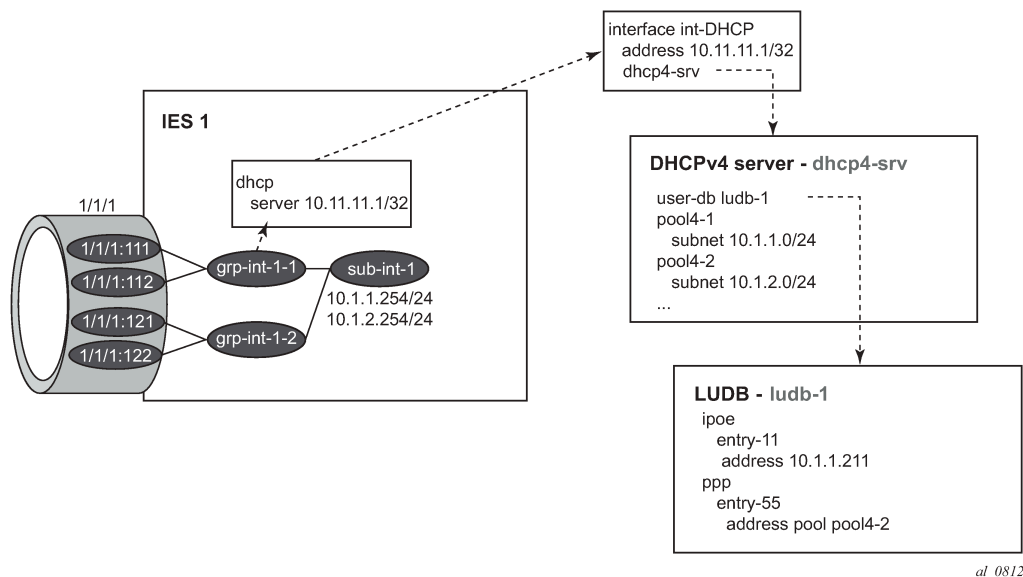
Other parameters

All other parameters in the host entry definition are silently ignored (for example, IPv6 related parameters, msap-defaults, retail-service-id, and so on) because they are not applicable for a DHCPv4 server-associated LUDB.

Configuration

Figure 129: Example configuration shows the example configuration used in this chapter.

Figure 129: Example configuration



An LUDB can be associated with a DHCPv4 server in the base router instance or in a VPRN service instance using the following commands:

```
configure router dhcp local-dhcp-server <name> user-db <name>
configure service vprn <service-id> dhcp local-dhcp-server <name> user-db <name>
```

The following example has the DHCPv4 server created in the base router instance.

```
configure
router
dhcp
    local-dhcp-server "dhcp4-srv" create
    user-db "ludb-1"
    use-gi-address
    use-pool-from-client
    pool "pool4-1" create
    subnet 10.1.1.0/24 create
```

```

        options
            subnet-mask 255.255.255.0
            default-router 10.1.1.254
        exit
        address-range 10.1.1.1 10.1.1.100
    exit
exit
pool "pool4-2" create
    subnet 10.1.2.0/24 create
        options
            subnet-mask 255.255.255.0
            default-router 10.1.2.254
        exit
        address-range 10.1.2.1 10.1.2.254
    exit
exit
---snip---
no shutdown
exit
exit

```

The server is then associated with the loopback interface at address 10.11.11.1, as follows:

```

configure
router
    interface "int-DHCP"
        address 10.11.11.1/32
        loopback
        local-dhcp-server "dhcp4-srv"
        no shutdown
    exit
exit

```

The following is a partial configuration of service IES 1. Note that the DHCPv4 relay agent is configured to include the service-ID VSO [82,9][3] because this option is used in the LUDB for matching purposes.

```

configure
service
    ies 1 customer 1 create
        subscriber-interface "sub-int-1" create
            address 10.1.1.254/24
            address 10.1.2.254/24
            group-interface "grp-int-1-1" create
                arp-populate
                dhcp
                option
                    action replace
                    circuit-id
                    remote-id
                    vendor-specific-option
                    pool-name
                    service-id
                exit
            exit
            server 10.11.11.1
            trusted
            lease-populate 100
            client-applications dhcp ppp
            gi-address 10.1.1.254
            no shutdown
        exit
    exit

```

```

        sap 1/1/1:111 create
        sub-sla-mgmt
            def-sub-profile "sub-prof-1"
            def-sla-profile "sla-prof-1"
            sub-ident-policy "sub-id-pol-1"
            multi-sub-sap
            no shutdown
        exit
    exit
    ---snip---
    pppoe
        session-limit 100
        sap-session-limit 100
        no shutdown
    exit
exit
exit
exit

```

In the following example, IPoE users are matched against the service-id, the MAC address and option 60. Host *entry-11* returns a fixed IP address, three identification strings, and a set of options. PPP users are matched against the MAC address. Host *entry-55* returns an address-pool, two identification strings, and two DNS servers as options.

```

configure
  subscriber-mgmt
    local-user-db "ludb-1" create
    description "example user-db"
    ipoe
      match-list service-id mac option60
      host "entry-11" create
        host-identification
          mac 00:00:00:11:11:11
          service-id 1
          option60 hex 0xaabb
        exit
      address 10.1.1.211
      identification-strings 254 create
        subscriber-id "sub-11"
        sla-profile-string "sla-profile-1"
        sub-profile-string "sub-profile-1"
      exit
      options
        subnet-mask 255.255.255.0
        default-router 10.1.1.251
        dns-server 2.2.2.2 2.2.2.1
        domain-name "domain.org"
        netbios-name-server 10.1.1.252
        netbios-node-type B
        lease-time hrs 12
        custom-option 251 hex 0x010203
      exit
    no shutdown
  exit
  ---snip---
exit
ppp
  match-list mac
  host "entry-55" create
    host-identification
      mac 00:00:00:55:55:55
    exit
  address pool "pool4-2"
  identification-strings 254 create

```

```

        subscriber-id "sub-55"
        sla-profile-string "sla-prof-3"
        sub-profile-string "sub-prof-3"
    exit
    options
        dns-server 2.2.2.2 2.2.2.1
    exit
    ---snip---
    no shutdown
exit
exit
no shutdown
exit

```

Entry-11 defines a fixed IP address in one of the subnets allowed on the group interface *grp-int-1-1* on IES 1, but out of the range defined in the DHCPv4 server.

Debugging

The following example debugs the local DHCP server and the LUDB *ludb-1*.

```

debug
  router "Base"
    local-dhcp-server "dhcp4"
    detail-level medium
    mode egr-ingr-and-dropped
  exit
exit
subscriber-mgmt
  local-user-db "ludb-1"
  detail all
exit
exit
exit

```

The following additional configuration ensures that the debug output is sent to the current login session.

```

configure
  log
    log-id 1
    from debug-trace
    to session
    no shutdown
  exit
exit
exit

```

IPoE users verification

The following command shows the DHCPv4 server lease state record for LUDB host *entry-11*. The address type is set to *fixed* because *ludb-1* returns a fixed IP address.

```

*A:BNG-1# show router dhcp local-dhcp-server "dhcp4-srv" leases 10.1.1.211 detail
=====
Lease for DHCP server dhcp4-srv router Base
=====

```

```

IP-address           : 10.1.1.211
Lease-state          : stable
Lease started        : 2018/11/21 13:14:34
Last renew           : N/A
Remaining LifeTime    : 11h59m31s
Remaining Potential Exp. Time: 0h0m0s
Sticky-lease Host Name : N/A
MAC address          : 00:00:00:11:11:11
Xid                   : 0x1
Failover Control      : local
Client Type         : dhcp
User-db Host Name   : entry-11
User-db Address Type : fixed
Persistence Key       : N/A
Opt82 Hex Dump        : (length=54)
                       : 52 34 01 1d 42 4e 47 2d 31 7c 31 7c 67 72 70 2d
                       : 69 6e 74 2d 31 2d 31 7c 31 2f 31 2f 31 3a 31 31
                       : 31 02 06 00 00 00 11 11 11 09 0b 00 00 19 7f 06
                       : 03 04 00 00 00 01
Opt82 Circuit Id      : BNG-1|1|grp-int-1-1|1/1/1:111
Opt82 Remote Id       : (hex) 00 00 00 11 11 11
Opt82 Subscr Id       :
Opt82 VS System       :
Opt82 VS Clnt MAC     :
Opt82 VS Service      : (hex) 00 00 00 01
Opt82 VS SAP          :
Opt82 VS String       :
Opt60 Hex Dump        : (length=2)
                       : aa bb
Lease Remaining Hold Time : 0h0m0s

=====
*A:BNG-1#

```

The debug output on the DHCPv4 server and the LUDB *ludb-1* shows that the LUDB is accessed for every incoming message. The identification-strings are returned to the relay agent in the Offer and Acknowledge messages through option [254]; see [Figure 130: Decoding the ESM user option](#) for the decoding.

```

1 2018/11/21 13:14:34.703 CET MINOR: DEBUG #2001 Base DHCP server
"DHCP server: dhcp4-srv
Rx DHCP Discover

ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
siaddr: 0.0.0.0          giaddr: 10.1.1.254
chaddr: 00:00:00:11:11:11 xid: 0x1

DHCP options:
[82] Relay agent information: len = 52
    [1] Circuit-id: BNG-1|1|grp-int-1-1|1/1/1:111
    [2] Remote-id: (hex) 00 00 00 11 11 11
    [9] Vendor-Specific info: len = 11
        Enterprise [6527] : len = 6
        [3] servId: 1
[53] Message type: Discover
[60] Class id: (hex) aa bb
[255] End
"

2 2018/11/21 13:14:34.703 CET MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
  svc-id: 1
  mac: 00:00:00:11:11:11
"

```

option60:

original: 0xaabb
masked: 0xaabb

Host entry-11 found in user data base lddb-1"

```
3 2018/11/21 13:14:34.703 CET MINOR: DEBUG #2001 Base DHCP server
"DHCP server:  dhcp4-srv
lease added for 10.1.1.211 state=offer
"
```

```
4 2018/11/21 13:14:34.703 CET MINOR: DEBUG #2001 Base DHCP server
"DHCP server:  dhcp4-srv
Tx DHCP Offer to local relay agent 10.1.1.254 vrId=1
```

```
ciaddr: 0.0.0.0          yiaddr: 10.1.1.211
siaddr: 10.11.11.1       giaddr: 10.1.1.254
chaddr: 00:00:00:11:11:11 xid: 0x1
```

DHCP options:

```
[82] Relay agent information: len = 52
    [1] Circuit-id: BNG-1|1|grp-int-1-1|1/1/1:111
    [2] Remote-id: (hex) 00 00 00 11 11 11
    [9] Vendor-Specific info: len = 11
        Enterprise [6527] : len = 6
        [3] servId: 1
[53] Message type: Offer
[54] DHCP server addr: 10.11.11.1
[51] Lease time: 43200
[1] Subnet mask: 255.255.255.0
[3] Router: 10.1.1.254
[6] Domain name server: length = 8
    2.2.2.2
    2.2.2.1
[15] Domain name: domain.org
[44] NETBIOS name server: 10.1.1.252
[46] NETBIOS type: 1
[251] Unknown option: len = 3, value = 01 02 03
[254] Unknown option: len = 38, value = 07 06 73 75 62 2d 31 31 08 0d 73
6c 61 2d 70 72 6f 66 69 6c 65 2d 31 09 0d 73 75 62 2d 70 72 6f 66 69 6c 65
2d 31
[60] Class id: (hex) aa bb
[255] End
"
```

```
5 2018/11/21 13:14:34.718 CET MINOR: DEBUG #2001 Base DHCP server
"DHCP server:  dhcp4-srv
Rx DHCP Request
```

```
ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
siaddr: 0.0.0.0          giaddr: 10.1.1.254
chaddr: 00:00:00:11:11:11 xid: 0x1
```

DHCP options:

```
[82] Relay agent information: len = 52
    [1] Circuit-id: BNG-1|1|grp-int-1-1|1/1/1:111
    [2] Remote-id: (hex) 00 00 00 11 11 11
    [9] Vendor-Specific info: len = 11
        Enterprise [6527] : len = 6
        [3] servId: 1
[53] Message type: Request
```



```
[50] Requested IP addr: 10.1.1.211
[60] Class id: (hex) aa bb
[54] DHCP server addr: 10.11.11.1
[255] End
"

6 2018/11/21 13:14:34.718 CET MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
  svc-id: 1
  mac: 00:00:00:11:11:11
  option60:
    original: 0xaabb
    masked: 0xaabb

  Host entry-11 found in user data base ludb-1"

7 2018/11/21 13:14:34.718 CET MINOR: DEBUG #2001 Base DHCP server
"DHCP server: dhcp4-srv
lease update for 10.1.1.211 state=stable
"

8 2018/11/21 13:14:34.718 CET MINOR: DEBUG #2001 Base DHCP server
"DHCP server: dhcp4-srv
Tx DHCP Ack to local relay agent 10.1.1.254 vrId=1

  ciaddr: 0.0.0.0          yiaddr: 10.1.1.211
  siaddr: 10.11.11.1       giaddr: 10.1.1.254
  chaddr: 00:00:00:11:11:11  xid: 0x1

DHCP options:
[82] Relay agent information: len = 52
  [1] Circuit-id: BNG-1|1|grp-int-1-1|1/1/1:111
  [2] Remote-id: (hex) 00 00 00 11 11 11
  [9] Vendor-Specific info: len = 11
    Enterprise [6527] : len = 6
    [3] servId: 1
[53] Message type: Ack
[54] DHCP server addr: 10.11.11.1
[51] Lease time: 43200
[1] Subnet mask: 255.255.255.0
[3] Router: 10.1.1.254
[6] Domain name server: length = 8
    2.2.2.2
    2.2.2.1
[15] Domain name: domain.org
[44] NETBIOS name server: 10.1.1.252
[46] NETBIOS type: 1
[251] Unknown option: len = 3, value = 01 02 03
[254] Unknown option: len = 38, value = 07 06 73 75 62 2d 31 31 08 0d 73
6c 61 2d 70 72 6f 66 69 6c 65 2d 31 09 0d 73 75 62 2d 70 72 6f 66 69 6c 65
2d 31
[60] Class id: (hex) aa bb
[255] End
"
```

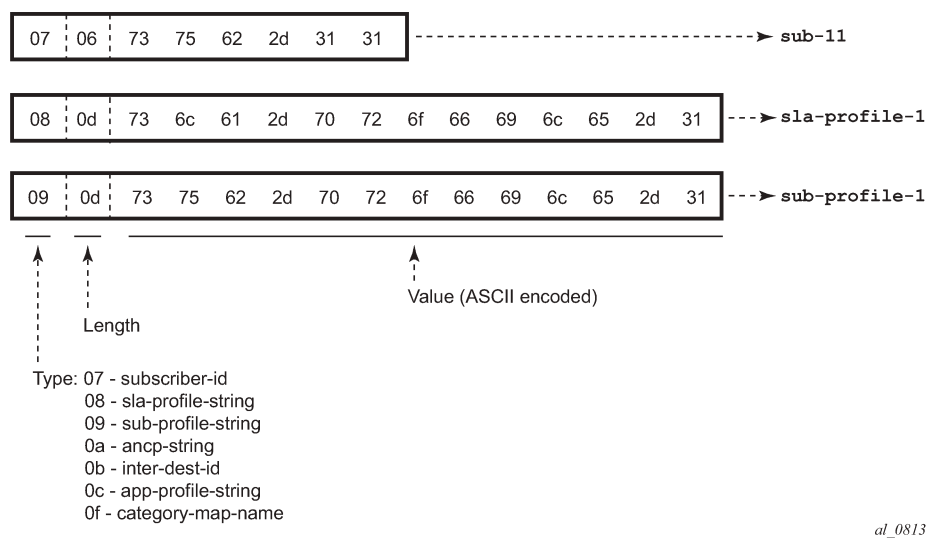
The user-defined option value for carrying the ESM strings and taken from the LUDB (identification-strings) should match the value defined in the strings-from-option parameter (value 254 in the following example) referenced in the subscriber identification policy on the relaying node.

```
configure
```

```
subscriber-mgmt
  sub-ident-policy "sub-id-pol-1" create
  sub-profile-map
    use-direct-map-as-default
  exit
  sla-profile-map
    use-direct-map-as-default
  exit
  strings-from-option 254
exit
```

The DHCP options in the debug output match the definition of the options in the LUBD.

Figure 130: Decoding the ESM user option



PPP users verification

The following command shows the DHCPv4 server lease state record for LUBD host *entry-55*. The client type is set to ppp, and the address type is set to dynamic as *ludb-1* returns address-pool *pool4-2* for this host.

```
*A:BNG-1# show router dhcp local-dhcp-server "dhcp4-srv" leases 10.1.2.2 detail

=====
Lease for DHCP server dhcp4-srv router Base
=====
IP-address           : 10.1.2.1
Lease-state          : stable
Lease started        : 2018/11/21 13:17:06
Last renew           : N/A
Remaining LifeTime    : 0h58m36s
Remaining Potential Exp. Time: 0h0m0s
Sticky-lease Host Name : N/A
MAC address          : 00:00:00:55:55:55
Xid                   : 0x2f1dd730
Failover Control      : local
Client Type           : ppp
```

```

User-db Host Name      : entry-55
User-db Address Type   : dynamic
Persistence Key        : N/A
Opt82 Hex Dump         : (length=83)
                       : 52 51 01 1d 42 4e 47 2d 31 7c 31 7c 67 72 70 2d
                       : 69 6e 74 2d 31 2d 31 7c 31 2f 31 2f 31 3a 31 31
                       : 31 02 06 00 00 00 55 55 55 09 28 00 00 19 7f 23
                       : 02 06 00 00 00 55 55 55 06 01 01 01 05 42 4e 47
                       : 2d 31 03 04 00 00 00 01 04 09 31 2f 31 2f 31 3a
                       : 31 31 31
Opt82 Circuit Id       : BNG-1|1|grp-int-1-1|1/1/1:111
Opt82 Remote Id        : (hex) 00 00 00 55 55 55
Opt82 VS System        : BNG-1
Opt82 VS Clnt MAC      : 00:00:00:55:55:55
Opt82 VS Service       : (hex) 00 00 00 01
Opt82 VS SAP           : 1/1/1:111
Opt82 VS String        :
Opt82 VS PPPoE Session ID :
Opt60 Hex Dump         : (length=10)
                       : 41 4c 55 37 58 58 58 53 42 4d
Lease Remaining Hold Time : 0h0m0s

=====
*A:BNG-1#

```

The debug output on the DHCPv4 server and the LUDB *ludb-1* shows that the LUDB is accessed for every incoming message, as follows. Again the identification-strings are returned to the relay agent in the Offer and Acknowledge messages using option [254].

```

12 2018/11/21 13:17:06.271 CET MINOR: DEBUG #2001 Base DHCP server
"DHCP server: dhcp4-srv
Rx DHCP Discover

  ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
  siaddr: 0.0.0.0          giaddr: 10.1.1.254
  chaddr: 00:00:00:55:55:55  xid: 0x2f1dd730

DHCP options:
[82] Relay agent information: len = 81
    [1] Circuit-id: BNG-1|1|grp-int-1-1|1/1/1:111
    [2] Remote-id: (hex) 00 00 00 55 55 55
    [9] Vendor-Specific info: len = 40
        Enterprise [6527] : len = 35
        [2] clntMac: 00:00:00:55:55:55
        [6] clntType: 1
        [1] systemId: BNG-1
        [3] servId: 1
        [4] sapId: 1/1/1:111
[51] Lease time: 3600
[53] Message type: Discover
[60] Class id: ALU7XXXSBM
[255] End
"

13 2018/11/21 13:17:06.271 CET MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
  mac: 00:00:00:55:55:55

  Host entry-55 found in user data base ludb-1"

14 2018/11/21 13:17:06.271 CET MINOR: DEBUG #2001 Base DHCP server

```

```
"DHCP server: dhcp4-srv
lease added for 10.1.2.1 state=offer
"

15 2018/11/21 13:17:06.271 CET MINOR: DEBUG #2001 Base DHCP server
"DHCP server: dhcp4-srv
Tx DHCP Offer to local client 10.1.1.254 vrId=1

ciaddr: 0.0.0.0          yiaddr: 10.1.2.1
siaddr: 10.11.11.1       giaddr: 10.1.1.254
chaddr: 00:00:00:55:55:55  xid: 0x2f1dd730

DHCP options:
[82] Relay agent information: len = 81
    [1] Circuit-id: BNG-1|1|grp-int-1-1|1/1/1:111
    [2] Remote-id: (hex) 00 00 00 55 55 55
    [9] Vendor-Specific info: len = 40
        Enterprise [6527] : len = 35
        [2] clntMac: 00:00:00:55:55:55
        [6] clntType: 1
        [1] systemId: BNG-1
        [3] servId: 1
        [4] sapId: 1/1/1:111
[53] Message type: Offer
[54] DHCP server addr: 10.11.11.1
[51] Lease time: 3600
[1] Subnet mask: 255.255.255.0
[6] Domain name server: length = 8
    2.2.2.2
    2.2.2.1
[254] Unknown option: len = 32, value = 07 06 73 75 62 2d 35 35 08 0a 73
6c 61 2d 70 72 6f 66 2d 33 09 0a 73 75 62 2d 70 72 6f 66 2d 33
[3] Router: 10.1.2.254
[60] Class id: ALU7XXXSBM
[255] End
"

16 2018/11/21 13:17:06.271 CET MINOR: DEBUG #2001 Base DHCP server
"DHCP server: dhcp4-srv
Rx DHCP Request

ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
siaddr: 0.0.0.0          giaddr: 10.1.1.254
chaddr: 00:00:00:55:55:55  xid: 0x2f1dd730

DHCP options:
[82] Relay agent information: len = 81
    [1] Circuit-id: BNG-1|1|grp-int-1-1|1/1/1:111
    [2] Remote-id: (hex) 00 00 00 55 55 55
    [9] Vendor-Specific info: len = 40
        Enterprise [6527] : len = 35
        [2] clntMac: 00:00:00:55:55:55
        [6] clntType: 1
        [1] systemId: BNG-1
        [3] servId: 1
        [4] sapId: 1/1/1:111
[50] Requested IP addr: 10.1.2.1
[51] Lease time: 3600
[53] Message type: Request
[54] DHCP server addr: 10.11.11.1
[60] Class id: ALU7XXXSBM
[255] End
```

```
"
17 2018/11/21 13:17:06.271 CET MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
  mac: 00:00:00:55:55:55

  Host entry-55 found in user data base ludb-1"

18 2018/11/21 13:17:06.271 CET MINOR: DEBUG #2001 Base DHCP server
"DHCP server:  dhcp4-srv
lease update for 10.1.2.1 state=stable
"

19 2018/11/21 13:17:06.271 CET MINOR: DEBUG #2001 Base DHCP server
"DHCP server:  dhcp4-srv
Tx DHCP Ack to local client 10.1.1.254 vrId=1

  ciaddr: 0.0.0.0          yiaddr: 10.1.2.1
  siaddr: 10.11.11.1      giaddr: 10.1.1.254
  chaddr: 00:00:00:55:55:55  xid: 0x2f1dd730

  DHCP options:
  [82] Relay agent information: len = 81
    [1] Circuit-id: BNG-1|1|grp-int-1-1|1/1/1:111
    [2] Remote-id: (hex) 00 00 00 55 55 55
    [9] Vendor-Specific info: len = 40
      Enterprise [6527] : len = 35
      [2] clntMac: 00:00:00:55:55:55
      [6] clntType: 1
      [1] systemId: BNG-1
      [3] servId: 1
      [4] sapId: 1/1/1:111
  [53] Message type: Ack
  [54] DHCP server addr: 10.11.11.1
  [51] Lease time: 3600
  [1] Subnet mask: 255.255.255.0
  [6] Domain name server: length = 8
    2.2.2.2
    2.2.2.1
  [254] Unknown option: len = 32, value = 07 06 73 75 62 2d 35 35 08 0a 73
6c 61 2d 70 72 6f 66 2d 33 09 0a 73 75 62 2d 70 72 6f 66 2d 33
  [3] Router: 10.1.2.254
  [60] Class id: ALU7XXSBM
  [255] End
"
```

Operational considerations

A DHCPv4 server with an LUDB cannot be used for supporting Local Address Assignment (LAA) scenarios. LAA can be used when the DHCPv4 relay agent and the DHCPv4 server are in the same node, and where IP address assignment for PPP users happens through the API directly into the local DHCPv4 server. See the [ESM SLAAC Prefix Assignment via Local Address Server](#) chapter for details on this topic.

Conclusion

This chapter explained and demonstrated the use of an LUDB in combination with a DHCPv4 server. For the DHCPv4 server to find matching entries in the LUDB, the relay agent has to be configured with the correct options and sub-options so that entries can be matched in the LUDB. It was noted that the DHCPv4 server does not require the relaying function to be located in the same node. The input and the output parameters of a DHCPv4 server-attached LUDB were listed. Carrying over the identification/ESM strings in a user-defined DHCP option was configured and demonstrated, and the decoding of this option was explained.

Local User Database for Enhanced Subscriber Management

This chapter provides information about local user database for enhanced subscriber management.

Topics in this chapter include:

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

Applicability

This chapter was initially written based on SR OS Release 13.0.R6. The CLI in the current edition corresponds to SR OS Release 16.0.R4.

Having knowledge of [ESM Basics](#), the [Routed CO](#) model, and [Local User Database Basics](#) are prerequisites for understanding this chapter.

Overview

A local user database (LUDB) is a data source providing enhanced subscriber management (ESM) data so that subscribers and subscriber hosts can be instantiated when end-users connect their devices. ESM data includes identification strings, IP address/prefix, profiles, and so on. See the [ESM Basics](#) chapter for more information.

LUDBs offer a self-contained method for providing the ESM data, so that no additional ESM data sources are needed.

Alternative ESM data sources are: RADIUS, Diameter NASREQ, Diameter Gx, DHCP-server, Python, and defaults.

Mixed scenarios, where part of the data is provided by an LUDB and the remaining part is provided through RADIUS, are the subject of the [Flexible Authentication Model in ESM](#) chapter.

LUDBs can be used for the following applications:

- assisting a DHCPv4 server in assigning fixed IP addresses to dedicated devices; see the [Local User Database for DHCPv4 Server](#) chapter.
- authenticating devices, so that ESM hosts and subscribers can be instantiated. This is supported for the Routed CO model only.
- authenticating devices as a fallback for RADIUS authentication, in case the RADIUS server is not available. This is supported for the Routed CO model only.

This chapter describes the use of LUDBs for authentication, including:

- parameters that can be returned by LUDBs
- contexts where LUDBs can be applied in the system


LUDB authentication is supported:

- for IPoE as well as for PPP
- for regular SAPs as well as for capture and managed SAPs
- for the proxy scenario as well as for the relay scenario

LUDB authentication can be started directly through one of the following protocol triggers; see [Figure 131: LUDB authentication](#):

- DHCPv4 Discover
- DHCPv6 Solicit
- PPPoE PADI
- PPPoA Conf Req
- Router Solicit [RS]

Figure 131: LUDB authentication



(*) – indirect trigger only, via RADIUS fallback

25536

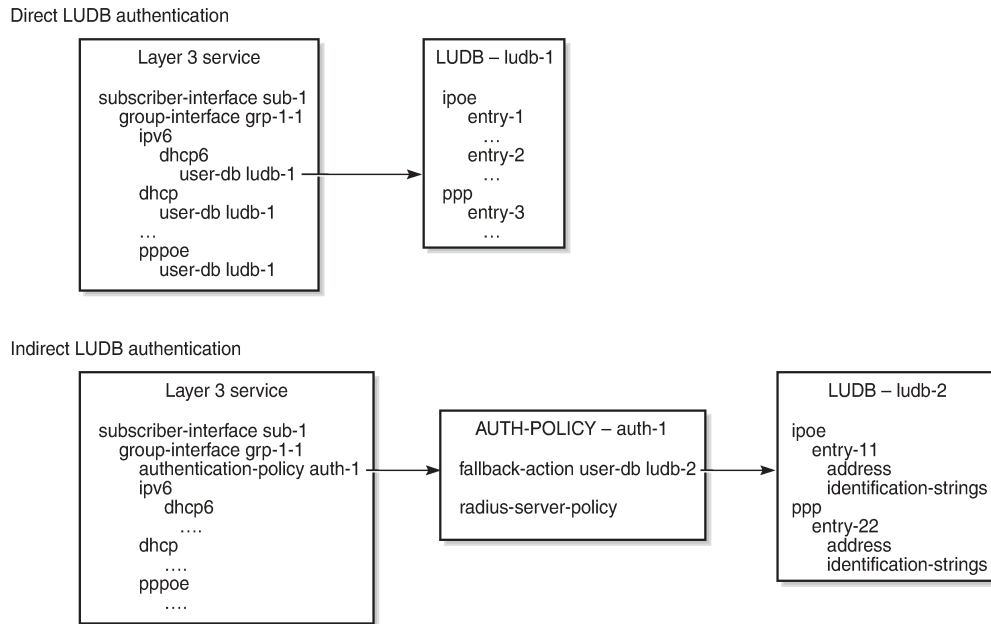
When triggered, ESM can directly access an LUDB because the LUDB is applied to the service directly (through one of its sub-contexts), or indirectly as a fallback action for RADIUS (through the authentication policy); see [Figure 132: Direct and indirect LUDB authentication](#). ARP requests can only trigger LUDB authentication indirectly.



Note:

An authentication policy can be referenced from a group interface and a capture SAP, or from an LUDB.

Figure 132: Direct and indirect LUDB authentication



25537

Three ESM scenarios in which an LUDB is accessed are as follows:

- ESM gets all the data needed for host creation directly from the LUDB.
- ESM gets some data from the LUDB and the remaining data from an AAA server (RADIUS, NASREQ, Gx). This requires the LUDB to provide an authentication or a Diameter application policy, and no authentication or Diameter policy at the group interface level.
- ESM tries to fetch the data from a RADIUS server, but because this server is not reachable, ESM falls back to an LUDB.

The examples in the [Configuration](#) section of this chapter describe the first and the last scenario. The second scenario is described in the [Flexible Authentication Model in ESM](#) chapter.

LUDB input and output parameters

As described in the [Local User Database Basics](#) chapter, when processing an LUDB lookup request, the input parameters are filtered and optionally masked before searching through the entries in the database. Every entry, except for the default, contains one or more host-identification fields that are used for matching purposes. As a result of the lookup process, these output parameters are then used for host creation.

The following IPoE host-identification fields are supported when accessing an LUDB for ESM; see [Figure 133: LUDB parameters for IPoE](#):

- mac
- circuit-id + remote-id
- option60 (excluded for the IPoE session model)
- sap-id + encap-tag-range

- service-id
- string
- system-id

The LUDB lookup process can take up to four IPoE match-criteria into account, as defined by the IPoE match-list.

The following PPP host-identification fields are supported when accessing an LUDB for ESM; see [Figure 134: LUDB parameters for PPPoE](#):

- mac
- circuit-id + remote-id
- sap-id + encap-tag-range
- service-name (PPPoE tag: service name)
- username (excluded for the RADIUS fallback scenario)

The LUDB lookup process can take up to three PPP match-criteria into account, as defined by the PPP match-list.

The fields output from the lookup process include the identification strings, options, and others.

See [Figure 133: LUDB parameters for IPoE](#) and [Figure 134: LUDB parameters for PPPoE](#) for the full list of input and output parameters for IPoE and PPPoE, respectively.

Figure 133: LUDB parameters for IPoE

Parameter		1	2	3	4
host-identification	circuit-id				
	derived-id				
	encap-tag-range				
	ip-prefix				
	mac				
	option60				
	remote-id				
	sap-id				
	service-id				
	string				
	system-id				
identification-strings	ancp-string				
	app-profile-string				
	category-map-name				
	inter-dest-id				
	sla-profile-string				
	spl-sharing-group-id				
	sub-profile-string				
	subscriber-id				
Parameter		1	2	3	4
acct-policy					
address	ip-address				
	gi-address				
	pool				
	use-from-pool-client				
-	auth-domain-name				
-	auth-policy				
-	ipv6-address				
-	ipv6-delegated-prefix				
-	ipv6-delegated-prefix-len				
-	ipv6-delegated-prefix-pool				
-	ipv6-slaac-prefix				
-	ipv6-wan-address-pool				
diameter-application-policy					
diameter-auth-policy					
gi-address					
link-address					
ipv6-lease-times	preferred-lifetime				
	rebind-timer				
	renew-timer				
	valid-lifetime				
ipv6-slaac-prefix-pool					
match-radius-proxy-cache					
msap-defaults	group-interface				
msap-defaults	policy				
msap-defaults	service				
options	custom-options				
	default-router				
	dns-server				
	domain-name				
	lease-rebind-time				
	lease-renew-time				
	lease-time				
	netbios-name-type				
options6	netbios-node-type				
	subnet-mask				
	dns-server				
-	retail-service-id				
rip-policy	server				
server6					
to-client-options	ipv4				
	ipv6				

	Supported
	Not Supported (ignored)
	Not Supported (error)
	Supported in proxy case, error in relay case
	Supported in proxy case, ignored in relay case
	Supported in relay case, ignored in proxy case

- 1 DHCPv4 proxy/relay
- 2 DHCPv6 proxy/relay
- 3 RADIUS fallback
- 4 DHCPv4 server

Figure 134: LUDB parameters for PPPoE

Parameter		1	2	3
host-identification	circuit-id			
	derived-id			
	encap-tag-range			
	mac			
	remote-id			
	sap-id			
	service-name			
	username			

Parameter		1	2	3
identification-strings	ancp-string			
	app-profile-string			
	category-map-name			
	inter-dest-id			
	sla-profile-string			
	spl-sharing-group-id			
	sub-profile-string			
	subscriber-id			

	Supported
	Not Supported (ignored)
	Not Supported (error)
	Supported in proxy case, error in relay case
	Supported in proxy case, ignored in relay case
	Supported in relay case, ignored in proxy case

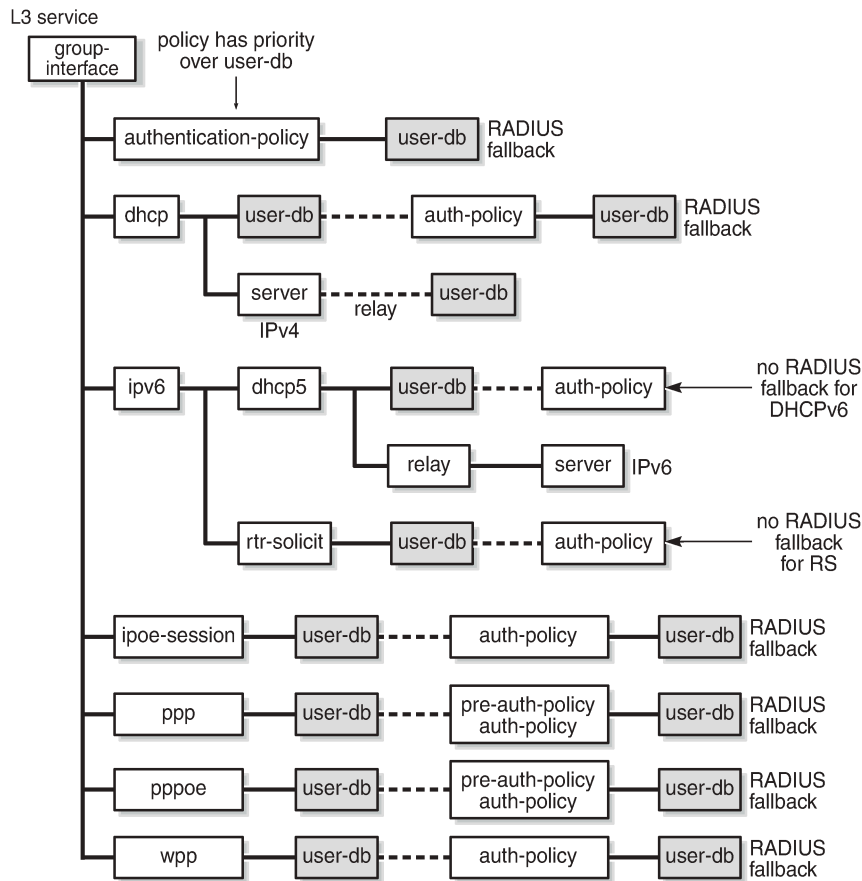
1 DHCPv4/6 PROXY / RELAY
2 RADIUS FALLBACK
3 DHCPv4 SERVER

Parameter		1	2	3
access-loop-encapsulation	encap-offset			
	rate-down			
access-loop-information	circuit-id			
	remote-id			
acct-policy				
address	ip-address			
	rt-address			
	pool			
	use-from-pool-client			
-	auth-policy			
	force-ipv6cp			
diameter-application-policy				
diameter-auth-policy				
ignore-df-bit				
ipv6-lease-times	preferred-lifetime			
	rebind-timer			
	renew-timer			
	valid-lifetime			
ipv6-slaac-prefix-pool				
interface	service-id			
-	ipv6-address			
-	ipv6-delegated-prefix			
-	ipv6-delegated-prefix-len			
-	ipv6-delegated-prefix-pool			
-	ipv6-slaac-prefix			
-	ipv6-wan-address-pool			
l2tp	group service-id			
msap-defaults	group-interface			
msap-defaults	policy			
msap-defaults	service			
options	custom-options			
	dns-server			
	netbios-name-server			
options6	dns-server			
	pad0-delay			
	password			
-	pre-auth-policy			
	retail-service-id			
rip-policy				
to-client-options	ipv4			
	ipv6			

Applying an LUDB for ESM

LUDB authentication for regular SAPs requires an LUDB to be applied at the group interface level in the Layer 3 service (VPRN or IES); see [Figure 135: LUDB authentication for regular SAPs](#). All the SAPs on that group interface share the same authentication configuration. See the [Local User Database for DHCPv4 Server](#) chapter for the scenario where a user database is attached to a DHCPv4 server.

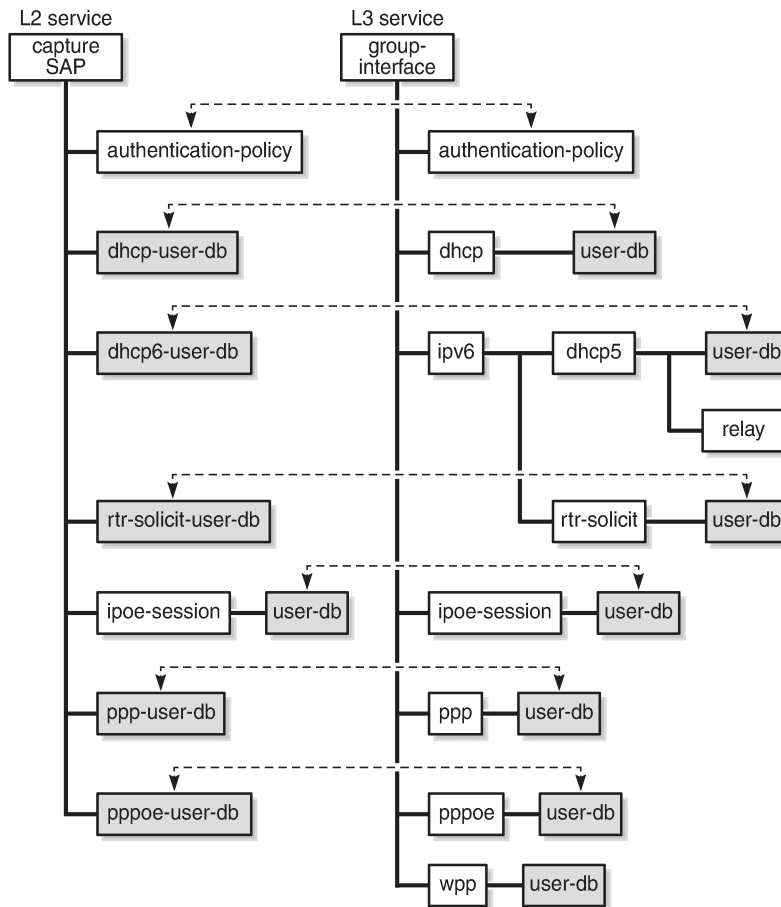
Figure 135: LUDB authentication for regular SAPs



25538

LUDB authentication for capture and managed SAPs requires an LUDB to be assigned at capture SAP level in the Layer 2 service (capture-VPLS), and at the group interface level in the Layer 3 service (VPRN or IES). Because the trigger messages to create the managed SAPs are received on the capture SAP and subsequent messages on the managed SAP, the authentication configurations for the Layer 2 and the Layer 3 service must align, including the LUDBs; see [Figure 136: LUDB authentication for capture and managed SAPs](#).

Figure 136: LUDB authentication for capture and managed SAPs



25539

The following CLI commands are available for applying LUDBs:

```

configure service vprn | ies subscriber-interface <x> group-interface <y>
  dhcp user-db <local-user-db-name>
  ipoe-session user-db <local-user-db-name>
  ipv6 dhcp6 user-db <local-user-db-name>
  ipv6 router-solicit user-db <local-user-db-name>
  ppp user-db <local-user-db-name>
  pppoe user-db <local-user-db-name>
  wpp user-db <local-user-db-name>
  
```

```

configure service vpls <x> sap <y>
  dhcp-user-db <local-user-db-name>
  dhcp6-user-db <local-user-db-name>
  ipoe-session user-db <local-user-db-name>
  ppp-user-db <local-user-db-name>
  pppoe-user-db <local-user-db-name>
  rtr-solicit-user-db <local-user-db-name>
  
```

An LUDB can be assigned in different contexts, and can be reused. Assuming an LUDB contains both IPoE as well as PPP entries, this LUDB is likely to be assigned in a dhcp context as well as in a ppp or a pppoe context.

Configuration guidelines

The following rules have to be observed when configuring authentication for regular, capture, and managed SAPs:

- If an authentication policy is applied at the capture SAP or group interface level, that authentication policy has priority, no matter whether or in which other sub-contexts an LUDB is assigned. Only when the AAA/RADIUS server referenced from the authentication policy is not available, can the SR OS rely on a fallback LUDB if configured. In that case, only a limited set of parameters are returned; see [Figure 133: LUDB parameters for IPoE](#) and [Figure 134: LUDB parameters for PPPoE](#).

This means that for an LUDB to provide ESM data, no authentication policy may be applied at the capture SAP or group interface level, provided that the LUDB is in the no shutdown state.

- An LUDB can return an authentication policy so that the ESM data can be partially provided by the LUDB, and partially by an AAA/RADIUS server. For this mixed scenario, RADIUS fallback is only possible for PPP, PPPoE, DHCPv4, IPoE sessions, and WPP, but not for DHCPv6 and IPv6 router solicitation. For more information, see the [Flexible Authentication Model in ESM](#) chapter. When the AAA/RADIUS server is defined but not available, the SR OS can rely on a fallback LUDB if configured.
- LUDB authentication for RADIUS fallback requires an LUDB to be applied to an authentication policy as a fallback action:

```
configure subscriber-mgmt authentication-policy <name>  
    fallback-action user-db <local-user-db-name>
```

- The DHCPv4 server referenced from a group interface in the dhcp context (for supporting the relay scenario) can have an LUDB assigned; see the [Local User Database for DHCPv4 Server](#) chapter. See [Figure 133: LUDB parameters for IPoE](#) and [Figure 134: LUDB parameters for PPPoE](#) for the parameters that this LUDB can return to the DHCPv4 server.

An LUDB cannot be assigned to a DHCPv6 server.

- If an LUDB is applied in the ipoe-session context of a group interface or capture SAP, the LUDBs assigned in the dhcp, dhcp6, and router-solicit contexts of the same group interface or capture SAP are ignored.

This avoids accessing the LUDB on every DHCPv4 DORA or DHCPv6 SARR message, which is the case when no IPoE sessions are used.

For IPoE sessions, the LUDB host identification cannot be based on option 60. Entries in the LUDB with host-identification option 60 strings are ignored. All the other LUDB entry match criteria are allowed.

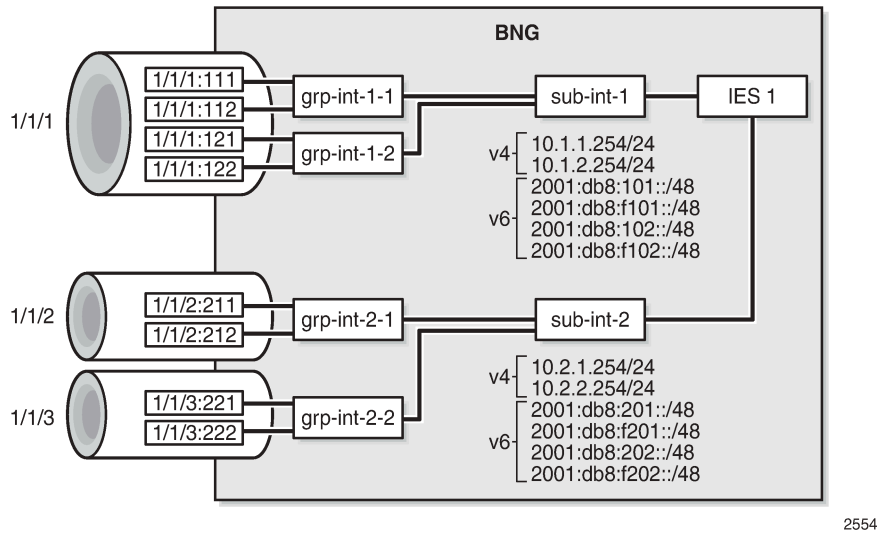
- If an LUDB is applied in the ppp or pppoe context of a group interface or capture SAP, PAP or CHAP authentication is based on the password configured in the entry. If no password is required, the password parameter in the LUDB entry must be explicitly set to ignore.

A password verification failure leads to a setup failure.

Configuration

Figure 137: Baseline setup shows the baseline configuration used in this chapter. Dual and single stack end-user devices supporting IPoE and PPPoE connect to the SAPs of IES-1. Different LUDBs are added to this baseline configuration later in this chapter, depending on the scenario.

Figure 137: Baseline setup



The following partial configuration applies to IES-1. This service is provisioned with ESM enabled on all of its SAPs, and supports proxy and relay scenarios on all group interfaces for both IPv4 and IPv6. Only the part relevant to subscriber interface *sub-int-1* and group interface *grp-int-1-1* is shown. The configurations for the other subscriber and group interfaces are similar. Check the [ESM Basics](#) and [Routed CO](#) chapters for more information.

```
configure
service
  ies 1 customer 1 create
    subscriber-interface "sub-int-1" create
      address 10.1.1.254/24
      ---snip---
      ipv6
        delegated-prefix-len 56
        link-local-address fe80::ea:4b:f1
        subscriber-prefixes
          prefix 2001:db8:101::/48 wan-host
          prefix 2001:db8:f101::/48 pd
          ---snip---
        exit
      exit
    group-interface "grp-int-1-1" create
      ipv6
        router-advertisements
          no shutdown
        exit
        dhcp6
          proxy-server
```



```

        client-applications dhcp ppp
        no shutdown
    exit
    relay
        link-address 2001:db8:101::1
        server 2001:db8::11
        client-applications dhcp ppp
        no shutdown
    exit
exit
router-solicit
    no shutdown
exit
exit
arp-populate
dhcp
    proxy-server
        emulated-server 10.1.1.254
        no shutdown
    exit
    option
        action keep
    exit
    server 10.11.11.1
    trusted
    lease-populate 100
    client-applications dhcp ppp
    gi-address 10.1.1.254
    no shutdown
exit
sap 1/1/1:111 create
    sub-sla-mgmt
        def-sub-profile "sub-prof-1"
        def-sla-profile "sla-prof-1"
        sub-ident-policy "sub-id-pol-1"
        multi-sub-sap
        no shutdown
    exit
exit
---snip---
exit
exit
---snip---
```

For brevity, the configurations of the local DHCPv4 and DHCPv6 servers are not shown.

An excerpt from the LUDB *ludb-rsap* follows. Host *entry-11* defines the settings for a dual stack IPoE host, and host *entry-55* the settings for a dual stack PPPoE host. For both hosts, the LUDB provides all the data needed to ensure host instantiation.

```

configure
    subscriber-mgmt
        local-user-db "ludb-rsap" create
        description "LUDB for Regular SAPs"
        ipoe
            match-list mac
            host "entry-11" create
            host-identification
                mac 00:00:00:11:11:11
            exit
            address 10.1.1.211
            identification-strings 254 create
            subscriber-id "sub-11"
```

```

        sla-profile-string "sla-prof-1"
        sub-profile-string "sub-prof-1"
    exit
    options
        subnet-mask 255.255.255.0
        default-router 10.1.1.254
        dns-server 2.2.2.2 2.2.2.1
        domain-name "domain.org"
        custom-option 251 hex 0x010203
    exit
    options6
        dns-server 2001:db8:ddd:1::1 2001:db8:ddd:2::1
    exit
    ipv6-address 2001:db8:102:11::11
    ipv6-delegated-prefix 2001:db8:f102:1100::/56
    ipv6-delegated-prefix-len 56
    no shutdown
exit
---snip---
exit
ppp
match-list username
host "entry-55" create
    host-identification
        username "sub55@domain1"
    exit
    address 10.1.1.225/24
    password chap letmein55
    identification-strings 254 create
        subscriber-id "sub-55"
        sla-profile-string "sla-prof-5"
        sub-profile-string "sub-prof-3"
    exit
    options
        dns-server 2.2.2.2
    exit
    options6
        dns-server 2001:db8:ddd:1::1 2001:db8:ddd:2::1
    exit
    ipv6-address 2001:db8:101:55::55
    ipv6-delegated-prefix 2001:db8:f101:5500::/56
    ipv6-delegated-prefix-len 56
    no shutdown
exit
---snip---
exit
no shutdown
exit

```

IPoE authentication - session model

In this example, the LUDB *ludb-rsap* is applied to the group interface in the ipoe-session context. This is the Nokia recommended way for supporting IPoE subscribers through an LUDB.

```

configure
service
    ies 1 customer 1 create
        subscriber-interface "sub-int-1"
        group-interface "grp-int-1-1"
        ipoe-session
            ipoe-session-policy "ipoe-sess-1"

```

```
session-limit 100
user-db "ludb-rsap"
no shutdown
exit
```

Use the following debug configuration for troubleshooting connection issues.

```
debug
router "Base"
ip
  dhcp
    detail-level low
    mode egr-ingr-and-dropped
  exit
  dhcp6
    mode egr-ingr-and-dropped
    detail-level low
  exit
exit
exit
subscriber-mgmt
  local-user-db "ludb-rsap"
  detail all
exit
exit
exit
```

The following trace appears when the user with MAC address 00:00:00:11:11:11 first connects using DHCPv4 and subsequently connects using DHCPv6 without removing the DHCPv4 connection. The LUDB is accessed just once, immediately after the DHCPv4 Discover message.

```
1 2018/11/22 12:45:34.750 CET MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 4 (grp-int-1-1),
  received DHCP Boot Request on Interface grp-int-1-1 (1/1/1:111) Port 67

H/W Type: Ethernet(10Mb)  H/W Address Length: 6
ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
siaddr: 0.0.0.0          giaddr: 0.0.0.0
chaddr: 00:00:00:11:11:11  xid: 0x1"

2 2018/11/22 12:45:34.750 CET MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
mac: 00:00:00:11:11:11

Host entry-11 found in user data base ludb-rsap"

3 2018/11/22 12:45:34.750 CET MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 4 (grp-int-1-1),
  transmitted DHCP Boot Reply to Interface grp-int-1-1 (1/1/1:111) Port 68

H/W Type: Ethernet(10Mb)  H/W Address Length: 6
ciaddr: 0.0.0.0          yiaddr: 10.1.1.211
siaddr: 10.1.1.254       giaddr: 10.1.1.254
chaddr: 00:00:00:11:11:11  xid: 0x1"

4 2018/11/22 12:45:34.772 CET MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
```

```
instance 1 (Base), interface index 4 (grp-int-1-1),
received DHCP Boot Request on Interface grp-int-1-1 (1/1/1:111) Port 67

H/W Type: Ethernet(10Mb) H/W Address Length: 6
ciaddr: 0.0.0.0 yiaddr: 0.0.0.0
siaddr: 0.0.0.0 giaddr: 0.0.0.0
chaddr: 00:00:00:11:11:11 xid: 0x1"

5 2018/11/22 12:45:34.774 CET MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 4 (grp-int-1-1),
transmitted DHCP Boot Reply to Interface grp-int-1-1 (1/1/1:111) Port 68

H/W Type: Ethernet(10Mb) H/W Address Length: 6
ciaddr: 0.0.0.0 yiaddr: 10.1.1.211
siaddr: 10.1.1.254 giaddr: 10.1.1.254
chaddr: 00:00:00:11:11:11 xid: 0x1"

6 2018/11/22 12:46:00.160 CET MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
Incoming DHCP6 Msg : SOLICIT (1)
on itf grp-int-1-1"

7 2018/11/22 12:46:00.160 CET MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
Outgoing DHCP6 Msg : ADVERTISE (2)
to itf grp-int-1-1"

8 2018/11/22 12:46:00.179 CET MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
Incoming DHCP6 Msg : REQUEST (3)
on itf grp-int-1-1"

9 2018/11/22 12:46:00.180 CET MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
Outgoing DHCP6 Msg : REPLY (7)
to itf grp-int-1-1"
```

The active subscriber hosts for service 1 are shown with the following command.

```
*A:BNG-1# show service id 1 subscriber-hosts

=====
Subscriber Host table
=====
```

Sap	Subscriber	IP Address	MAC Address	PPPoE-SID Origin	Fwding State
1/1/1:111	sub-11	10.1.1.211	00:00:00:11:11:11	N/A	DHCP
1/1/1:111	sub-11	2001:db8:102:11::11/128	00:00:00:11:11:11	N/A	IPoE-DHCP6
1/1/1:111	sub-11	2001:db8:f102:1100::/56	00:00:00:11:11:11	N/A	IPoE-DHCP6

```
-----
Number of subscriber hosts : 3
=====
*A:BNG-1#
```

The following command shows the session details for MAC address 00:00:00:11:11:11. This information aligns with the LUDB configuration of *ludb-rsap*, and the origin codes are set to *UserDb*.

```
*A:BNG-1# show service id 1 ipoe session mac 00:00:00:11:11:11 detail

=====
IPoE sessions for service 1
=====

SAP                          : 1/1/1:111
Mac Address                  : 00:00:00:11:11:11
Circuit-Id                   : 11
Remote-Id                    : AA
Session Key                   : sap-mac

MC-Standby                   : No

Subscriber-interface         : sub-int-1
Group-interface              : grp-int-1-1

Termination Type             : local
Up Time                      : 0d 00:01:25
Session Time Left            : N/A
Last Auth Time               : 11/22/2018 12:45:35
Min Auth Intvl (left)       : infinite (N/A)
Persistence Key              : N/A

Subscriber                   : "sub-11"
Sub-Profile-String           : "sub-profile-1"
SLA-Profile-String           : "sla-profile-1"
SPI group ID                 : (Not Specified)
ANCP-String                  : ""
Int-Dest-Id                  : ""
App-Profile-String           : ""
Category-Map-Name            : ""
Acct-Session-Id              : "0217FF000000315BF696DE"
Sap-Session-Index            : 1

IP Address                   : 10.1.1.211/24
IP Origin                     : UserDb
Primary DNS                   : 2.2.2.2
Secondary DNS                 : 2.2.2.1
Primary NBNS                  : N/A
Secondary NBNS                : N/A
Address-Pool                  : N/A

IPv6 Prefix                  : N/A
IPv6 Prefix Origin            : None
IPv6 Prefix Pool              : ""
IPv6 Del.Pfx.                : 2001:db8:f102:1100::/56
IPv6 Del.Pfx. Origin          : UserDb
IPv6 Del.Pfx. Pool            : ""
IPv6 Address                  : 2001:db8:102:11::11
IPv6 Address Origin            : UserDb
IPv6 Address Pool              : ""
Primary IPv6 DNS              : 2001:db8:ddd:1::1
Secondary IPv6 DNS            : 2001:db8:ddd:2::1
Router adv. policy            : N/A
```

```

Radius sub-if prefix      : N/A
Radius Session-T0        : N/A
Radius Class              :
Radius User-Name         :

GTP IMSI                 :
GTP APN                   : (Not Specified)
-----
Number of sessions : 1
=====
*A:BNG-1#

```

The commands for showing the IPv4 and IPv6 lease states display the lease origin codes too, as follows:

```

*A:BNG-1# show service id 1 dhcp lease-state mac 00:00:00:11:11:11

=====
DHCP lease state table, service 1
=====
IP Address      Mac Address      Sap/Sdp Id      Remaining   Lease   MC
                  LeaseTime   Origin          Stdby
-----
10.1.1.211      00:00:00:11:11:11  1/1/1:111      06d23h57m  UserDb
-----
Number of lease states : 1
=====
*A:BNG-1#

```

```

*A:BNG-1# show service id 1 dhcp6 lease-state mac 00:00:00:11:11:11

=====
DHCP lease state table, service 1
=====
IP Address      Mac Address      Sap/Sdp Id      Remaining   Lease   MC
                  LeaseTime   Origin          Stdby
-----
2001:db8:102:11::11/128
                  00:00:00:11:11:11  1/1/1:111      23h57m48s  UserDb
2001:db8:f102:1100::/56
                  00:00:00:11:11:11  1/1/1:111      23h57m48s  UserDb
-----
Number of lease states : 2
=====
*A:BNG-1#

```

IPoE authentication - host model

In this example, the LUDB *ludb-rsap* is applied to the group interface in the dhcp6, router-solicit, and dhcp contexts, but not in the ipoe-session context.

```

configure
  service
    ies 1
      subscriber-interface "sub-int-1"
      group-interface "grp-int-1-1"
      ipv6
        dhcp6
          user-db "ludb-rsap"

```

```

        exit
        router-solicit
        user-db "ludb-rsap"
        no shutdown
    exit
exit
dhcp
    user-db "ludb-rsap"
    no shutdown
exit
exit

```

With the same debug configuration as for the IPoE session model, the LUDB is accessed multiple times when devices connect, as shown in the following trace.

```

13 2018/11/22 12:50:55.275 CET MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 4 (grp-int-1-1),
    received DHCP Boot Request on Interface grp-int-1-1 (1/1/1:111) Port 67

H/W Type: Ethernet(10Mb)  H/W Address Length: 6
ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
siaddr: 0.0.0.0           giaddr: 0.0.0.0
chaddr: 00:00:00:11:11:11  xid: 0x1"

14 2018/11/22 12:50:55.275 CET MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
mac: 00:00:00:11:11:11

Host entry-11 found in user data base ludb-rsap"

15 2018/11/22 12:50:55.275 CET MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 4 (grp-int-1-1),
    transmitted DHCP Boot Reply to Interface grp-int-1-1 (1/1/1:111) Port 68

H/W Type: Ethernet(10Mb)  H/W Address Length: 6
ciaddr: 0.0.0.0           yiaddr: 10.1.1.211
siaddr: 10.1.1.254        giaddr: 10.1.1.254
chaddr: 00:00:00:11:11:11  xid: 0x1"

16 2018/11/22 12:50:55.286 CET MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 4 (grp-int-1-1),
    received DHCP Boot Request on Interface grp-int-1-1 (1/1/1:111) Port 67

H/W Type: Ethernet(10Mb)  H/W Address Length: 6
ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
siaddr: 0.0.0.0           giaddr: 0.0.0.0
chaddr: 00:00:00:11:11:11  xid: 0x1"

17 2018/11/22 12:50:55.286 CET MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
mac: 00:00:00:11:11:11

Host entry-11 found in user data base ludb-rsap"

18 2018/11/22 12:50:55.288 CET MINOR: DEBUG #2001 Base PIP

```

```
"PIP: DHCP
instance 1 (Base), interface index 4 (grp-int-1-1),
transmitted DHCP Boot Reply to Interface grp-int-1-1 (1/1/1:111) Port 68
H/W Type: Ethernet(10Mb) H/W Address Length: 6
ciaddr: 0.0.0.0 yiaddr: 10.1.1.211
siaddr: 10.1.1.254 giaddr: 10.1.1.254
chaddr: 00:00:00:11:11:11 xid: 0x1"

19 2018/11/22 12:51:20.248 CET MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
Incoming DHCP6 Msg : SOLICIT (1)
on itf grp-int-1-1"

20 2018/11/22 12:51:20.248 CET MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
mac: 00:00:00:11:11:11

Host entry-11 found in user data base ladb-rsap"

21 2018/11/22 12:51:20.248 CET MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
Outgoing DHCP6 Msg : ADVERTISE (2)
to itf grp-int-1-1"

22 2018/11/22 12:51:20.261 CET MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
Incoming DHCP6 Msg : REQUEST (3)
on itf grp-int-1-1"

23 2018/11/22 12:51:20.261 CET MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
mac: 00:00:00:11:11:11

Host entry-11 found in user data base ladb-rsap"

24 2018/11/22 12:51:20.262 CET MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
Outgoing DHCP6 Msg : REPLY (7)
to itf grp-int-1-1"
```

The LUDB is accessed for every incoming message. In a proxy case, the LUDB is accessed two times per host because the downstream messages (Offer and Reply for IPv4, Solicit and Reply for IPv6) are generated by ESM. In a relay case, where an IP address or an IP prefix is allocated by the DHCP server, the LUDB is accessed four times per host.

The command to list the active subscriber hosts is the same as for the IPoE session model, and is not repeated here. The same applies to the other commands providing origin codes.

PPPoE authentication

In this example, the LUDB *ladb-rsap* is applied to the group interface in the pppoe context.

```
configure
service
```



```

    ies 1
      subscriber-interface "sub-int-1"
        group-interface "grp-int-1-1"
          pppoe
            user-db "ludb-rsap"
            no shutdown
          exit
        exit
      exit

```

The following debug configuration applies for this example.

```

debug
  service
    id 1
      ppp
        packet
          mode egr-ingr-and-dropped
          detail-level high
          discovery
          ppp
            dhcp-client
          exit
        exit
      exit
    exit
  subscriber-mgmt
    local-user-db "ludb-rsap"
    detail all
  exit
exit

```

The trace shows that the LUDB *ludb-rsap* is accessed once when user *sub55@domain1* connects. In this example, the LUDB is accessed in the middle of the CHAP authentication.

```

---snip---

37 2018/11/22 12:52:59.419 CET MINOR: DEBUG #2001 Base PPPoE
"PPPoE: TX Packet
  IES 1, SAP 1/1/1:111

  DMAC: 00:00:00:55:55:55
  SMAC: 02:17:01:01:00:01
  Ether Type: 0x8864 (Session)

  PPPoE Header:
  Version: 1                      Type      : 1
  Code   : 0x00                  Session-Id: 0x0001 (1)
  Length : 50

  PPP:
  Protocol : 0xc223 (CHAP)
  Code     : 1 (Challenge)
  Identifier: 1                      Length   : 48

  Value-Size: 38
  Value      : 3c d2 f7 9c 6b d5 9d 12 0e d7 96 8e ac d8 61 b5 e2 d2 8c 06 8a
  8b 50 b3 10 f4 d3 81 80 f8 ca 3d 4b 42 d9 b6 98 78
  Name       : "BNG-1"

  Hex Packet Dump:
  11 00 00 01 00 32 c2 23 01 01 00 30 26 3c d2 f7 9c 6b d5 9d 12 0e d7 96 8e

```



```
---snip---
```

With this dual stack PPP user connected, the subscriber hosts created are:

```
*A:BNG-1# show service id 1 subscriber-hosts

=====
Subscriber Host table
=====
Sap      Subscriber
IP Address
MAC Address      PPPoE-SID Origin      Fwding State
-----
1/1/1:111      sub-55
10.1.1.225
00:00:00:55:55:55      1      IPCP      Fwding
1/1/1:111      sub-55
2001:db8:101:55::55/128
00:00:00:55:55:55      1      PPP-DHCP6      Fwding
1/1/1:111      sub-55
2001:db8:f101:5500::/56
00:00:00:55:55:55      1      PPP-DHCP6      Fwding
-----
Number of subscriber hosts : 3
=====
*A:BNG-1#
```

Detailed session information for PPP user *sub55@domain55* shows the origin codes.

```
*A:BNG-1# show service id 1 ppp session user-name "sub55@domain1" detail

=====
PPP sessions for service 1
=====
User-Name      : sub55@domain1

Description     : svc:1 sap:1/1/1:111 mac:00:00:00:55:55:55 sid:1
Up Time        : 0d 00:01:00
Type           : oE
Termination    : local
IP/L2TP-Id/If-Id : 10.1.1.225 02:00:00:FF:FE:55:55:55
MC-Standby     : No
Session Time Left : N/A

LCP State      : Opened
IPCP State     : Opened
IPv6CP State   : Opened
PPP MTU        : 1492
PPP Auth-Protocol : CHAP
PPP User-Name  : sub55@domain1

Subscriber-interface : sub-int-1
Group-interface    : grp-int-1-1

IP Origin       : local-user-db
DNS Origin      : local-user-db
NBNS Origin     : none

Subscriber      : "sub-55"
Sub-Profile-String : "sub-prof-3"
SLA-Profile-String : "sla-prof-5"
SPI group ID    : (Not Specified)
```

```

---snip---

IP Address      : 10.1.1.225/32
Primary DNS     : 2.2.2.2
Secondary DNS   : N/A
Primary NBNS    : N/A
Secondary NBNS  : N/A
Address-Pool    : N/A
IPv6 Prefix     : N/A
IPv6 Prefix Origin : none
IPv6 Prefix Pool : ""
IPv6 Del.Pfx.   : 2001:db8:f101:5500::/56
IPv6 Del.Pfx. Origin : local-user-db
IPv6 Del.Pfx. Pool : ""
IPv6 Address    : 2001:db8:101:55::55
IPv6 Address Origin : local-user-db
IPv6 Address Pool : ""
Primary IPv6 DNS : 2001:db8:ddd:1::1
Secondary IPv6 DNS : 2001:db8:ddd:2::1
Router adv. policy : N/A

---snip---

-----
No. of sessions: 1
=====
*A:BNG-1#

```

The following command shows the lease origin.

```

*A:BNG-1# show service id 1 dhcp6 lease-state session ppp

=====
DHCP lease state table, service 1
=====

```

IP Address	Mac Address	Sap/Sdp Id	Remaining LeaseTime	Lease Origin	MC Stdby
2001:db8:101:55::55/128	00:00:00:55:55:55	1/1/1:111	23h58m07s	UserDb	
2001:db8:f101:5500::/56	00:00:00:55:55:55	1/1/1:111	23h58m07s	UserDb	

```

-----
Number of lease states : 2
=====
*A:BNG-1#

```

Regular SAPs versus capture and managed SAPs

When an LUDB is to be used for regular SAPs, the LUDB must be assigned at the group interface level of a Layer 3 service (IES or VPRN). This LUDB is then used for all SAPs on that group interface, as described in the section [Applying an LUDB for ESM](#).

When an LUDB is to be used for capture and managed SAPs, the LUDB must be assigned at the capture SAPs of the Layer 2 (VPLS) service and at the group interface level of the corresponding Layer 3 service (IES or VPRN).

Because the managed SAPs are dynamically created at the group interface of a Layer 3 service, this service must have its authentication configuration aligned with the Layer 2 service; see [Figure 136: LUDB authentication for capture and managed SAPs](#).

Capture and managed SAPs support IPoE (session and host model) and PPP.

The capture VPLS is defined as follows.

```
configure
service
  vpls 3 customer 1 create
  stp
    shutdown
  exit
  sap 1/1/2:* capture-sap create
    trigger-packet arp dhcp dhcp6 pppoe rtr-solicit
    dhcp-user-db "ludb-cmsap"
    pppoe-user-db "ludb-cmsap"
    ipoe-session
      ipoe-session-policy "ipoe-sess-1"
      user-db "ludb-cmsap"
      no shutdown
    exit
    msap-defaults
      group-interface "grp-int-1-1"
      policy "msap-pol-1"
      service 2
    exit
  exit
  no shutdown
exit
```

The VPRN on which the managed SAPs are created is defined as follows.

```
configure
service
  vprn 2 customer 1 create
  ---snip---
  subscriber-interface "sub-int-1" create
    address 10.111.1.254/24
    ipv6
      delegated-prefix-len 56
      subscriber-prefixes
        prefix 2001:db8:901::/48 wan-host
        prefix 2001:db8:f901::/48 pd
    exit
  exit
  group-interface "grp-int-1-1" create
  ---snip---
  ipoe-session
    ipoe-session-policy "ipoe-sess-1"
    sap-session-limit 100
    user-db "ludb-cmsap"
    no shutdown
  exit
  oper-up-while-empty
  pppoe
    session-limit 100
    user-db "ludb-cmsap"
    no shutdown
  exit
  exit
exit
```

```
exit
```

The msap-defaults needed for creation of the managed SAPs can be taken from the capture SAP, but can also be obtained from an LUDB, as the following example shows. In that case, they overrule the capture SAP msap-defaults.

```
configure
subscriber-mgmt
local-user-db "ludb-cmsap" create
description "LUDB for capture/managed SAPs"
ipoe
match-list mac
host "entry-1" create
host-identification
mac 00:00:00:01:01:01
exit
address 10.111.1.101
identification-strings 254 create
subscriber-id "sub-priv-1"
sla-profile-string "sla-prof-3"
sub-profile-string "sub-prof-4"
exit
msap-defaults
group-interface "grp-int-1-1"
policy "msap-pol-1"
service 2
exit
options
subnet-mask 255.255.255.0
exit
ipv6-address 2001:db8:901:11::11
ipv6-delegated-prefix 2001:db8:f901:1100::/56
ipv6-delegated-prefix-len 56
no shutdown
exit
exit
ppp
match-list mac
host "entry-1" create
host-identification
mac 00:00:00:05:05:05
exit
address 10.111.1.105/32
identification-strings 254 create
subscriber-id "sub-05"
sla-profile-string "sla-prof-2"
sub-profile-string "sub-prof-4"
exit
msap-defaults
group-interface "grp-int-1-1"
policy "msap-pol-1"
service 2
exit
ipv6-address 2001:db8:901:5::5
ipv6-delegated-prefix 2001:db8:f901:500::/56
ipv6-delegated-prefix-len 56
no shutdown
exit
exit
```

Detailed information on managed and capture SAPs is in the [Managed SAPs with Routed CO](#) chapter.

The commands to display the subscribers, lease, and session states with the origin codes are the same as in the section [PPPoE authentication](#), so these are not repeated.

LUDB for ESM as RADIUS fallback

RADIUS fallback can be triggered in the following situations; see also [Figure 135: LUDB authentication for regular SAPs](#) and [Figure 136: LUDB authentication for capture and managed SAPs](#):

- with the authentication policy directly assigned at the group interface level
- with the authentication policy referenced from an LUDB

For the second case, first-level authentication is performed by the LUDB, and second-level authentication should be performed by the RADIUS server. For both cases, when the RADIUS server is not reachable, fallback happens.



Note:

RADIUS fallback is not supported when the LUDB is attached to the group interface or capture SAP via the `ipv6 dhcp6` and `rtr-solicit` contexts.

Although RADIUS fallback applies to both IPoE and PPP, only IPoE is shown in the example that follows.

To demonstrate the use of an LUDB for RADIUS fallback, the configuration of the previous example with capture and managed SAPs is modified, as follows.

```
# the (capture-)VPLS
configure
  service
    vpls 3 customer 1 create
    sap 1/1/2:* capture-sap create
    authentication-policy "auth-pol-1"
    exit
  exit
exit
exit
```

```
# the VPRN
configure
  service
    vprn 2 customer 1 create
    subscriber-interface "sub-int-1"
    group-interface "grp-int-1-1"
    authentication-policy "auth-pol-1"
    exit
  exit
exit
exit
```

The authentication policy is applied in the VPLS at the SAP level, and in the VPRN at the group interface level. Even with LUDBs assigned in other contexts at that group interface, the authentication policy takes higher priority.

The LUDB used for RADIUS fallback is defined as follows, and both the `ipoe` and the `ppp` sections contain a default host entry.

```
configure
  subscriber-mgmt
    local-user-db "ludb-radiusfb" create
```

```
description "LUDB for RADIUS fallback"
ipoe
  match-list mac
  host "default" create
  msap-defaults
    group-interface "grp-int-1-1"
    policy "msap-pol-1"
    service 2
  exit
  no shutdown
exit
ppp
  match-list username
  host "default" create
  msap-defaults
    group-interface "grp-int-1-1"
    policy "msap-pol-1"
    service 2
  exit
  no shutdown
exit
no shutdown
exit
```

The authentication policy from which this LUDB is referenced is defined as follows.

```
configure
subscriber-mgmt
  authentication-policy "auth-pol-1" create
  fallback-action user-db "ludb-radiusfb"
  radius-server-policy "rsp-1"
exit
```

The definition of the RADIUS server policy is not relevant so it is not shown.

The following debug configuration applies.

```
debug
  router "Base"
    radius
      packet-type authentication accounting coa
      detail-level high
    exit
  exit
  router "2"
    ip
      dhcp
        detail-level medium
        mode egr-ingr-and-dropped
      exit
    exit
  exit
  service
    id 3
      dhcp
        mode egr-ingr-and-dropped
      exit
    exit
  exit
subscriber-mgmt
```



```
        local-user-db "ludb-radiusfb"
        detail all
    exit
exit
exit
```

The following partial debug output shows that when a DHCPv4 user connects, the LUDB *ludb-radiusfb* is accessed after failing to connect to the RADIUS server. Similar debug output appears when connecting through DHCPv6 via IPE sessions, or PPP.

```
62 2018/11/22 13:03:15.510 CET MINOR: DEBUG #2001 Base SVCNMR
"SVCNMR: RX DHCP Packet
  VPLS 3, SAP 1/1/2:*

  BootRequest to UDP port 67
  ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
  siaddr: 0.0.0.0          giaddr: 0.0.0.0
  chaddr: 00:00:00:01:01:01  xid: 0x3

  DHCP options:
  [82] Relay agent information: len = 8
    [1] Circuit-id: 11
    [2] Remote-id: AA
  [53] Message type: Discover
  [255] End"

63 2018/11/22 13:03:15.510 CET MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  server 192.168.66.66:1812 not reachable"

64 2018/11/22 13:03:15.510 CET MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Access-Request
  user 00:00:00:01:01:01  policy rsp-1
  send failed"

65 2018/11/22 13:03:15.510 CET MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
  mac: 00:00:00:01:01:01

  Host default found in user data base ludb-radiusfb"

66 2018/11/22 13:03:15.513 CET MINOR: DEBUG #2001 vprn2 PIP
"PIP: DHCP
instance 2 (2), interface index 10 (grp-int-1-1),
received DHCP Boot Request on Interface grp-int-1-1 (1/1/2:123) Port 67
H/W Type: Ethernet(10Mb)  H/W Address Length: 6
ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
siaddr: 0.0.0.0          giaddr: 0.0.0.0
chaddr: 00:00:00:01:01:01  xid: 0x3
DHCP options:
[82] Relay agent information: len = 8
  [1] Circuit-id: 11
  [2] Remote-id: AA
[53] Message type: Discover
[255] End"

---snip---
```

```
72 2018/11/22 13:03:15.524 CET MINOR: DEBUG #2001 vprn2 PIP
"PIP: DHCP
instance 2 (2),
received DHCP Boot Reply on 10.111.111.1 Port 67

H/W Type: Ethernet(10Mb) H/W Address Length: 6
ciaddr: 0.0.0.0 yiaddr: 10.111.1.3
siaddr: 10.111.111.1 giaddr: 10.111.1.254
chaddr: 00:00:00:01:01:01 xid: 0x3

DHCP options:
[82] Relay agent information: len = 8
[1] Circuit-id: 11
[2] Remote-id: AA
[53] Message type: Ack
[54] DHCP server addr: 10.111.111.1
[51] Lease time: 864000
[1] Subnet mask: 255.255.255.0
[255] End"

73 2018/11/22 13:03:15.525 CET MINOR: DEBUG #2001 vprn2 PIP
"PIP: DHCP
instance 2 (2), interface index 10 (grp-int-1-1),
transmitted DHCP Boot Reply to Interface grp-int-1-1 (1/1/2:123) Port 68

H/W Type: Ethernet(10Mb) H/W Address Length: 6
ciaddr: 0.0.0.0 yiaddr: 10.111.1.3
siaddr: 10.111.111.1 giaddr: 10.111.1.254
chaddr: 00:00:00:01:01:01 xid: 0x3
DHCP options:
[82] Relay agent information: len = 8
[1] Circuit-id: 11
[2] Remote-id: AA
[53] Message type: Ack
[54] DHCP server addr: 10.111.111.1
[51] Lease time: 864000
[1] Subnet mask: 255.255.255.0
[255] End"
```

In this example, the LUDB accessed (on RADIUS fallback) defines a default host for ipoe as well as for ppp with msap-defaults only, which means relaying applies where the DHCPv4 and DHCPv6 servers provide the IP addresses and prefixes.

See [Figure 133: LUDB parameters for IPoE](#) and [Figure 134: LUDB parameters for PPPoE](#) for the list of supported parameters for IPoE and PPP in the RADIUS fallback scenario.

Operational considerations and remarks

The operational considerations listed in the [Local User Database Basics](#) chapter still apply.

To maintain backward compatibility with previous software releases, LUDB informational and error messages are sent to the error logs as if they are originating from the DHCPS application (DHCPS #xyz in the preceding outputs).

Conclusion

LUDBs offer a self-contained method of providing ESM data locally stored on the router, so that no external database is needed for supporting authentication. In case authentication relies on an AAA/RADIUS server that fails, an LUDB can provide the ESM data instead through RADIUS fallback. LUDBs can be used on regular, managed, and capture SAPs.

Managed SAPs with Routed CO

This chapter provides information about Managed SAPs with Routed CO.

Topics in this chapter include:

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

Applicability

This chapter is applicable to SR OS routers and was initially written for Release 8.0.R1. The CLI in this edition corresponds to Release 15.0.R2.

Overview

Managed SAPs (MSAP) are SAPs dynamically created after the reception of a trigger packet on a capture SAP. The creation of the MSAP is controlled through an MSAP policy, which is defined during the authentication phase along with the subscriber host parameters required for host instantiation.

Following trigger packet types can lead to the creation of an MSAP:

- arp
- dhcp
- dhcp6
- rtr-solicit
- pppoe
- ppp
- data

Multiple trigger packet types can be enabled for a single capture SAP.

MSAP creation takes several steps:

- Reception of a trigger packet on the capture SAP.
- Authentication, for example via RADIUS, LUDB, NASREQ, etc.
Authentication provides the MSAP policy and the target service context required in the next step.
- The actual creation of the MSAP in the service defined during the authentication step, taking the MSAP policy into account.

MSAPs are supported in the Bridged Central Office model and the Routed Central Office (RCO) model. For the bridged model, the service context returned by authentication is the service ID of a VPLS. For the

routed model, the service context is the service ID of a routed service (IES or VPRN) plus the name of a group-interface in the target service. Only the RCO model is explained in this chapter.

The capture SAP receives trigger packets and initiates authentication. The capture SAP is defined in a VPLS, and does not forward traffic.

The MSAP is created in the target service, and the VLAN of the MSAP is the same as the VLAN of the trigger packet. The MSAP behaves as a regular SAP, but its configuration is not user editable and not maintained in the configuration file. The MSAP remains active as long as the session is active. MSAPs and regular SAPs can co-exist on the same port and in the same service.

MSAPs can be created in a wholesale VPRN service while the corresponding subscriber host or session is terminated in a retail VPRN or IES service. Both wholesale MSAP data (service, group-interface, and policy) and retail service id must be provided during authentication.

Knowledge of TPSDA (Triple Play Service Delivery Architecture) and functionality is assumed throughout this chapter.

Capture SAP

The IOM classifies traffic based on the tags present in the incoming packets, and sends traffic to existing SAPs if the tag or tag combination in the incoming packet is known to the IOM.

The capture SAP is used if a more specific match for the Q or Q-in-Q tags is not found by the traffic classification on the IOM.

Trigger packets received on the capture SAP are sent to the CPM, non-trigger packets received on the capture SAP are dropped.

Following formats are allowed on the capture SAP:

SAP 1/2/2:*	for dot1Q
SAP 1/2/2:.*	for QinQ
SAP 1/2/2:Q1.*	for QinQ
SAP 1/2/2:*.Q1	for QinQ (inverse capture SAP)

By default, the MSAP created will have one q-tag (for dot1q) or two q-tags (for qinq), and these are taken from the original trigger packet. The optional **allow-dot1q-msaps** command additionally enables single tagged trigger packet support for QinQ capture SAPs. See the user manual for a full description.

MSAP with Redundant Configurations

MSAPs are High Availability (HA) enabled (there is no service impact following a CPM failover). In addition, the MSAPs are also stored in the subscriber management persistence file (if enabled), allowing the MSAPs to be recreated after a reboot.

MSAPs can be used in dual-homed BNG scenarios with multi-chassis LAG, multi-chassis ring and subscriber router redundancy protocol.

RADIUS Authentication and Vendor Specific Attributes (VSAs) for MSAP

The Alc-MSAP-Serv-Id attribute returned by the RADIUS server defines the service in which the MSAP must be created.

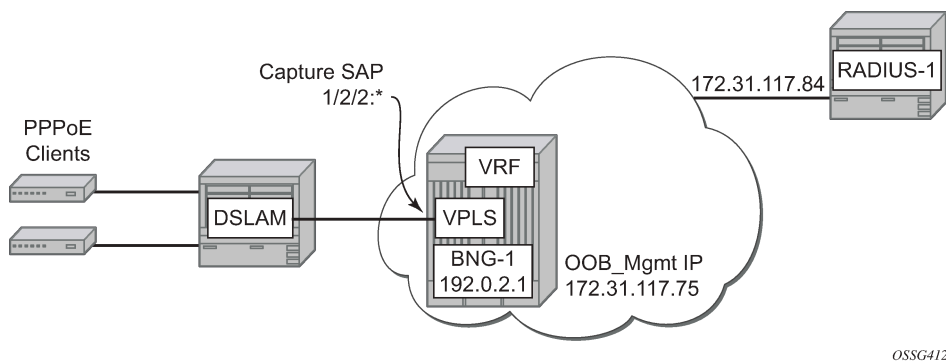
In the RCO scenario, the MSAP is created in a group-interface context. The Alc-MSAP-Interface attribute returned by the RADIUS server defines the group-interface where the MSAP must be installed, and must exist in the target service.

The Alc-MSAP-Policy attribute returned by the RADIUS server defines the MSAP parameters required for creating the MSAP.

Topology

The network topology is displayed in [Figure 138: Network Topology](#). This chapter uses the RCO model with PPPoE, IPv4, and RADIUS authentication for demonstrating MSAPs.

Figure 138: Network Topology



Configuration

RADIUS

In this chapter the management router is used for RADIUS communication, and the configuration used is as follows:

```
configure
  router "management"
    radius-server
      server "radius-138.203.10.250" address 172.31.117.84
      secret vsecret1 create
      description "Management router is used for RADIUS"
      accept-coa
    exit
  exit
exit
```

```
exit
```

```
configure
  aaa
    radius-server-policy "rad-serv-pol-1" create
      servers
        router "management"
        source-address 172.31.117.75
        server 1 name "radius-172.31.117.84"
      exit
    exit
  exit
exit
```

```
configure
  subscriber-mgmt
    authentication-policy "authentication-1" create
      description "RADIUS authentication policy"
      password "letmein"
      pppoe-access-method pap-chap
      include-radius-attribute
        remote-id
        nas-identifier
        mac-address
      exit
    radius-server-policy "rad-serv-pol-1"
  exit
exit
exit
```

The value of the secret is defined as *vsecret1*. The secret is a case sensitive character string of 20 characters maximum, which must be configured in the clients.conf file on the RADIUS server.

The management routing instance with the out-of-band 172.31.117.75 IP address is used as the source to communicate authentication messages between the BNG and the RADIUS server. The RADIUS server IP address is 172.31.117.84. Up to sixteen servers can be configured in the RADIUS server policy. When multiple servers are defined, the access algorithm can be set to **direct**, or **round-robin**.

The authentication method used in this example is PAP/CHAP, so the pap-chap value is used for the pppoe-access-method.

The user's remote-id and mac-address are included with the nas-identifier into the access request message sent to the RADIUS.

QoS SAP Policies

The following QoS SAP ingress and egress policies are used later in this chapter. The dot1p and dscp values used are examples:

```
configure
  qos
    sap-ingress 20 create
      description "64K_upstream"
      queue 1 create
        rate 64
      exit
    queue 11 multipoint create
  exit
```

```
exit
---snip---
sap-egress 50 create
  description "2M_downstream"
  queue 1 create
    rate 2048
  exit
  fc be create
    queue 1
    dot1p 3
    dscp cs1
  exit
exit
exit
exit
```

Enhanced Subscriber Management Parameters

SLA profiles are configured where the downstream speed is four times the upstream speed and the SLA profile will be named with the downstream speed. A subscriber profile is configured to initiate RADIUS accounting. A subscriber identification profile is configured for direct mapping subscriber and SLA profiles, as follows:

```
configure
  subscriber-mgmt
    sla-profile "sla-profile-1M" create
      ingress
        qos 40 shared-queuing
      exit
    exit
    egress
      qos 40
    exit
    no qos-marking-from-sap
  exit
exit
---snip---
sub-profile "sub-profile-default" create
  radius-accounting
    policy "accounting-11"
  exit
  sla-profile-map
    use-direct-map-as-default
  exit
exit
sub-ident-policy "sub-id-default" create
  sub-profile-map
    use-direct-map-as-default
  exit
  sla-profile-map
    use-direct-map-as-default
  exit
exit
```


MSAP Policy

MSAP policies contain the parameters which are used for MSAP creation and the information required to complete the subscriber identification process.

Creation of an MSAP requires an MSAP policy. The MSAP policy to be used can be defined during authentication. If authentication does not return an MSAP policy, then the default MSAP policy configured in the capture-sap as msap-defaults is used instead.

```
configure
  subscriber-mgmt
    msap-policy "msap-ISP1" create
      sub-sla-mgmt
        def-sub-id use-sap-id
        def-sub-profile "sub-profile-default"
        def-sla-profile "sla-profile-512K"
        sub-ident-policy "sub-id-default"
        single-sub-parameters
          profiled-traffic-only
        exit
      exit
    exit
  msap-policy "msap-default" create
    sub-sla-mgmt
      def-sub-id use-sap-id
      def-sub-profile "sub-profile-default"
      def-sla-profile "sla-profile-256K"
      sub-ident-policy "sub-id-default"
      single-sub-parameters
        profiled-traffic-only
      exit
    exit
  exit
exit
exit
exit
```

If managed routes are required for some subscribers, then the anti-spoof command is required in the msap-policy. The default value for anti-spoof is **ip-mac**. Managed routes are out of the scope of this chapter.

```
configure
  subscriber-mgmt
    msap-policy "msap-ISP1" create
      ies-vprn-only-sap-parameters
        anti-spoof nh-mac
      exit
    exit
  exit
exit
```

VPLS Service with a Capture SAP

Configure a VPLS service with capture SAP and define the triggering packet types. The **trigger-packet** is mandatory. In case of RADIUS authentication, an **authentication-policy** is required. Additionally, the **cpu-protection** command can be added to enable CPU protection policies, as follows:

```
configure
  service
    vpls 1 customer 1 create
      description "VPLS for Capture SAPs"
      stp
        shutdown
      exit
      sap 1/2/2:* capture-sap create
        description "capture SAP for MSAP creation on port 1/2/2"
        trigger-packet arp dhcp pppoe
        msap-defaults
          policy "msap-default"
        exit
        authentication-policy "authentication-1"
      exit
      no shutdown
    exit
  exit
exit
```

Verify the details of capture SAP:

```
*A:BNG-1# show service id 1 sap 1/2/2:* detail

=====
Service Access Points(SAP)
=====
Service Id       : 1
SAP              : 1/2/2:*           Encap              : q-tag
Description      : capture SAP for MSAP creation on port 1/2/2
Admin State      : Up                Oper State         : Up
Flags           : None
Multi Svc Site   : None
Last Status Change : 05/18/2017 15:44:05
Last Mgmt Change  : 05/22/2017 15:38:49
Sub Type         : capture
Triggers         : arp dhcp pppoe
Dot1Q Ethertype  : 0x8100           QinQ Ethertype     : 0x8100
Split Horizon Group: (Not Specified)

---snip---

Auth Policy      : authentication-1
DHCP User Db     : None
PPP Policy       : None
PPP User Db      : None
PPPoE Policy     : default
PPPoE User Db    : None
DHCPv6 User Db   : None
IPoE Policy      : None
IPoE User Db     : None
Rtr-Sol User Db  : None
DHCP Python policy : None
DHCP6 Python policy: None
PPPoE Python policy: None
```

```

Diameter auth plcy : None
Dynamic svc plcy   : None
Allow dot1q msap   : Disabled
DestMac Rewrite    : Disabled
SendBvplsEvpnFlush : Enabled

---snip---

-----
Sap Statistics
-----
Last Cleared Time      : N/A

CPM Ingress            Packets            Octets
: 474539                33476253

Forwarding Engine Stats
Dropped                : 9                842

DHCP Capture Stats
Received                : 0
Redirected              : 0
Dropped                 : 0

PPPoE Capture Stats
Received                : 406735
Redirected              : 0
Dropped                 : 0

ARP Capture Stats
Received                : 0
Redirected              : 0
Dropped                 : 0

DHCP6 Capture Stats
Received                : 0
Redirected              : 0
Dropped                 : 0

PPP Capture Stats
Received                : 0
Redirected              : 0
Dropped                 : 0

Rtr-Sol Capture Stats
Received                : 0
Redirected              : 0
Dropped                 : 0

Unknown Capture Stats
Received                : 0
Redirected              : 0
Dropped                 : 0

-----
Sap per Queue stats
-----
Packets            Octets

No entries found
=====
* indicates that the corresponding row element may have been truncated.
*A:BNG-1#

```

The Sap Statistics section provides statistics for the capture SAP per trigger type, which can help troubleshooting the service. The dropped packet counter indicates the amount of non triggering packets received on the capture SAP. No SAP queues are instantiated for a capture SAP.

VRPN Service - VLAN-Per-Subscriber (PPPOE)

The following output shows an RCO configuration example. No static SAPs are defined in this example, but it is allowed.

```
configure
  service
    vprn 2 customer 1 create
      route-distinguisher 64496:2
      subscriber-interface "sub-int-1" create
        address 10.255.255.254/8
        group-interface "grp-int-1" create
          description "ROUTED CO MSAP VLAN X"
          authentication-policy "authentication-1"
          pppoe
            session-limit 2000
            no shutdown
          exit
        exit
      exit
    exit
  exit
  no shutdown
exit
exit
exit
exit
```

Initially, no MSAPs are present, so the operational state of both the subscriber interface and group interface context are down.

```
*A:BNG-1# show router 2 interface

=====
Interface Table (Service: 2)
=====
Interface-Name      Adm      Opr(v4/v6)  Mode      Port/SapId
IP-Address                                     PfxState
-----
grp-int-1           Up        Down/Down   VPRN G*    n/a
sub-int-1           Up        Down/Down   VPRN S*    subscriber
10.255.255.254/8                                     n/a
-----
Interfaces : 2
=====
* indicates that the corresponding row element may have been truncated.
*A:BNG-1#
```

To allow the subscriber interface to consider this group interface to be operationally enabled without any active MSAPs, the following command can be added to the configuration (this would be useful in order to propagate the subnet interface address into a routing protocol):

```
configure
  service
    vprn 2
      subscriber-interface "sub-int-1"
      group-interface "grp-int-1"
```

```

oper-up-while-empty
exit
exit
exit
exit
exit

```

The status of the interfaces then is as follows:

```

*A:BNG-1# show router 2 interface

=====
Interface Table (Service: 2)
=====
Interface-Name      Adm      Opr(v4/v6)  Mode      Port/SapId
IP-Address          PfxState
-----
grp-int-1           Up       Down/Down   VPRN G*   n/a
sub-int-1           Up       Up/Down     VPRN S*   subscriber
10.255.255.254/8                    n/a
-----
Interfaces : 2
=====
* indicates that the corresponding row element may have been truncated.
*A:BNG-1#

```

Check the status of the group interface once the first MSAP is created.

RADIUS User File

The following entry is an example of a user entry in the RADIUS users file for the FreeRadius server:

```

"user1@ISP1.com" Cleartext-Password := "letmein"
  Alc-Subsc-ID-Str := "%{ADSL-Agent-Remote-Id}",
  Alc-SLA-Prof-Str == "sla-profile-2M",
  Alc-MSAP-Serv-ID = 2,
  Alc-MSAP-Policy == "msap-ISP1",
  Alc-MSAP-Interface == "grp-int-1",
  Framed-IP-Address = 10.255.0.1,
  Alc-Primary-DNS = 172.31.31.31,
  Alc-Secondary-DNS = 172.31.31.32,

```

So when the PPPoE user sends the correct username and password, the RADIUS accepts the access message and returns the correct VPRN service id 2, the correct group interface *group-int-1*, the MSAP policy to use *msap-ISP1*.

In case no MSAP policy is returned by the RADIUS server, the default MSAP policy *msap-default* under the capture SAP is used instead.

In the preceding entry, the PPPoE user will have its IP address and DNS assigned by RADIUS as well.

Connect PPPoE user

Connect PPPoE user *user1*, initiate a PPPoE session on VLAN 1, and verify PPPoE session establishment.

```

*A:BNG-1# show service id 2 pppoe session

```

```
=====
PPPoE sessions for svc-id 2
=====
Sap Id          Mac Address      Sid   Up Time          Type
  IP/L2TP-Id/Interface-Id          MC-Stdb
-----
[1/2/2:1]      00:00:00:01:01:01 1    0d 00:01:12     local
  10.255.0.1
-----
Number of sessions : 1
=====
*A:BNG-1#
```

The PPPoE session is established successfully and the IP address and subscriber strings obtained from the RADIUS server are used.

In order to differentiate between the MSAP and the normal SAP, the MSAP will be shown between square brackets [1/2/2:1] in the show commands.

Verify Subscriber Values

Verify subscriber values returned from RADIUS for user1.

```
*A:BNG-1# show service id 2 pppoe session ip-address 10.255.0.1 detail

=====
PPPoE sessions for svc-id 2
=====
Sap Id          Mac Address      Sid   Up Time          Type
  IP/L2TP-Id/Interface-Id          MC-Stdb
-----
[1/2/2:1]      00:00:00:01:01:01 1    0d 00:00:51     local
  10.255.0.1

LCP State       : Opened
IPCP State      : Opened
IPv6CP State    : Closed
PPP MTU         : 1492
PPP Auth-Protocol : CHAP
PPP User-Name   : user1@ISP1.com

Subscriber-interface : sub-int-1
Group-interface    : grp-int-1

IP Origin       : radius
DNS Origin      : radius
NBNS Origin     : none

Subscriber      : "user1"
Sub-Profile-String : ""
SLA-Profile-String : "sla-profile-2M"
ANCP-String     : ""
Int-Dest-Id     : ""
App-Profile-String : ""
Category-Map-Name : ""
Acct-Session-Id : "14F2FF00000006591EA903"
Sap-Session-Index : 1

IP Address      : 10.255.0.1/32
```

```

Primary DNS      : 172.31.31.31
Secondary DNS    : 172.31.31.32
Primary NBNS     : N/A
Secondary NBNS   : N/A
Address-Pool     : N/A

IPv6 Prefix      : N/A
IPv6 Prefix Origin : none
IPv6 Prefix Pool  : ""
IPv6 Del.Pfx.    : N/A
IPv6 Del.Pfx. Origin : none
IPv6 Del.Pfx. Pool  : ""
IPv6 Address     : N/A
IPv6 Address Origin : none
IPv6 Address Pool  : ""
Primary IPv6 DNS  : N/A
Secondary IPv6 DNS : N/A
Router adv. policy : N/A

Ignoring DF bit   : false
Radius sub-if prefix : N/A

Circuit-Id       : DSLAM1_1/1/1/1:0.35
Remote-Id        : user1

Radius Session-T0 : N/A
Radius Class      :
Radius User-Name  : user1@ISP1.com
Logical-Line-Id   :
Service-Name      :
-----
Number of sessions : 1
=====
*A:BNG-1#

```

Check Actual Values

Check the actual values used by *user1*, including the subscriber profile, SLA profile, VPRN and group interface association, the subscriber queues statistics and others.

```

*A:BNG-1# show service active-subscribers subscriber "user1" detail

=====
Active Subscribers
=====
-----
Subscriber user1 (sub-profile-default)
-----
I. Sched. Policy : N/A
E. Sched. Policy : N/A
I. Policer Ctrl. : N/A
E. Policer Ctrl. : N/A
I. vport-hashing : Disabled
I. sec-sh-hashing: Disabled
Q Frame-Based Ac*: Disabled
Acct. Policy      : N/A
ANCP Pol.        : N/A
HostTrk Pol.     : N/A
IGMP Policy      : N/A
MLD Policy       : N/A
PIM Policy       : N/A
E. Agg Rate Limit: Max
Collect Stats    : Disabled

```

```

Sub. MCAC Policy : N/A
NAT Policy      : N/A
Firewall Policy : N/A
UPnP Policy     : N/A
NAT Prefix List : N/A
Def. Encap Offset: none                      Encap Offset Mode: none
Avg Frame Size  : N/A
Vol stats type  : full
Preference      : 5
LAG hash class  : 1
LAG hash weight : 1
Sub. ANCP-String : "user1"
Sub. Int Dest Id : ""
Igmp Rate Adj   : N/A
RADIUS Rate-Limit: N/A
Oper-Rate-Limit : Maximum
-----
Radius Accounting
-----
Policy          : accounting-1
Session Opti.Stop: False
* indicates that the corresponding row element may have been truncated.
-----
(1) SLA Profile Instance
    - sap:[1/2/2:1] (VPRN 2 - grp-int-1)
    - sla:sla-profile-2M
-----
Description      : (Not Specified)
Host Limits      : No Limit
Egr Sched-Policy : N/A
Ingress Qos-Policy : 50                      Egress Qos-Policy : 50
Ingress Queuing Type : Shared-queuing (Not Applicable to Policer)
Ingr IP Fltr-Id   : N/A                      Egr IP Fltr-Id   : N/A
Ingr IPv6 Fltr-Id : N/A                      Egr IPv6 Fltr-Id : N/A
Ingress Report-Rate : Maximum
Egress Report-Rate  : Maximum
Egress Remarking    : from SLA Profile Qos
Credit Control Pol. : N/A
Category Map        : (Not Specified)
Use ing L2TP DSCP    : false
Hs-Agg-Rate-Limit   : Maximum
Hs-Oper-Rate-Limit  : Maximum
Egr hqos mgmt status : disabled
-----
IP Address      MAC Address      Session      Origin      Svc      Fwd
-----
10.255.0.1      00:00:00:01:01:01      PPP 1      IPCP      2      Y
-----
SLA Profile Instance statistics
-----
                Packets                Octets
Off. HiPrio      : 0                    0
Off. LowPrio     : 0                    0
Off. Uncolor     : 0                    0
Off. Managed     : 0                    0
Queueing Stats (Ingress QoS Policy 50)
Dro. HiPrio      : 0                    0
Dro. LowPrio     : 0                    0

```



```

For. InProf      : 0          0
For. OutProf     : 0          0

Queueing Stats (Egress QoS Policy 50)
Dro. In/InplusProf : 0          0
Dro. Out/ExcProf   : 0          0
For. In/InplusProf : 0          0
For. Out/ExcProf   : 2         128

-----
SLA Profile Instance per Queue statistics
-----
                          Packets          Octets

Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio        : 0          0
Off. LowPrio       : 0          0
Dro. HiPrio        : 0          0
Dro. LowPrio       : 0          0
For. InProf        : 0          0
For. OutProf       : 0          0

Egress Queue 1
Dro. In/InplusProf : 0          0
Dro. Out/ExcProf   : 0          0
For. In/InplusProf : 0          0
For. Out/ExcProf   : 2         128

=====
*A:BNG-1#

```

Where, the subscriber id is *user1*, and the subscriber profile is *sub-profile-default*.

Because the RADIUS server did not return a subscriber profile string, the system uses the **def-sub-profile** configured under the msap-policy *msap-ISP1*.

Another command can also be used to show less detail in a hierarchical form.

```

*A:BNG-1# show service active-subscribers hierarchy subscriber "user1"

=====
Active Subscribers Hierarchy
=====
-- user1 (sub-profile-default)
  |
  +-- sap:[1/2/2:1] - sla:sla-profile-2M
    |
    +-- PPP-session - mac:00:00:00:01:01:01 - sid:1 - svc:2
      |   circuit-id:DSLAM1_1/1/1/1:0.35
      |   remote-id:user1
      |
      +-- 10.255.0.1 - IPCP

=====
*A:BNG-1#

```

Verify that the IPv4 state of the group interface now is up, as follows:

```

*A:BNG-1# show router 2 interface

=====
Interface Table (Service: 2)

```

```
=====
Interface-Name      Adm      Opr(v4/v6)  Mode      Port/SapId
IP-Address          PfxState
-----
grp-int-1           Up        Up/Down      VPRN G*  1/2/2
sub-int-1           Up        Up/Down      VPRN S*  subscriber
10.255.255.254/8    n/a
-----
Interfaces : 2
=====
* indicates that the corresponding row element may have been truncated.
*A:BNG-1#
```

The MSAP details display the capture service id, capture SAP and MSAP policy, as follows:

```
*A:BNG-1# show service id 2 sap 1/2/2:1 detail

=====
Service Access Points(SAP)
=====
Service Id          : 2
SAP                 : 1/2/2:1                      Encap           : q-tag
Description         : Managed SAP - Capture Svc 1 1/2/2:*
Admin State         : Up                          Oper State      : Up
Flags               : None
Multi Svc Site      : None
Last Status Change  : 05/18/2017 15:43:43
Last Mgmt Change    : 05/19/2017 10:12:51
Sub Type            : managed
Capture Service Id  : 1                          Capture SAP     : 1/2/2:*
MSAP Policy         : msap-ISP1
Idle                : no                          Sticky          : no
Dot1Q Ethertype     : 0x8100                      QinQ Ethertype  : 0x8100
Split Horizon Group: (Not Specified)

---snip---

-----
Sap per Queue stats
-----
Packets      Octets
No entries found
=====
* indicates that the corresponding row element may have been truncated.
*A:BNG-1#
```

The **Sub Type** shows "managed" for MSAPs, or "regular" for normal SAPs (a SAP created manually under a group-interface).

MSAP QoS

By default an MSAP is created with default QoS policies.

```
*A:BNG-1# show service id 2 sap 1/2/2:1 detail

=====
Service Access Points(SAP)
=====
Service Id          : 2
SAP                 : 1/2/2:1                      Encap           : q-tag
```

```

Description      : Managed SAP - Capture Svc 1 1/2/2:*
Admin State      : Up                               Oper State      : Up

---snip---

-----
QoS
-----
Ingress qos-policy : 1                               Egress qos-policy : 1
Ingress FP QGrp    : (none)                           Egress Port QGrp  : (none)
Ing FP QGrp Inst   : (none)                           Egr Port QGrp Inst: (none)
Shared Q plcy      : default                           Multipoint shared : Disabled
I. Sched Pol       : (Not Specified)
E. Sched Pol       : (Not Specified)
I. Policer Ctl Pol : (Not Specified)
E. Policer Ctl Pol : (Not Specified)
E. HS Sec. Shaper  : (Not Specified)
I. QGrp Redir. List: (Not Specified)
E. QGrp Redir. List: (Not Specified)
-----
Subscriber Management
-----
Admin State      : Up                               MAC DA Hashing    : False
Def Sub-Id       : Use sap-id (1/2/2:1)
Def Sub-Profile   : sub-profile-default
Def SLA-Profile   : sla-profile-512K
Def Inter-Dest-Id : None
Def App-Profile   : None
Sub-Ident-Policy  : sub-id-default

---snip---

=====
* indicates that the corresponding row element may have been truncated.
*A:BNG-1#

```

The default QoS policy associated with MSAPs can be changed:

- To save queue resources when profiled-traffic-only cannot be used, for example when more than one subscriber is active on an MSAP. See further.
- To provide adequate QoS treatment for multicast traffic in a per MSAP replication mode.

Egress multicast traffic in per MSAP replication mode is forwarded via the MSAP queues or policers. Multicast traffic can be mapped into a dedicated queue or policer. The MSAP queue can be port-parented to provide scheduling priority at port level. The QoS policies associated with an MSAP are configured in the MSAP policy.

QoS Egress Remarking

For remarking to apply to MSAP egress traffic the SLA profile must include the **no qos-marking-from-sap** command, as follows:

```

configure
  subscriber-mgmt
    sla-profile "sla-profile-512K" create
    ---snip---
    egress
      qos 30
      exit
      no qos-marking-from-sap

```

```
        exit
    exit
exit
```

By default, the egress QoS marking for subscriber-host traffic is derived from the SAP-egress QoS policy associated with the corresponding SAP rather than the SLA profile associated with the corresponding subscriber-host. As a consequence, no egress QoS marking (for example, dot1p marking was set to 0, DSCP/PREC field is unchanged) is performed for traffic transmitted on an MSAP because by default, SAP-egress policy one (1) was attached to every MSAP.

MSAP Queue Optimization

For single subscriber SAPs, where the multi-sub-sap limit equals 1, the SAP queues will not be instantiated when using the **profiled-traffic-only** option in the msap-policy. This parameter is ignored when the multi-sub-sap limit is different from 1.

```
configure
  subscriber-mgmt
    msap-policy "msap-ISP1" create
      sub-sla-mgmt
        def-sub-id use-sap-id
        def-sub-profile "sub-profile-default"
        def-sla-profile "sla-profile-512K"
        sub-ident-policy "sub-id-default"
        single-sub-parameters
          profiled-traffic-only
      exit
    exit
  exit
exit
```

For multi subscriber MSAPs, a QoS policy can be associated with an MSAP in which all forwarding classes are mapped to a policer. In that case, a single ingress and egress policer is instantiated per MSAP (instead of ingress and egress queues). QoS policies associated with an MSAP are configured in the MSAP policy:

```
configure
  subscriber-mgmt
    msap-policy "msap-ISP2" create
      ies-vprn-only-sap-parameters
        egress
          qos 10
        exit
        ingress
          qos 10 shared-queuing
        exit
      exit
    exit
  exit
exit
```

Troubleshooting

The authentication policy used on the capture SAP must be the same as the policy used on the managed SAP.

The managed SAP will not be created if the authentication policy on the group-interface is different from the authentication policy defined on the capture SAP.

```
configure
  service
    vpls 1
      ---snip---
      sap 1/2/2:* capture-sap create
      ---snip---
      authentication-policy "authentication-1"
    exit
  no shutdown
  exit

configure
  service
    vprn 2
      subscriber-interface "sub-int-1" create
      ---snip---
      group-interface "group-int-1" create
      authentication-policy "authentication-2"
      ---snip---
    exit
  exit
  no shutdown
  exit
```

This can be seen in log 99:

```
*A:BNG-1# show log log-id 99

8 2017/05/19 10:50:43.70 CEST MINOR: SVCMMGR #2214 Base Managed SAP creation failure
"The system could not create Managed SAP:1/2/2:1, MAC:00:00:00:01:01:01, Capturing
SAP:1/2/2:*, Service:1. Description: MSAP group-interface "grp-int-1" RADIUS auth
-policy "authentication-2" differs from capture SAP"

7 2017/05/19 10:50:30.28 CEST WARNING: SVCMMGR #2501 Base Subscriber deleted
"Subscriber user1 has been removed from the system"

6 2017/05/19 10:50:29.68 CEST WARNING: SNMP #2004 vprn2 sub-int-1
"Interface sub-int-1 is not operational"

---snip---

*A:BNG-1#
```

Enable debug for PPPoE and RADIUS packets for troubleshooting purposes:

```
debug
  router "management"
    radius
      packet-type authentication accounting coa
      detail-level medium
    exit
  exit
```

```

service
  id 1
    ppp
      packet
        mode egr-ingr-and-dropped
        detail-level medium
        discovery
        ppp
      exit
    exit
  exit
  id 2
    ppp
      packet
        mode egr-ingr-and-dropped
        detail-level medium
        discovery
        ppp
        dhcp-client
      exit
    exit
  exit
exit

```

```

configure
  log
    log-id 1
      from debug-trace
      to session
    exit
  exit
exit

```

Disconnect/connect *user1*, then check the RADIUS access request/accept and accounting messages from the debug output.

```

11 2017/05/19 10:58:55.13 CEST MINOR: DEBUG #2001 management RADIUS
"RADIUS: Transmit
  Access-Request(1) 172.31.117.84:1812 id 202 len 174 vrid 4095 pol authenticat
ion-1
    USER NAME [1] 14 user1@ISP1.com
    NAS IP ADDRESS [4] 4 172.31.117.75
    SERVICE TYPE [6] 4 Framed(2)
    FRAMED PROTOCOL [7] 4 PPP(1)
    CHAP PASSWORD [3] 17 1 0x39721157837095dd2dc4a9351670e543
    CHAP CHALLENGE [60] 39 0x9e0eb2baf4c436f2f9a364ac0eb43cc6446943f5912d2c96570
ffd572732b245416501b5a9b6a8
    VSA [26] 7 DSL(3561)
      AGENT REMOTE ID [2] 5 user1
    NAS PORT TYPE [61] 4 PPPoEoVLAN(33)
    NAS PORT ID [87] 7 1/2/2:1
    NAS IDENTIFIER [32] 5 BNG-1
    VSA [26] 19 Nokia(6527)
      CHADDR [27] 17 00:00:00:01:01:01
"

12 2017/05/19 10:58:55.14 CEST MINOR: DEBUG #2001 management RADIUS
"RADIUS: Receive
  Access-Accept(2) id 202 len 131 from 172.31.117.84:1812 vrid 4095 pol authent
ication-1

```

```
VSA [26] 7 Nokia(6527)
  SUBSC ID STR [11] 5 user1
VSA [26] 16 Nokia(6527)
  SLA PROF STR [13] 14 sla-profile-2M
VSA [26] 6 Nokia(6527)
  MSAP SERVICE ID [31] 4 2
VSA [26] 11 Nokia(6527)
  MSAP POLICY [32] 9 msap-ISP1
VSA [26] 11 Nokia(6527)
  MSAP INTERFACE [33] 9 grp-int-1
FRAMED IP ADDRESS [8] 4 10.255.0.1
VSA [26] 6 Nokia(6527)
  PRIMARY DNS [9] 4 172.31.31.31
VSA [26] 6 Nokia(6527)
  SECONDARY DNS [10] 4 172.31.31.32
"
```

The MSAP policies can be checked as follows:

```
*A:BNG-1# show subscriber-mgmt msap-policy
```

```
=====
Managed SAP Policies
=====
Name                               Num      Description
                                MSAPs
-----
msap-ISP1                          1      (Not Specified)
msap-default                       0      (Not Specified)
-----
Number of MSAP Policies : 2
Number of MSAPs         : 1
=====
*A:BNG-1#
```

The MSAP policy associations can be checked as follows:

```
*A:BNG-1# show subscriber-mgmt msap-policy "msap-ISP1" association
```

```
=====
MSAP Policy Associations
=====
Service-Id : 2 (VPRN)
- SAP : [1/2/2:1]
-----
Number of associated MSAPs: 1
Flags: (I) = Idle MSAP
=====
*A:BNG-1#
```

All MSAPs created and associations with the services can be checked as follows:

```
*A:BNG-1# show service sap-using msap
```

```
=====
Service Access Points
=====
PortId                               SvcId      Ing.   Ing.   Egr.   Egr.   Adm   Opr
                                QoS       Fltr   QoS   Fltr
-----
[1/2/2:1]                           2          1     none   1     none   Up    Up
-----
```

```
Number of SAPs : 1
-----
Number of Managed SAPs : 1, indicated by [<sap-id>]
Flags : (I) = Idle MSAP
-----
=====
*A:BNG-1#
```

It is possible to use a **tools** command to update an existing MSAP when a specific msap-policy has changed.

```
*A:BNG-1# tools perform subscriber-mgmt eval-msap ?
- eval-msap { policy <msap-policy-name> | msap <sap-id> }

<msap-policy-name> : [32 chars max]
<sap-id>             : dot1q          - <port-id|lag-id>:qtag1
                      qtag1          - [0..4094]
                      qinq           - <port-id|lag-id>:qtag1.qtag2
                      qtag1          - [0..4094]
                      qtag2          - [0..4094]
                      atm            - <port-id>:vpi/vci
                      vpi            - [0..4095] (NNI)
                                   [0..255] (UNI)
                      vci            - [1..65535]
                      port-id        - slot/mda/port
                      lag-id         - lag-<id>
                      lag            - keyword
                      id              - [1..800]

*A:BNG-1#
```

An MSAP can be deleted as follows:

```
*A:BNG-1# clear service id 2 msap 1/2/2:1
```

This event is recorded in log 99 as follows:

```
*A:BNG-1# show log log-id 99

=====
Event Log 99
=====
Description : Default System Log
Memory Log contents [size=500 next event=55 (not wrapped)]

54 2017/05/19 11:24:04.29 CEST WARNING: SVCMMGR #2501 Base Subscriber deleted
"Subscriber user1 has been removed from the system"

53 2017/05/19 11:24:04.03 CEST INDETERMINATE: LOGGER #2010 Base Clear SVCMMGR
"Clear function clearSvcIdMsap has been run with parameters: svc-id="2" sap-id="1/2
/2:1". The completion result is: success. Additional error text, if any, is: "

---snip---

*A:BNG-1#
```

To delete all MSAPs associated with a certain MSAP policy use the following command:

```
*A:BNG-1# clear service id 2 msap-policy msap-ISP1
```


This event is recorded in log 99 as follows:

```
*A:BNG-1# show log log-id 99

=====
Event Log 99
=====
Description : Default System Log
Memory Log contents [size=500  next event=74  (not wrapped)]

67 2017/05/19 11:29:15.28 CEST WARNING: SVCMMGR #2501 Base Subscriber deleted
"Subscriber user1 has been removed from the system"

66 2017/05/19 11:29:14.54 CEST INDETERMINATE: LOGGER #2010 Base Clear SVCMMGR
"Clear function clearSvcIdMsapPlcy has been run with parameters: svc-id="2" policy
-name="msap-ISP1". The completion result is: success. Additional error text, if any,
is: "

65 2017/05/19 11:29:14.54 CEST MINOR: SVCMMGR #2213 vprn2 MSAP delete
"Managed SAP, 1/2/2:1 in service 2, has been deleted."

---snip---

*A:BNG-1#
```

Conclusion

MSAP allows dynamic creation of SAPs which results in:

- Less provisioning.
- Less possibility for introducing provisioning errors.
- Reduced configuration file size.

Multi-Chassis Ring Layer 2 with Enhanced Subscriber Management

This chapter provides information about MC-ring layer 2 with enhanced subscriber management (ESM).

Topics in this chapter include:

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

Applicability

The information and configuration in this chapter is based on SR OS Release 7.0.R5.

Overview

Multi-Chassis Ring (MC-ring) is an extension for dual homing support in TPSDA (Triple Play Service Delivery Architecture) networks based on Layer 2 CO (Layer 2 Central Office) model. The extension addresses networks where multiple access nodes (for example, DSLAMs, GPON OLT) are connected in a single ring.

MC Ring Layer 2 ESM is considered an extension or evolution for ring topologies of the Multi-Chassis LAG dual-homing solution used for directly connected access nodes.


MC Ring Layer 2 CO is documented in the Triple Play Enhanced Subscriber Management / Dual Homing section of the 7450 ESS, 7750 SR, and VSR Triple Play Service Delivery Architecture Guide.

MC-Ring

[Figure 139: MC-Ring Layer 2 CO Dual Homing](#) shows a typical ring-based configuration of network model based on the Layer 2 CO model.

Individual rings of access nodes are aggregated at the Broadband Service Aggregator (BSA) level in one (or multiple) Virtual Private LAN Service (VPLS) service(s). At higher aggregation levels, Broadband Service Router (BSR) individual BSAs are connected to a Layer 3 interface (Internet Enhanced Service (IES) or Virtual Private Routed Network (VPRN)) by means of a spoke SDP (service destination point) termination. Every Layer 3 interface at the BSR level aggregates all subscribers in one or more subnets. ESM is performed in the BSAs.

Figure 139: MC-Ring Layer 2 CO Dual Homing



OSSG424

The key functional components of the MC-Ring Layer 2 CO dual-homing redundancy solution are:

1. Mirroring of the subscriber state between the two BSAs performing subscriber management using the multi-chassis synchronization (MCS) protocol (BSA-1 and BSA-2 in [Figure 139: MC-Ring Layer 2 CO Dual Homing](#)).
2. Ring control between the two BSAs, using the following mechanisms:
 - In-band bi-directional forwarding detection (BFD) between the BSAs over the ring to monitor ring integrity and detect failures. A BFD session between BSA-1 and BSA-2 runs through the access ring using a dedicated IES/VPDN interface configured on both BSAs. This connection uses a separate VLAN throughout the ring (access nodes provide transparent bridging for this VLAN).
 - Out-of-band communication between the BSA nodes to exchange information about the reachability of individual access nodes, and to verify the configuration consistency of the ring. The information on configuration is synchronized through MCS (this use of MCS is related to, but independent to, MCS for subscriber-state synchronization mentioned above).
3. Ring Node Connectivity Verification (RNCV). Each BSA uses RNCV to detect the reachability of individual Access Nodes. It is used after a ring failure to determine which BSA should handle the traffic of each Access Node. RNCV uses ARP requests to "ping" individual ANs, which must be configured with an IP address for this purpose.
4. Per-subscriber attribute (intermediate destination ID) for the BSA to know the Access Node each subscriber is connected to (assigned through RADIUS or DHCP/Python).
5. VRRP in the BSRs to provide a redundant default gateway to the CPEs/Home Gateways. BSRs do not perform any subscriber management functions.

The operation of dual homing at the BSA level will be covered based on two underlying mechanisms.

MC-Ring Layer 2 CO Operation

To describe the functional behavior and operation of the dual-homing concept in a ring, it is best to sub-divide the description into the following three parts:

- Steady-state operation with ring fully closed
- Transition to ring-broken state
- Transition from ring-broken to steady state operation

Steady-State Operation of Dual-Homed Ring

Figure 140: Dual homing Under Steady-State Condition illustrates the detailed operation of the dual-homed ring. The steady-state is achieved under following conditions:


- Both nodes are configured in a consistent way
- The MCS peering relation is up
- In-Band Ring Control Connection (IB-RCC) is in an operationally UP state.



Note:

This connection is set-up using BFD session between IP interfaces on BSA-1 and BSA-2

Figure 140: Dual homing Under Steady-State Condition



OSSG425

Under the above conditions, the ring is fully closed and every access node (the **ring-node**) has two possible paths toward the VPLS core. **Figure 140: Dual homing Under Steady-State Condition** refers to them as **path-a** and **path-b**. In order to avoid the loop created by the ring, only one of the paths may be used by any given ring-node for any given VLAN. The assignment of the individual VLANs to path-a or path-b, respectively, has to be provisioned on both BSAs in a consistent manner. The BSA with a lower IP address in the interface used for BFD will be the master for the VLANs associated with path-a and standby for the VLANs associated with path-b.

The following summarizes the behavior of the master and standby BSA for a given path (a or b) and the VLANs configured for that path:

Master BSA for the VLANs/SAPs associated with a given path (a or b):

- SAPs associated with the path are operationally up and FDB entries for sub-hosts point to the corresponding SAP
- Master BSA for a path performs RNCV checks to all ring nodes. RNCV failures trigger an alarm but do **not** trigger a switchover
- The ARP reply agent replies to ARP requests for subscriber-hosts associated with ring-nodes for which the BSA is master

Standby BSA:

- All SAPs associated with the path for which the BSA is standby will be operationally down and all FDB entries of subscriber hosts associated with those SAPs will point towards an SDP connecting to the master BSA.

Broken-Ring Operation and the Transition to this State


Figure 141: Broken Ring State illustrates the scenario with the broken ring (link failure or ring-node failure). This state is reached under following conditions:

- Both nodes are configured in a consistent way
- The MCS peering relation is up
- IB-RCC is in operationally down state

In this situation, every ring node has only one access path towards the VPLS core and hence the path-a and path-b notion has no real meaning in this situation.

From a functional point of view, both BSAs are now the master for the ring-nodes they can reach. For all hosts behind unreachable ring-nodes, the corresponding subscriber host FDB (Forwarding DataBase) entries will point to SDP pointing to the other head-end ring PE.

Figure 141: Broken Ring State



OSSG426

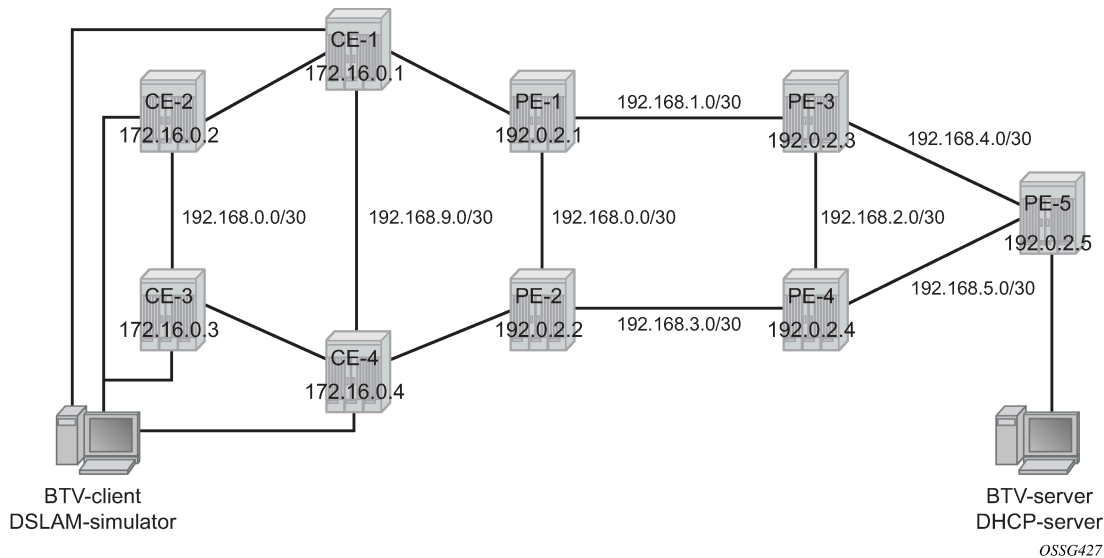
Transition from Broken to Closed Ring State

MC-ring operates in a revertive mode. This means that whenever the ring connectivity is restored, the BSA with the lower IP address in the IB-RCC communication channel will become master of path-a and slave for ring-b with vice versa operations on the other BSA.

The actions in this case are straightforward. After restoration of BFD session, the master functionality is assumed by respective BSAs. The FDB tables are updated according to the master/standby role of the given BSA and FDB population messages will be sent accordingly.

Configuration

Figure 142: Network Topology



The network topology is displayed in [Figure 142: Network Topology](#). The setup consists of two BSA nodes (PE-1 and PE-2), two BSR nodes (PE-3, PE-4) and another PE router (PE-5). A setup with one BSR node and two BSA nodes can also be used, but in this example, two (2) BSR nodes were used to show the typical Layer 2-CO setup with VRRP and Protocol Independent Multicast (PIM) on the BSRs. The access ring consists of four CE nodes.

Base Topology

This section assumes that the following base configuration has been implemented on the PE:

- Cards, MDAs and ports configured
- Interfaces configured
- IGP (Interior Gateway Protocol) configured and converged LDP (Label Distribution Protocol) configured on the interfaces between PE-3-PE-1, PE-1-PE-2 and PE-2-PE-4
- T-LDP (Targeted-LDP) configured on PE-1, PE-2, PE-3, PE-4

- SDPs configured between PE-3-PE-1, PE-1-PE-2 and PE-2-PE-4



Note:

You can choose between OSPF and IS-IS as the IGP. Both LDP and RSVP (Resource Reservation Protocol) can be used for signaling the transport MPLS labels. Alternatively, GRE (Generic Routing Encapsulation) can be used for the transport tunnels. In this example, OSPF is configured and LDP SDPs are used.

7750 SR routers are used as ring-nodes to simulate Access Nodes. Ring-nodes can be any L2 device that support VLANs and have the ability to connect an IP interface to one VLAN (required for RNCV).

NTP Configuration

Time must be the same on the redundant NSAs (PE-1 and PE-2); otherwise, the lease times will be different on both nodes. NTP can be used:

```
configure system time
    ntp
        server 10.30.30.30 prefer
        no shutdown
    exit
exit all
```

MC-Ring Configuration

Configure a VPLS service on the CE routers for the In-Band Ring Control connection (BFD packets between PE-1 and PE-2 traversing the ring). This VPLS service will also be used for RNCV.



Note:

An Epipe service can also be used in case the service is only used for BFD. In that case, a separate service is required for the RNCV.

In this example, QinQ encapsulation is used and BFD and RNCV packets will use VLAN tag 1. Note that dot1q encapsulation can also be used.

The configuration of CE-1 is shown below. A similar configuration is required on the other CE routers.

```
configure
    port 1/1/1
        ethernet
            mode access
            encap-type qinq
        exit
        no shutdown
    exit
    port 1/1/2
        ethernet
            mode access
            encap-type qinq
        exit
        no shutdown
    exit
    service
        vpls 1 customer 1 create
        interface "lo1" create
```

```
        address 172.16.0.1/24
    exit
    sap 1/1/1:1.0 create
    exit
    sap 1/1/2:1.0 create
    exit
    no shutdown
    exit
exit all
```

An interface *lo1* is created in the VPLS. This interface will be used for RNCV.

On the BSA nodes (PE-1 and PE-2), configure an IES services that will originate BFD and RNCV messages. On PE-1:

```
configure
  port 1/1/1
    ethernet
      mode access
      encap-type qinq
    exit
    no shutdown
  exit
  service
    ies 1 customer 1 create
      interface "bfd-rncv1" create
        address 172.16.0.251/24
        bfd 100 receive 100 multiplier 3
        sap 1/1/1:1.0 create
      exit
    exit
    no shutdown
  exit
exit all
```

On PE-2:

```
configure
  port 1/1/2
    ethernet
      mode access
      encap-type qinq
    exit
    no shutdown
  exit
  service
    ies 1 customer 1 create
      interface "bfd-rncv1" create
        address 172.16.0.252/24
        bfd 100 receive 100 multiplier 3
        sap 1/1/2:1.0 create
      exit
    exit
    no shutdown
  exit
exit all
```

In-Band Ring Control Connection BFD is always originated on an IES or VPRN service on the BSA nodes. RNCV messages can be originated from the same IES/VPRN service or from a separate service.

Configure Multi-Chassis Synchronization (MCS) on the BSA nodes. Enable MCS for **igmp-snooping**, **mc-ring** and **sub-mgmt**. The configuration of PE-1 is shown below. The configuration of PE-2 is similar.

```
configure redundancy multi-chassis peer 192.0.2.2 create
    sync
        igmp-snooping
        mc-ring
        sub-mgmt
        port 1/1/1 sync-tag "l2ring1" create
        exit
        no shutdown
    exit
    no shutdown
exit all
```

Add the MC-ring configuration on the BSA nodes and link the Ring Control Connection BFD session to the IES service created before.

```
configure redundancy multi-chassis peer 192.0.2.2 create
    mc-ring
        ring "l2ring1" create
            in-band-control-path
                service-id 1
                interface "bfd-rncv1"
                dst-ip 172.16.0.252
            exit
            no shutdown
        exit
    exit
    no shutdown
exit all
```

Note that the ring name is exactly the same as the sync-tag already configured.

At this moment, the MC-ring should be up:

```
A:PE-1# show redundancy multi-chassis mc-ring peer 192.0.2.2 ring l2ring1 detail
=====
Multi-Chassis MC-Ring Detailed Information
=====
Peer          : 192.0.2.2
Ring Type     : Layer 2
Sync Tag      : l2ring1
Port ID       : 1/1/1
Admin State    : inService
Oper State     : connected
Admin Change   : 11/05/2009 21:17:54
Oper Change    : 11/05/2009 21:17:54
Failure Reason : None
Control B Path : No
-----
In Band Control Path
-----
Service ID     : 1
Interface Name  : bfd-rncv1
Oper State      : connected
Dest IP        : 172.16.0.252
Src IP         : 172.16.0.251
...
```

Next, configure under MCS the ring nodes on PE-1 and PE-2. This configuration will be used for the RNCV. The ring node configuration for CE-1 is shown below:

```
configure redundancy multi-chassis peer 192.0.2.2 mc-ring ring l2ring1
    ring-node "CE-1" create
    connectivity-verify
        dst-ip 172.16.0.1
        interval 1
        service-id 1
        vlan 1
        no shutdown
    exit
exit
exit all
```



Note:

The **interval** parameter is the interval used to check the reachability of the CE nodes (configurable from 1 to 6000 minutes). A BFD failure will also be a trigger for a reachability check.

The configuration on PE-2 is identical and the configuration of the other ring nodes is similar.

Configure VLAN 3 to take path-b (default is path-a).

```
configure redundancy multi-chassis peer 192.0.2.2 mc-ring ring l2ring1
    path-b
        range 3-3
    exit
exit all
```

MC-Ring Verification

Verify that MCS is up and running:

```
A:PE-1# show redundancy multi-chassis all
=====
Multi-Chassis Peers
=====
Peer IP      Src IP      Auth      Peer Admin  MC-Ring Oper MC-EP Adm
MCS Admin   MCS Oper   MCS State MC-LAG Adm  MC-LAG Oper
-----
192.0.2.2    192.0.2.1  None      Enabled     inService    --
Enabled      Enabled    inSync    Disabled    Disabled
=====
```

The output shows that MCS is administrative and operational up and that both peers are synchronized. MC-ring is operational in service.

The following output shows more detailed information about the configured ring:

```
A:PE-1# show redundancy multi-chassis mc-ring peer 192.0.2.2 ring l2ring1 detail
=====
Multi-Chassis MC-Ring Detailed Information
=====
Peer      : 192.0.2.2
Ring Type : Layer 2
Sync Tag  : l2ring1
Port ID   : 1/1/1
Admin State : inService
```

```
Oper State      : connected
Admin Change   : 11/05/2009 21:22:29
Oper Change    : 11/05/2009 21:22:29
Failure Reason : None
Control B Path : No
```

In Band Control Path

```
Service ID     : 1
Interface Name : bfd-rncv1
Oper State     : connected
Dest IP       : 172.16.0.252
Src IP        : 172.16.0.251
Debounce State : inService
Debounce Max   : 10 s
Debounce Guard : 60 s
```

VLAN Managed Range

```
full range
```

VLAN Map B Path Provisioned

```
range 3-3
```

VLAN Map Excluded Path Provisioned

```
no range
```

VLAN Map B Path Operational

```
range 3-3
```

VLAN Map Excluded Path Operational

```
no range
```

The output above shows that the ring *l2ring1* on port 1/1/1 is in service and connected. Ring-node Connectivity Check (BFD) is running on **interface bfd-rncv1** in service 1. All VLANs on port 1/1/1 use path-a (default) except VLAN 3, which uses path-b.

The BSA peer with the lower IP address (the master) will put the SAPs configured for path-a in operational up state while the SAPs configured in path-b will be put in operational down state. The other BSA peer will do the reverse. This is done to prevent loops in the ring.



Note:

No VLANs are configured to be excluded from the paths. If a VLAN is configured to be excluded from the paths, the respective SAPs of this VLAN on both BSA nodes will be operationally up.

Check which ring-nodes are configured and connected:


```
A:PE-1# show redundancy multi-chassis mc-ring peer 192.0.2.2 ring l2ring1 ring-node
```

MC-Ring Node entries

Name	Loc Oper St.	Failure Reason
In Use	Rem Oper St.	
CE-1	connected	None
No	notTested	
CE-2	connected	None
No	notTested	

```
CE-3          connected      None
No            notTested
CE-4          connected      None
No            notTested
-----
No. of MC-Ring Node entries: 4
=====
A:PE-1#
```

The output above shows that four ring-nodes are connected. Only the master will send RNCV messages to the ring-nodes. As soon as a ring failure occurs, the BFD session goes down and both BSA nodes send out RNCV messages to see which ring-nodes are connected.



Note:
When the reachability of the CE nodes is determined, the RNCV messages will be sent at a continuous interval (see above).

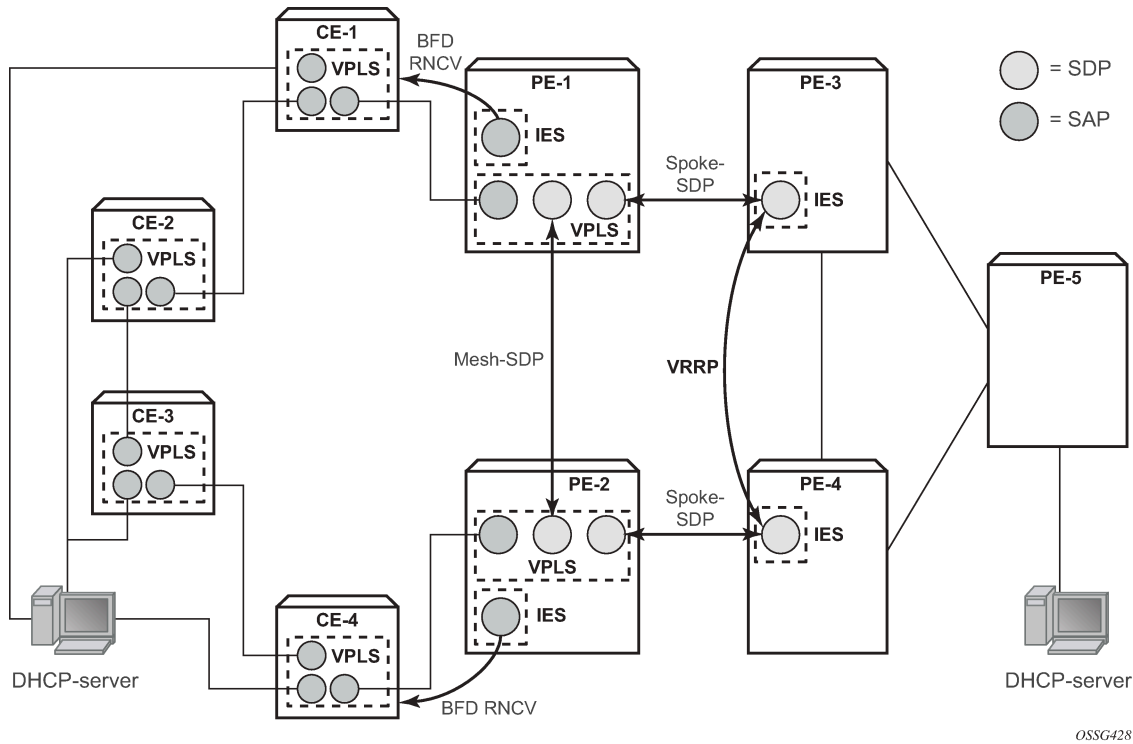
More detail about each ring-node can be provided with following command:

```
A:PE-1# show redundancy multi-chassis mc-ring peer 192.0.2.2 ring l2ring1 ring-node CE-1
detail
=====
Multi-Chassis MC-Ring Node Detailed Information
=====
Peer           : 192.0.2.2
Sync Tag       : l2ring1
Node Name      : CE-1
Oper State Loc : connected
Oper State Rem : notTested
In Use         : False
Admin Change   : 11/05/2009 21:21:30
Oper Change    : 11/05/2009 21:22:32
Failure Reason : None
-----
Ring Node Connectivity Verification
-----
Admin State    : inService
Service ID     : 1
Encap Value    : 1
Dest IP        : 172.16.0.1
Src IP         : None
Interval       : 1 minutes
Src MAC        : None
=====
A:PE-1#
```

Unicast Services Configuration

Figure 143: Unicast Services — Logical Topology shows the logical setup of the services that will be created. A mesh SDP is used between PE-1 and PE-2. A spoke SDP could also be used.

Figure 143: Unicast Services — Logical Topology



In this setup, two unicast services (VPLS 2 and VPLS 3) are created. VPLS 2 uses path-a and VPLS 3 uses path-b.

The services use a mesh SDP between PE-1 and PE-2 and a spoke-SDP between PE-1/PE-3 and another spoke SDP between PE-2/PE-4. On PE-3 and PE-4 a spoke SDP terminated IES is configured with a VRRP default gateway address. The VRRP packets are switched through the BSAs.

IGMP snooping and ESM are configured on both services. ESM is required since the BSA node must know which ring-node each subscriber is connected to. In this setup, the intermediate destination identifier (int_dest_id) will be returned by Option 254 in x Dynamic Host Control Protocol (DHCP). ESM configuration details are outside the scope of this document. Refer to the 7450 ESS, 7750 SR, and VSR Triple Play Service Delivery Architecture Guide.

The following output shows the configuration of VPLS 2 on PE-1:

```
configure service
  vpls 2 customer 1 create
    description "VLAN_2"
    split-horizon-group "RSHG" residential-group create
  exit
  sap 1/1/1:2.* split-horizon-group "RSHG" create
    dhcp
      snoop
      lease-populate 10
      no shutdown
    exit
  anti-spoof ip-mac
  sub-sla-mgmt
    def-sub-profile "initial"
    def-sla-profile "initial"
```

```

        sub-ident-policy "speedy"
        multi-sub-sap 10
        no shutdown
    exit
exit
mesh-sdp 12:2 create
    dhcp
        snoop
    exit
exit
spoke-sdp 13:2 create
    dhcp
        snoop
    exit
exit
no shutdown
exit
exit all

```

The configuration of VPLS 3 is similar. The configuration of VPLS 2 and 3 are similar on PE-2.

The output below shows the subscriber management configuration on PE-1. Similar configuration is required on PE-2.

```

configure subscriber-mgmt
    sla-profile "initial" create
    exit
    sub-profile "initial" create
    exit
    sub-ident-policy "speedy" create
        strings-from-option 254
    exit
exit all

```

Also, configure VLAN 2 and 3 on all the ring-nodes. This configuration is straightforward. Below is an example of the VLAN 2 configuration on CE-1. In this example, the client is connect to port 1/1/3 and should send packets with an outer tag of 2:

```

configure service
    vpls 2 customer 1 create
        sap 1/1/1:2.* create
        exit
        sap 1/1/2:2.* create
        exit
        sap 1/1/3:2.* create
        exit
        no shutdown
    exit
exit all

```

In the example, VLAN 2.* and 3.* is used to allow for transparent transport of the customer VLAN.

The IES service where VPLS 2 terminates looks like this on BSR PE-3:

```

configure service
    ies 2 customer 1 create
        interface "VLAN_2" create
            address 10.0.2.3/24
            dhcp
                server 10.10.10.10
            trusted

```

```
        no shutdown
        exit
        ip-mtu 1500
        vrrp 2
            backup 10.0.2.254
            ping-reply
        exit
        local-proxy-arp
        spoke-sdp 31:2 create
        exit
    exit
    no shutdown
    exit
exit all
```

And on PE-4:

```
configure service
    ies 2 customer 1 create
        interface "VLAN_2" create
            address 10.0.2.4/24
            dhcp
                server 10.10.10.10
                trusted
                no shutdown
            exit
            ip-mtu 1500
            vrrp 2
                backup 10.0.2.254
                ping-reply
            exit
            local-proxy-arp
            spoke-sdp 42:2 create
            exit
        exit
        no shutdown
    exit
exit all
```

Notice that the ip-mtu must be set to match the vc-mtu signaled by the other side of the spoke-SDP. Otherwise, the service will be operationally down with a ServiceMTUMismatch.

Note also in the configuration that DHCP relay is done by configuring a DHCP server under the IES interface.

The configuration of IES 3 on PE-3 and PE-4 is similar..

On PE-5 an IES interface is configured to the DHCP-server. The interface is configured with a Dot1Q encapsulated port because this port will be also be used for an interface to the multicast server.

```
configure service
    ies 2 customer 1 create
        interface "dhcp-server" create
            address 192.168.6.1/30
            sap 1/1/3:2 create
            exit
        exit
        no shutdown
    exit
exit all
```

Unicast Services Verification

Request an IP address on VLAN 2 on CE-1. On BSA routers PE-1 and PE-2 following DHCP info can be checked:

```
A:PE-1# show service id 2 dhcp lease-state
=====
DHCP lease state table, service 2
=====
IP Address      Mac Address      Sap/Sdp Id      Remaining   Lease    MC
                  LifeTime        Origin          Stdby
-----
10.0.2.107      00:00:64:01:01:02 1/1/1:2.*      09d07h38m  DHCP
-----
Number of lease states : 1
=====
```

ESM show commands can be used to obtain the subscriber identity, IP address, MAC address and SLA-profile:

```
A:PE-1# show service active-subscribers
=====
Active Subscribers
=====
Subscriber subscriber_1_vlan_2 (initial)
-----
(1) SLA Profile Instance sap:1/1/1:2.* - sla:initial
-----
IP Address      MAC Address      PPPoE-SID Origin
-----
10.0.2.107      00:00:64:01:01:02 N/A          DHCP
-----
Number of active subscribers : 1
=====
```

The following command gives more information about a specific subscriber:

```
A:PE-1# show service active-subscribers subscriber subscriber_1_vlan_2 detail
=====
Active Subscribers
=====
Subscriber subscriber_1_vlan_2 (initial)
-----
I. Sched. Policy : N/A
E. Sched. Policy : N/A
Q Frame-Based Ac*: Disabled
Acct. Policy      : N/A
Rad. Acct. Pol.   : N/A
Dupl. Acct. Pol.  : N/A
ANCP Pol.         : N/A
HostTrk Pol.      : N/A
Sub. ANCP-String  : "subscriber_1_vlan_2"
Sub. Int Dest Id  : "CE-1"
Host Trk Rate Adj: N/A
E. Agg Rate Limit: Max
Collect Stats     : Disabled
```



```

-----
(1) SLA Profile Instance
  - sap:1/1/1:2.* (VPLS 2)
  - sla:initial
-----
Description      : (Not Specified)
Host Limit       : No Limit
Ingress Qos-Policy : 1                      Egress Qos-Policy : 1
Ingress Queuing Type : Service-queuing
Ingress Filter-Id  : N/A                    Egress Filter-Id   : N/A
Ingress Report-Rate : N/A
Egress Report-Rate  : N/A
Egress Remarking    : from Sap Qos
Credit Control Pol. : N/A
-----
-----
IP Address      MAC Address      PPPoE-SID Origin
-----
10.0.2.107      00:00:64:01:01:02 N/A      DHCP
-----

```

The output above gives more details about which ring-node the customer is connected to (Sub. Int Dest Id : CE-1), which QoS policies are applied, statistics of each queue,

MCS Verification

Check if the two redundant peers are in sync and check the detailed MCS info:

```

A:PE-1# show redundancy multi-chassis sync peer 192.0.2.2
=====
Multi-chassis Peer Table
=====
Peer
-----
Peer IP Address      : 192.0.2.2
Description          : (Not Specified)
Authentication       : Disabled
Source IP Address    : 192.0.2.1
Admin State          : Enabled
-----
Sync-status
-----
Client Applications  : IGMP Snooping SUBMGMT RING
Sync Admin State     : Up
Sync Oper State      : Up
DB Sync State        : inSync
Num Entries          : 11
Lcl Deleted Entries  : 0
Alarm Entries        : 0
Rem Num Entries      : 11
Rem Lcl Deleted Entries : 0
Rem Alarm Entries    : 0
=====
MCS Application Stats
=====
Application          : igmp
Num Entries          : 0
Lcl Deleted Entries  : 0
Alarm Entries        : 0
-----
Rem Num Entries      : 0
Rem Lcl Deleted Entries : 0

```

```

Rem Alarm Entries      : 0
-----
Application           : igmpSnooping
Num Entries           : 0
Lcl Deleted Entries    : 0
Alarm Entries         : 0
-----
Rem Num Entries       : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries     : 0
-----
Application           : subMgmt
Num Entries           : 1
Lcl Deleted Entries    : 0
Alarm Entries         : 0
-----
Rem Num Entries       : 1
Rem Lcl Deleted Entries : 0
Rem Alarm Entries     : 0
-----
Application           : srrp
Num Entries           : 0
Lcl Deleted Entries    : 0
Alarm Entries         : 0
-----
Rem Num Entries       : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries     : 0
-----
Application           : mcRing
Num Entries           : 10
Lcl Deleted Entries    : 0
Alarm Entries         : 0
-----
Rem Num Entries       : 10
Rem Lcl Deleted Entries : 0
Rem Alarm Entries     : 0
-----
Application           : mldSnooping
Num Entries           : 0
Lcl Deleted Entries    : 0
Alarm Entries         : 0
-----
Rem Num Entries       : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries     : 0
-----
Application           : dhcpServer
Num Entries           : 0
Lcl Deleted Entries    : 0
Alarm Entries         : 0
-----
Rem Num Entries       : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries     : 0
-----
Application           : subHostTrk
Num Entries           : 0
Lcl Deleted Entries    : 0
Alarm Entries         : 0
-----
Rem Num Entries       : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries     : 0

```

```
=====
A:PE-1#
```

The output shows that both MCS peers are in sync and that entries are populated for MC-ring and Subscriber Management (DHCP lease states).

Notice that the lease states are also populated on PE-2:

```
A:PE-2# show service id 2 dhcp lease-state
=====
DHCP lease state table, service 2
=====
IP Address      Mac Address      Sap/Sdp Id      Remaining   Lease   MC
LifeTime        Origin          Stdbby
-----
10.0.2.107      00:00:64:01:01:02 1/1/2:2.*      09d06h16m  DHCP    Yes
-----
Number of lease states : 1
=====
A:PE-2#
```

The output is similar to the output on PE-1 except that on PE-2 the flag MC-Stdbby is set to yes, which implies that this node is in standby mode for this VLAN.

This can be verified by looking at the status of the SAP on PE-1 and PE-2. On PE-1 the SAP is operationally up:

```
A:PE-1# show service id 2 sap 1/1/1:2.*
=====
Service Access Points(SAP)
=====
Service Id      : 2
SAP             : 1/1/1:2.*      Encap           : qinq
QinQ Dot1p     : Default
Description     : (Not Specified)
Admin State     : Up            Oper State      : Up
Flags           : None
Multi Svc Site  : None
Last Status Change : 11/06/2009 17:17:26
Last Mgmt Change  : 11/04/2009 22:51:15
=====
```

On PE-2 the situation is different:

```
A:PE-2# show service id 2 sap 1/1/2:2.*
=====
Service Access Points(SAP)
=====
Service Id      : 2
SAP             : 1/1/2:2.*      Encap           : qinq
QinQ Dot1p     : Default
Description     : (Not Specified)
Admin State     : Up            Oper State      : Down
Flags           : StandByForMcRing
Multi Svc Site  : None
Last Status Change : 11/06/2009 18:07:30
Last Mgmt Change  : 11/04/2009 23:43:16
=====
```

Notice that the SAP on PE-2 is operationally down and that a flag is set: StandByForMcRing.

The situation is reversed for VLAN 3 since it was configured for path-b; here the SAP on PE-1 is operationally down and the SAP on PE-2 is operationally up.

Ring Failure Verification

In case a ring failure occurs (either ring link failure or ring-node failure), the IB-RCC BFD session between PE-1 and PE-2 will go down and both nodes will put the SAP in operational up state.

Break the link between CE-2 and CE-1.

Observe that the ring is now in a *broken* state:

```
A:PE-1# show redundancy multi-chassis mc-ring peer 192.0.2.2
=====
MC-Ring entries
=====
Sync Tag                Oper State      Failure Reason
-----
l2ring1                 broken         None
-----
No. of MC-Ring entries: 1
=====
```

The following command shows the BSA to ring nodes connections:

```
A:PE-1# show redundancy multi-chassis mc-ring peer 192.0.2.2 ring l2ring1 ring-node
=====
MC-Ring Node entries
=====
Name                    Loc Oper St.   Failure Reason
In Use                 Rem Oper St.
-----
CE-1                    connected      None
Yes                     disconnected
CE-2                    disconnected    None
No                       connected
CE-3                    disconnected    None
No                       connected
CE-4                    disconnected    None
No                       connected
-----
No. of MC-Ring Node entries: 4
=====
```

The output shows that CE-1 is connected to PE-1. CE-2, CE-3 and CE4 are connected to PE-2.

Notice that the SAPs on both PE-1 and PE-2 are now operationally up. This can be checked with the **show service id 2 sap 1/1/1:2.*** command on PE1 and the **show service id 2 sap 1/1/2:2.*** command on PE2.

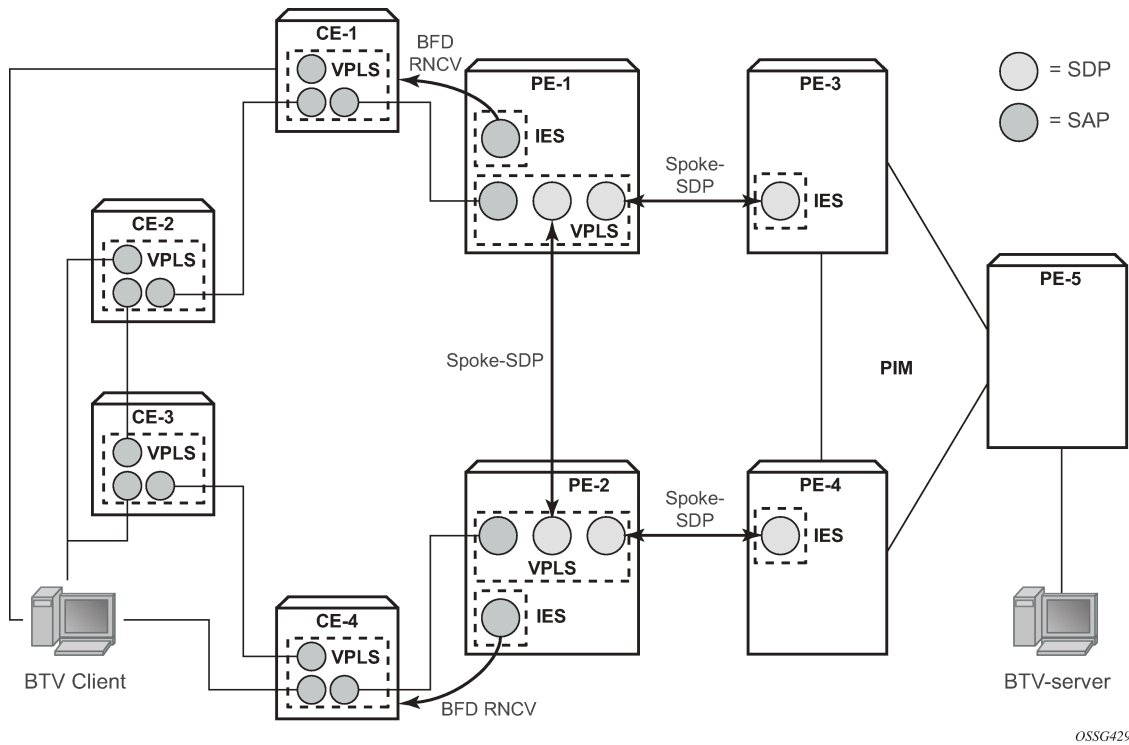
Following the ring failure, the BSA nodes send a MAC address withdrawal message to all the SDP peers configured in the VPLS.

Multicast Service Configuration

In general, BTV (Broadcast TV) services are delivered to the DSLAM on a different VLAN using a different VPLS service. The DSLAM can send/relay IGMP group membership messages if it wants to receive a multicast stream. At the BSA level (PE-1 and PE-2), IGMP snooping is configured. At the BSR level (PE-3,

PE-4) and the PE where the video source is located, PIM is configured and IGMP is configured on the IES service facing the BSA ring. The BSA ring consists of concatenated spoke SDPs. The spoke SDP ring is not closed between PE-3 and PE-4 to avoid a loop. [Figure 144: Multicast Service — Logical Setup](#) shows the logical setup for the multicast service with a MC-ring.

Figure 144: Multicast Service — Logical Setup



Configure an interface to the multicast server on PE-5:

```
configure service
  ies 3 customer 1 create
    interface "mcast-server" create
      address 192.168.7.1/30
      sap 1/1/3:3 create
    exit
  exit
  no shutdown
exit
exit all
```

Configure a VPLS service on PE-1 and PE-2. The configuration on PE-1 is shown below. The VPLS configuration on PE-2 is similar.

```
configure service
  vpls 4 customer 1 create
    description "Multicast VLAN"
    igmp-snooping
    no shutdown
  exit
  sap 1/1/1:4.* create
exit
```

```
        spoke-sdp 12:4 create
            igmp-snooping
            mrouter-port
        exit
    exit
    spoke-sdp 13:4 create
        igmp-snooping
        mrouter-port
    exit
    exit
    no shutdown
    exit
exit all
```

Notice that igmp-snooping has been enabled and that the spoke SDPs are configured as mrouter-ports in order to forward the IGMP joins to the BSRs.

Configure a spoke SDP terminated IES service on PE-3:

```
configure service
    ies 4 customer 1 create
        interface "btv-dst" create
            address 10.0.4.3/24
            ip-mtu 1500
            spoke-sdp 31:4 create
        exit
    exit
    no shutdown
    exit
exit all
```

On PE-4:

```
configure service
    ies 4 customer 1 create
        interface "btv-dst" create
            address 10.0.4.4/24
            ip-mtu 1500
            spoke-sdp 42:4 create
        exit
    exit
    no shutdown
    exit
exit all
```

Notice that also here the **ip-mtu** must be configured to bring the service up. The **ip-mtu** is required to have an MTU match on the spoke SDP between the IES service and the VPLS service.

Configure PIM on PE-3, PE-4 and PE-5. The configuration on PE-3 is shown below. The PIM configuration on PE-4 and PE-5 is similar.

```
configure router pim
    interface "system"
    exit
    interface "int-PE-3-PE-4"
        priority 10
    exit
    interface "int-PE-3-PE-5"
    exit
    interface "btv-dst"
    exit
```

```

rp
  static
    address 192.0.2.5
    group-prefix 224.0.0.0/4
  exit
exit
exit
exit all

```



Note:

PE-5 is statically configured as the RP. This is just an example. Different configurations can be used.

Configure IGMP on PE-3/PE-4:

```
configure router igmp interface btv-dst no shutdown
```

The service should also be configured on all ring nodes. Below, the configuration on CE-2 is shown. Similar configurations are required on the other ring-nodes.

```

configure service
  vpls 4 customer 1 create
    sap 1/1/1:4.* create
    exit
    sap 1/1/2:4.* create
    exit
    sap 1/1/3:4.* create
    exit
    no shutdown
  exit
exit all

```

Multicast Service Verification

Configure the multicast server to send one or more multicast streams. Have the BTV client connected to CE-2 send an IGMP join message for this multicast stream.

On the BSA routers (PE-1/PE-2), IGMP snooping can be checked:

```

A:PE-1# show service id 4 mfib
=====
Multicast FIB, Service 4
=====
Source Address  Group Address      Sap/Sdp Id          Svc Id  Fwd/Blk
-----
*               *               sdp:12:4            Local   Fwd
                  sdp:13:4            Local   Fwd
*               225.1.1.1  sap:1/1/1:4.*       Local   Fwd
                  sdp:12:4            Local   Fwd
                  sdp:13:4            Local   Fwd
-----
Number of entries: 2
=====

```

On PE-3/PE-4 the IGMP groups can be checked:

```
A:PE-3# show router igmp group
```

```
=====
IGMP Groups
=====
(*,225.1.1.1)                               Up Time : 0d 00:01:03
  Fwd List  : btv-dst
-----
(*,G)/(S,G) Entries : 1
=====
```

Following command shows that the MCS peers are synchronized and that there is one IGMP entry on both peers:

```
A:PE-1# show redundancy multi-chassis sync peer 192.0.2.2 detail
=====
Multi-chassis Peer Table
=====
Peer
-----
Peer IP Address      : 192.0.2.2
Description          : (Not Specified)
Authentication       : Disabled
Source IP Address    : 192.0.2.1
Admin State          : Enabled
-----
Sync-status
-----
Client Applications  : IGMP Snooping SUBMGMT RING
Sync Admin State     : Up
Sync Oper State      : Up
DB Sync State        : inSync
Num Entries          : 12
Lcl Deleted Entries  : 0
Alarm Entries        : 0
Rem Num Entries      : 12
Rem Lcl Deleted Entries : 0
Rem Alarm Entries    : 0
=====
MCS Application Stats
=====
Application          : igmp
Num Entries          : 1
Lcl Deleted Entries  : 0
Alarm Entries        : 0
-----
Rem Num Entries      : 1
Rem Lcl Deleted Entries : 0
Rem Alarm Entries    : 0
=====
```

Notice that the MFIB on PE-2 has also been updated:

```
A:PE-2# show service id 4 mfib
=====
Multicast FIB, Service 4
=====
Source Address  Group Address      Sap/Sdp Id          Svc Id  Fwd/Blk
-----
*               *               sdp:21:4            Local    Fwd
                  sdp:24:4            Local    Fwd
*               225.1.1.1    sap:1/1/2:4.*       Local    Fwd
                  sdp:21:4            Local    Fwd
                  sdp:24:4            Local    Fwd
```



```
-----
Number of entries: 2
=====
A:PE-2#
```

The SAP on PE-2 is down to avoid duplicated traffic and loops:

```
A:PE-2# show service id 4 sap 1/1/2:4.*
=====
Service Access Points(SAP)
=====
Service Id      : 4
SAP             : 1/1/2:4.*          Encap           : qinq
QinQ Dot1p     : Default
Description     : (Not Specified)
Admin State    : Up                 Oper State      : Down
Flags          : StandByForMcRing
Multi Svc Site : None
Last Status Change : 11/06/2009 18:56:25
Last Mgmt Change  : 11/04/2009 23:43:16
=====
```

If a failure occurs in the ring-node, the IB-RCC BFD session between PE-1 and PE-2 will go down and the SAP on both PE-1 and PE-2 will be put in operational up state.

Configuration Notes

RNCV (used for ring-node connectivity check) and BFD (used for ring control connection) can either run on the same VLAN in the same IES or VPRN service or can run on different VLAN.

MCS for IGMP or DHCP states on a MC-ring requires ESM since the BSA nodes must know which ring node a subscriber is connected to in case a ring failure occurs. The ring-node name is returned through a Python script, a RADIUS server or through a local user database. This string (int_dest_id) must match one of the ring nodes defined in the redundancy configuration.

Convergence time after a ring failure should be 3 * BFD timer + MCS convergence time. Convergence time after a BSA failure should be likewise.

Note that a debounce timer runs on the MC-ring peers. After a ring failure, the MC-ring converges immediately (after the BFD session times out) and the debounce timer is started. After the ring is fixed and the BFD session is up the MC-ring converges immediately again. If another ring failure occurs before the debounce timer expires, convergence will be slowed down by two (2) seconds. If a third ring failure occurs before the debounce timer expires, four second delays are introduced. In case of a fourth failure, an eight second delay is introduced. 200 seconds of delay is the maximum. The debounce timer can be configured under the MC-ring.

Conclusion

This chapter covers an extension of dual homing support in TPSDA networks based on Layer 2 CO model. The extension addresses networks where multiple access nodes (for example, DSLAMs) are connected in a single ring. The examples show the use of a ring with four access nodes in a ring. The behavior is described in normal operation and in case a failure occurs in the access ring.

Python Cache Support for ESM Applications

This chapter provides information about Python cache support for ESM applications.

Topics in this chapter include:

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

Applicability

The information and configuration in this chapter is based on SR OS Release 12.0.R4.

Overview

SR OS sports an embedded Python scripting engine which can be used to manipulate selected messages of protocols including DHCPv4, DHCPv6, RADIUS and Diameter.

Python cache provides a central key-value memory cache with a set of APIs allowing Python scripts to store and retrieve strings across different runs of the same or even different Python scripts.

The following are some basic concepts of Python cache:

- Multiple Python policies can be defined, and each Python policy has a separate cache; only scripts configured in the Python policy can access and share that policy's cache.
- A cache entry consists of a key and a value, both of them being strings. The key is used to search and fetch the cache entry.
- A cache entry has a lifetime and the system automatically removes expired entries.


The following are the Python cache APIs:

- `alc.cache.save(val,key)`: Save the value identified by the key into the cache.
- `alc.cache.retrieve(key)`: Return the cached entry's value identified by the key.
- `alc.cache.clear(key)`: Remove the entry identified by the key from the cache.
- `cache.get_lifetime(key)`: The system returns an integer indicating the remaining lifetime of the specified entry expressed in seconds.
- `cache.set_lifetime(key,new_lifetime)`: `new_lifetime` is an integer; the system sets the remaining lifetime of the specified entry (in seconds).

Configuration

The test topology is shown in [Figure 145: Test topology](#).

Figure 145: Test topology



The example shows how the Python cache can be used to store multiple class attributes from the RADIUS access-accept packets and reflect them into RADIUS accounting request packets (start/interim-update/stop).

Test setup:

- A PC is used as DHCPv4 host, connect to SAP 1/1/9 of BNG-1
- SAP 1/1/9 is attached to group interface grp-int-1, which is under subscriber interface sub-int-1 of IES 1
- DHCP host is authenticated via RADIUS server, which resides in management routing instance of BNG-1
- A Python script python-script-1 stores all class attributes in access-accept and add stored attributes into RADIUS accounting requests.
- RADIUS user-name is used as cache key, which is the MAC address of the PC.

The Python cache configuration commands are shown below.

```
config>python>py-policy>
  [no] cache [create]
    [no] entry-size <size>
    [no] max-entries <count>
    [no] max-entry-lifetime [days <days>] [hrs <hours>]
                                [min <minutes>] [sec <seconds>]
    [no] mcs-peer <ip-address> sync-tag <[32 chars max]>
    [no] minimum-lifetimes
    [no] high-availability <seconds>
    [no] multi-chassis-redundancy <seconds>
    [no] persistence <seconds>
  [no] persistence
  [no] shutdown

config>system>persistence>
  python-policy-cache
    [no] description <desc>
    [no] location <cflash-id>

config>redundancy>multi-chassis>peer>sync>
  [no] python
```

Refer to the *7450 ESS, 7750 SR, and VSR Triple Play Service Delivery Architecture Guide* for details of above commands. The basic configuration of the Python cache is the **cache create** statement in the python-policy as shown below:

```
config>python>py-policy>
-----
    cache create
        no shutdown
    exit
    radius access-accept direction ingress script "python-script-1"
    radius accounting-request direction egress script "python-script-1"
-----
```

The **cache create** configuration enables the cache support for the Python policy. The system's behavior can be tuned in the following aspects:

- Configure **entry-size** and **max-entries** to limit the memory usage.
- Configure **max-entry-lifetime** to specify the maximum lifetime.
- Enable **persistence** to make cache entries persistent across reboot.
- Configure **mcs-peer** to enable Multi-Chassis Synchronization.

In this example only the basic cache configuration is used.

Step 0. Configuring ESM

As a prerequisite ESM must be enabled and as such BNG-1 is configured as follows:

- An **authentication-policy radius-auth-policy-1** is used to authenticate DHCPv4 hosts on group interface grp-int-1.
- A **radius-accounting-policy radius-acct-policy-1** is configured in the **sub-profile sub-profile-1** to enable RADIUS accounting.
- The **user-name** is included in *radius-acct-policy-1* since the User-Name is used as cache entry key.
- Interim-update is enabled in the *radius-acct-policy-1* with an interval 5 minutes.
- The local DHCPv4 server *dhcpv4-svr* is used to assign address to hosts.

```
#-----
echo "Management Router Configuration"
#-----
    router management
        radius-server
            server "radius-svr-1" address 192.168.1.1 secret
                                "iaCuILBunKirJurE4jK2URAnzip6nK32" hash2 create
        exit
    exit
exit
#-----
echo "Router (Network Side) Configuration"
#-----
    router
        dhcp
            local-dhcp-server "dhcpv4-svr" create
        exit
    exit
    interface "system"
        address 192.0.2.1/32
        local-dhcp-server "dhcpv4-svr"
```

```
        no shutdown
    exit
exit

#-----
echo "Subscriber-mgmt Configuration"
#-----
    subscriber-mgmt
        authentication-policy "radius-auth-policy-1" create
            password "mcgLj0q0695Dp.pD5DthrCv9Bu8X2qPVSvGYWQmCgUg" hash2
            radius-server-policy "radius-svr-policy-1"
        exit
        radius-accounting-policy "radius-acct-policy-1" create
            update-interval 5
            update-interval-jitter absolute 0
            include-radius-attribute
                user-name
            exit
            radius-server-policy "radius-svr-policy-1"
        exit
        sla-profile "sla-profile-1" create
        exit
        sub-profile "sub-profile-1" create
            radius-accounting-policy "radius-acct-policy-1"
        exit
    exit
#-----
echo "Service Configuration"
#-----
    service
        ies 1 customer 1 create
            subscriber-interface "sub-int-1" create
                address 10.0.0.254/24
                group-interface "grp-int-1" create
                    dhcp
                        server 192.0.2.1
                        gi-address 10.0.0.254
                        no shutdown
                    exit
                    authentication-policy "radius-auth-policy-1"
                sap 1/1/9 create
                    sub-sla-mgmt
                        def-sub-id use-auto-id
                        def-sub-profile "sub-profile-1"
                        def-sla-profile "sla-profile-1"
                        no shutdown
                    exit
                exit
            exit
        exit
        exit
        no shutdown
    exit
exit
#-----
echo "Local DHCP Server (Base Router) Configuration"
#-----
    router
        dhcp
            local-dhcp-server "dhcpv4-svr" create
                use-gi-address
                pool "addr-pool-1" create
                subnet 10.0.0.0/24 create
                options
                    subnet-mask 255.255.255.0
```

```

                                default-router 10.0.0.254
                                exit
                                address-range 10.0.0.1 10.0.0.100
                                exit
                                exit
                                no shutdown
                                exit
                                exit
                                exit
                                exit
                                exit
                                #-----
                                echo "AAA Configuration"
                                #-----
                                aaa
                                radius-server-policy "radius-svr-policy-1" create
                                servers
                                router "management"
                                server 1 name "radius-svr-1"
                                exit
                                exit
                                exit
                                exit

```

1. Create the Python script file.

A Python script is created and stored on the local storage, for example as CF3:\python_cache.py. This script handles the RADIUS packets listed below:

- Access-Accept — All class attributes from the access-accept packets are stored in the cache by combining them into a single string. The user-name (MAC address) is used as the key, and the format of this string is:
 - 1st byte is the number of class attributes in this string
 - 2nd – nth bytes: each byte holds the number of bytes for class-n attributes
 - Rest of bytes: combined string of class-1 to class-n
- Acct-Start — Retrieves the stored combined class string using the user name as key. The combined string is parsed and split into the individual class attributes and then inserted into the packet.
- Interim-Update — The cached entry is parsed and the stored class attributes are inserted, then the lifetime is reset to 15 minutes. This is greater than the interim-update interval (5 minutes) so the cached entry will not expire before next interim-update arrives.
- Acct-Stop — The cached entry is parsed and the stored class attributes are inserted, then the cached entry is removed.



Note:

There is some error checking and exception handling logic in the script which causes the script to drop the packet if certain errors occur. This script is only an overview example so the error handling logic should be added according to real application requirements. In addition to the exception handling logic included in the script, the **action-on-fail** command could be used in the **python-script** command to define system's action upon failed execution, for example when an un-captured exception is encountered.

```

#-----
# Name:      Nokia SR OS Python Cache Support Example
# Purpose:   This script is used to demonstrate SR OS python cache support for
#            the following use case:
#            RADIUS sever returns multiple Class attributes in
#            access-accept, these Class attributes need to be reflected
#            in accounting request packets
#-----

```

```
#-----
from alc import cache
from alc import radius
import struct
def main():
    radius_header = radius.header()
    if radius_header['code'] == 2: # in case of access-accept
        entry_key = radius.attributes.get(1) # use user-name as the cache entry
                                           # key
        if entry_key == "": #drop the packet if there is no user-name present
            radius.drop()
            return
        class_list = radius.attributes.getTuple(25) # get a list of Class
                                                    # attributes
        if class_list == (): #drop the packet if there is no class present
            radius.drop()
            return
        class_len_str = ''
        class_str = ''
        class_count = 0
        for radius_class in class_list:
            class_len_str += chr(len(radius_class))
            class_str += radius_class
            class_count += 1
        entry_val = chr(class_count)+class_len_str+class_str
        try:
            cache.save(entry_val,entry_key)
        except:
            radius.drop() #drop the packet if cache.save fails
            return

    elif radius_header['code'] == 4: # in case of acct-request
        entry_key = radius.attributes.get(1)
        if entry_key == "": #drop the packet if there is no user-name present
            radius.drop()
            return
        try:
            entry_val = cache.retrieve(entry_key)
        except:#drop the packet if cache.retrieve fails
            radius.drop()
            return
        class_count = ord(entry_val[0])
        pos = class_count+1
        class_list = []
        for i in range(class_count):
            class_len = ord(entry_val[i+1])
            class_list.append(entry_val[pos:pos+class_len])
            pos += class_len
        radius.attributes.set(25,tuple(class_list))
        acct_type=struct.unpack('>I',radius.attributes.get(40))[0]
        if acct_type == 2: # in case of acct-stop
            cache.clear(entry_key) # remove the cache entry
        elif acct_type == 3: # in case of interim-update
            try:
                cache.set_lifetime(entry_key,900) # reset the lifetime
            except:#drop the packet if cache.set_lifetime fails
                radius.drop()
                return

    main()
```

2. Configure **python-script** and **python-policy**.

- Enable *python-script-1* to process received (ingress) access-accept packets and transmitted (egress) accounting-request packets.
- Enable the Python cache by configuring **cache create** in *python-policy-1*.
- Reference *python-policy-1* in the *radius-server-policy-1*.

```
#-----  
echo "PYTHON Configuration"  
#-----  
python  
    python-script "python-script-1" create  
        primary-url "cf3:/python_cache.py"  
        no shutdown  
    exit  
    python-policy "python-policy-1" create  
        cache create  
        no shutdown  
    exit  
    radius access-accept direction ingress script "python-script-1"  
    radius accounting-request direction egress script "python-script-1"  
    exit  
exit  
#-----  
echo "AAA Configuration"  
#-----  
aaa  
    radius-server-policy "radius-svr-policy-1" create  
        python-policy "python-policy-1"
```

3. Configure the RADIUS server so that:

- The DHCP host is authenticated via its MAC address.
- The RADIUS server returns the following class attributes and values in access-accept packets:
 - "Class-1"
 - "Class-22"
 - "Class-333"

4. Enable debug on BNG-1 to observe the Python cache in action.

- Enable the following debug on BNG-1:

```
debug  
    router "Base"  
        ip  
            dhcp  
                detail-level low  
                mode egr-ingr-and-dropped  
            exit  
        exit  
    exit  
    router "management"  
        radius  
            packet-type authentication accounting coa  
            detail-level medium  
        exit  
    exit  
python
```



```
python-script "python-script-1"
  script-all-info
  exit
exit
exit

A:BNG-1>config>log# info
-----
  log-id 10
    from debug-trace
    to session
  exit
-----
```

5. Initiate DHCPv4 on the PC.

- Initiate the DHCPv4 process on the PC. When the PC sends out a DHCPv4 discover message, BNG-1 contacts the RADIUS server to authenticate the user and a DHCPv4 ESM host is created.
- RADIUS accounting-start will be sent upon host creation and interim-update will be sent every 5 minutes.

The following describes the debug output:

- The system initiates RADIUS authentication upon receipt of a DHCP discovery message.
- The RADIUS server returns an access-accept with the three class attributes.
- The system executes python-script-1, stores the three class attributes.
- After the DHCP host is created, the system sends RADIUS accounting-start and interim-update messages in which python-script-1 adds the three stored class attributes.

```
26 2014/06/26 03:02:18.34 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 3 (grp-int-1),
  received DHCP Boot Request on Interface grp-int-1 (1/1/9) Port 67

H/W Type: Ethernet(10Mb)  H/W Address Length: 6
ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
siaddr: 0.0.0.0          giaddr: 0.0.0.0
chaddr: 00:20:fc:1e:cd:53  xid: 0x1e866b74
"

27 2014/06/26 03:02:18.34 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Transmit
  Access-Request(1) 192.168.1.1:1812 id 7 len 63 vrid 4095 pol radius-svr-policy -1
    USER NAME [1] 17 00:20:fc:1e:cd:53
    PASSWORD [2] 16 R/mc867ChKkdx50PJauB5U
    NAS IP ADDRESS [4] 4 192.168.1.254
"

28 2014/06/26 03:02:18.34 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Receive
  Access-Accept(2) id 7 len 69 from 192.168.1.1:1812 vrid 4095 pol radius-svr-policy-1
    CLASS [25] 7 0x436c6173732d31
    CLASS [25] 8 0x436c6173732d3232
    CLASS [25] 9 0x436c6173732d333333
    USER NAME [1] 17 00:20:fc:1e:cd:53
"

29 2014/06/26 03:02:18.34 UTC MINOR: DEBUG #2001 Base Python Output
```

```
"Python Output: python-script-1
"

30 2014/06/26 03:02:18.34 UTC MINOR: DEBUG #2001 Base Python Result
"Python Result: python-script-1
"

31 2014/06/26 03:02:18.34 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Script
Access-Accept(2) id 7 len 69 from 192.168.1.1:1812 policy python-policy-1 stat us success
"

32 2014/06/26 03:02:18.34 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base),
  transmitted DHCP Boot Request to 192.0.2.1 Port 67

  H/W Type: Ethernet(10Mb)  H/W Address Length: 6
  ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
  siaddr: 0.0.0.0           giaddr: 10.0.0.254
  chaddr: 00:20:fc:1e:cd:53  xid: 0x1e866b74
"

33 2014/06/26 03:02:18.34 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base),
  received DHCP Boot Reply on 192.0.2.1 Port 67

  H/W Type: Ethernet(10Mb)  H/W Address Length: 6
  ciaddr: 0.0.0.0           yiaddr: 10.0.0.1
  siaddr: 192.0.2.1         giaddr: 10.0.0.254
  chaddr: 00:20:fc:1e:cd:53  xid: 0x1e866b74
"

34 2014/06/26 03:02:18.34 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 3 (grp-int-1),
  transmitted DHCP Boot Reply to Interface grp-int-1 (1/1/9) Port 68

  H/W Type: Ethernet(10Mb)  H/W Address Length: 6
  ciaddr: 0.0.0.0           yiaddr: 10.0.0.1
  siaddr: 192.0.2.1         giaddr: 10.0.0.254
  chaddr: 00:20:fc:1e:cd:53  xid: 0x1e866b74
"

35 2014/06/26 03:02:18.34 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 3 (grp-int-1),
  received DHCP Boot Request on Interface grp-int-1 (1/1/9) Port 67

  H/W Type: Ethernet(10Mb)  H/W Address Length: 6
  ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
  siaddr: 0.0.0.0           giaddr: 0.0.0.0
  chaddr: 00:20:fc:1e:cd:53  xid: 0x1e866b74
"

36 2014/06/26 03:02:18.35 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base),
  transmitted DHCP Boot Request to 192.0.2.1 Port 67

  H/W Type: Ethernet(10Mb)  H/W Address Length: 6
  ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
  siaddr: 0.0.0.0           giaddr: 10.0.0.254
```

```
    chaddr: 00:20:fc:1e:cd:53    xid: 0x1e866b74
"

37 2014/06/26 03:02:18.35 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base),
    received DHCP Boot Reply on 192.0.2.1 Port 67

    H/W Type: Ethernet(10Mb)  H/W Address Length: 6
    ciaddr: 0.0.0.0            yiaddr: 10.0.0.1
    siaddr: 192.0.2.1          giaddr: 10.0.0.254
    chaddr: 00:20:fc:1e:cd:53    xid: 0x1e866b74
"

38 2014/06/26 03:02:18.38 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 3 (grp-int-1),
    transmitted DHCP Boot Reply to Interface grp-int-1 (1/1/9) Port 68

    H/W Type: Ethernet(10Mb)  H/W Address Length: 6
    ciaddr: 0.0.0.0            yiaddr: 10.0.0.1
    siaddr: 192.0.2.1          giaddr: 10.0.0.254
    chaddr: 00:20:fc:1e:cd:53    xid: 0x1e866b74
"

39 2014/06/26 03:02:18.38 UTC MINOR: DEBUG #2001 Base Python Output
"Python Output: python-script-1
"

40 2014/06/26 03:02:18.38 UTC MINOR: DEBUG #2001 Base Python Result
"Python Result: python-script-1
RADIUS Attribute: Type 25, SET
    'Class-1'
    'Class-22'
    'Class-333'
"

41 2014/06/26 03:02:18.38 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Transmit
    Accounting-Request(4) 192.168.1.1:1813 id 8 len 152 vrid 4095 pol radius-svr-policy-1
    STATUS TYPE [40] 4 Start(1)
    NAS IP ADDRESS [4] 4 192.168.1.254
    USER NAME [1] 17 00:20:fc:1e:cd:53
    SESSION ID [44] 63 00:20:fc:1e:cd:53|1/1/9@1/1/9@sla-profile-1_2014/06/26 03
:02:18
    EVENT TIMESTAMP [55] 4 1403751738
    CLASS [25] 7 0x436c6173732d31
    CLASS [25] 8 0x436c6173732d3232
    CLASS [25] 9 0x436c6173732d333333
"

42 2014/06/26 03:02:18.38 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Receive
    Accounting-Response(5) id 8 len 20 from 192.168.1.1:1813 vrid 4095 pol radius-svr-
policy-1
"

43 2014/06/26 03:02:18.38 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Script
    Accounting-Response(5) id 8 len 20 from 192.168.1.1:1813 policy python-policy-1 status
success
"

44 2014/06/26 03:07:18.71 UTC MINOR: DEBUG #2001 Base Python Output
"Python Output: python-script-1
```

```
"
45 2014/06/26 03:07:18.71 UTC MINOR: DEBUG #2001 Base Python Result
"Python Result: python-script-1
RADIUS Attribute: Type 25, SET
    'Class-1'
    'Class-22'
    'Class-333'
"

46 2014/06/26 03:07:18.71 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Transmit
  Accounting-Request(4) 192.168.1.1:1813 id 9 len 302 vrid 4095 pol radius-svr-policy-1
    STATUS TYPE [40] 4 Interim-Update(3)
    NAS IP ADDRESS [4] 4 192.168.1.254
    USER NAME [1] 17 00:20:fc:1e:cd:53
    SESSION ID [44] 63 00:20:fc:1e:cd:53|1/1/9@1/1/9@sla-profile-1_2014/06/26 03
:02:18
    SESSION TIME [46] 4 300
    EVENT TIMESTAMP [55] 4 1403752038
    VSA [26] 12 Alcatel(6527)
      INPUT_INPROF_OCTETS_64 [19] 10 0x00010000000000000000
    VSA [26] 12 Alcatel(6527)
      INPUT_OUTPROF_OCTETS_64 [20] 10 0x00010000000000000000
    VSA [26] 12 Alcatel(6527)
      INPUT_INPROF_PACKETS_64 [23] 10 0x00010000000000000000
    VSA [26] 12 Alcatel(6527)
      INPUT_OUTPROF_PACKETS_64 [24] 10 0x00010000000000000000
    VSA [26] 12 Alcatel(6527)
      OUTPUT_INPROF_OCTETS_64 [21] 10 0x000100000000000000bea
    VSA [26] 12 Alcatel(6527)
      OUTPUT_OUTPROF_OCTETS_64 [22] 10 0x00010000000000000000
    VSA [26] 12 Alcatel(6527)
      OUTPUT_INPROF_PACKETS_64 [25] 10 0x00010000000000000019
    VSA [26] 12 Alcatel(6527)
      OUTPUT_OUTPROF_PACKETS_64 [26] 10 0x00010000000000000000
    CLASS [25] 7 0x436c6173732d31
    CLASS [25] 8 0x436c6173732d3232
    CLASS [25] 9 0x436c6173732d333333
"

47 2014/06/26 03:07:18.71 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Receive
  Accounting-Response(5) id 9 len 20 from 192.168.1.1:1813 vrid 4095 pol radius-svr-
policy-1
"

48 2014/06/26 03:07:18.71 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Script
  Accounting-Response(5) id 9 len 20 from 192.168.1.1:1813 policy python-policy-1 status
success
"

44 2014/06/26 03:07:18.71 UTC MINOR: DEBUG #2001 Base Python Output
"Python Output: python-script-1
"

45 2014/06/26 03:07:18.71 UTC MINOR: DEBUG #2001 Base Python Result
"Python Result: python-script-1
RADIUS Attribute: Type 25, SET
    'Class-1'
    'Class-22'
    'Class-333'
"
```

```

46 2014/06/26 03:07:18.71 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Transmit
  Accounting-Request(4) 192.168.1.1:1813 id 9 len 302 vrid 4095 pol radius-svr-policy-1
    STATUS TYPE [40] 4 Interim-Update(3)
    NAS IP ADDRESS [4] 4 192.168.1.254
    USER NAME [1] 17 00:20:fc:1e:cd:53
    SESSION ID [44] 63 00:20:fc:1e:cd:53|1/1/9@1/1/9@sla-profile-1_2014/06/26 03
:02:18
  SESSION TIME [46] 4 300
  EVENT TIMESTAMP [55] 4 1403752038
  VSA [26] 12 Alcatel(6527)
    INPUT_INPROF_OCTETS_64 [19] 10 0x00010000000000000000
  VSA [26] 12 Alcatel(6527)
    INPUT_OUTPROF_OCTETS_64 [20] 10 0x00010000000000000000
  VSA [26] 12 Alcatel(6527)
    INPUT_INPROF_PACKETS_64 [23] 10 0x00010000000000000000
  VSA [26] 12 Alcatel(6527)
    INPUT_OUTPROF_PACKETS_64 [24] 10 0x00010000000000000000
  VSA [26] 12 Alcatel(6527)
    OUTPUT_INPROF_OCTETS_64 [21] 10 0x000100000000000000bea
  VSA [26] 12 Alcatel(6527)
    OUTPUT_OUTPROF_OCTETS_64 [22] 10 0x00010000000000000000
  VSA [26] 12 Alcatel(6527)
    OUTPUT_INPROF_PACKETS_64 [25] 10 0x00010000000000000019
  VSA [26] 12 Alcatel(6527)
    OUTPUT_OUTPROF_PACKETS_64 [26] 10 0x00010000000000000000
  CLASS [25] 7 0x436c6173732d31
  CLASS [25] 8 0x436c6173732d3232
  CLASS [25] 9 0x436c6173732d333333
"

47 2014/06/26 03:07:18.71 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Receive
  Accounting-Response(5) id 9 len 20 from 192.168.1.1:1813 vrid 4095 pol radius-svr-
policy-1
"

48 2014/06/26 03:07:18.71 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Script
  Accounting-Response(5) id 9 len 20 from 192.168.1.1:1813 policy python-policy-1 status
success
"

```

As the debug output shows, *python-script-1* stores the three class attributes from the access-accept message which then are reflected into the accounting-start and interim-update messages.

The **tools dump python python-policy <name> cache** command can be used to show the existing cached entries in the specified python-policy:

```

A:BNG-1# tools dump python python-policy "python-policy-1" cache
=====
Python policy cache "python-policy-1" entries
=====
Key       : 00:20:fc:1e:cd:53
Value     : (hex) 03 07 08 09 43 6c 61 73 73 2d 31 43 6c 61 73 73 2d 32 32 43 6c 61 73 73 2d
33 33 33
Time Left : 0d 00:12:08
DDP Key   : N/A
=====

```

6. Release DHCPv4 lease on the PC

- Release DHCPv4 lease on the PC which is sent to BNG-1.

- DHCPv4 release message from the PC will trigger BNG-1 to remove the ESM host on BNG-1.
- BNG-1 will send an accounting-stop packet to the RADIUS server.

The following is the debug output:

```
"Python Output: python-script-1
"

67 2014/06/26 03:26:14.84 UTC MINOR: DEBUG #2001 Base Python Result
"Python Result: python-script-1
RADIUS Attribute: Type 25, SET
    'Class-1'
    'Class-22'
    'Class-333'
"

68 2014/06/26 03:26:14.85 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Transmit
    Accounting-Request(4) 192.168.1.1:1813 id 13 len 308 vrid 4095 pol radius-svr-
policy-1
    STATUS TYPE [40] 4 Stop(2)
    NAS IP ADDRESS [4] 4 192.168.1.254
    USER NAME [1] 17 00:20:fc:1e:cd:53
    SESSION ID [44] 63 00:20:fc:1e:cd:53|1/1/9@1/1/9@sla-profile-1_2014/06/26 03
:02:18
    SESSION TIME [46] 4 1436
    TERMINATE CAUSE [49] 4 User Request(1)
    EVENT TIMESTAMP [55] 4 1403753174
    VSA [26] 12 Alcatel(6527)
        INPUT_INPROF_OCTETS_64 [19] 10 0x00010000000000000000
    VSA [26] 12 Alcatel(6527)
        INPUT_OUTPROF_OCTETS_64 [20] 10 0x00010000000000000000
    VSA [26] 12 Alcatel(6527)
        INPUT_INPROF_PACKETS_64 [23] 10 0x00010000000000000000
    VSA [26] 12 Alcatel(6527)
        INPUT_OUTPROF_PACKETS_64 [24] 10 0x00010000000000000000
    VSA [26] 12 Alcatel(6527)
        OUTPUT_INPROF_OCTETS_64 [21] 10 0x0001000000000000001a36
    VSA [26] 12 Alcatel(6527)
        OUTPUT_OUTPROF_OCTETS_64 [22] 10 0x00010000000000000000
    VSA [26] 12 Alcatel(6527)
        OUTPUT_INPROF_PACKETS_64 [25] 10 0x0001000000000000000037
    VSA [26] 12 Alcatel(6527)
        OUTPUT_OUTPROF_PACKETS_64 [26] 10 0x00010000000000000000
    CLASS [25] 7 0x436c6173732d31
    CLASS [25] 8 0x436c6173732d3232
    CLASS [25] 9 0x436c6173732d333333
"

69 2014/06/26 03:26:14.85 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Receive
    Accounting-Response(5) id 13 len 20 from 192.168.1.1:1813 vrid 4095 pol radius
-svr-policy-1
"

70 2014/06/26 03:26:14.85 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Script
    Accounting-Response(5) id 13 len 20 from 192.168.1.1:1813 policy python-policy
-1 status success
"
```

As the debug output shows, *python-script-1* inserts three RADIUS class attributes in the accounting-stop message.

The **tools dump python python-policy <name> cache** command can be used to verify the cached entry has been removed:

```
A:BNG-1# tools dump python python-policy "python-policy-1" cache
=====
Python policy cache "python-policy-1" entries
=====
```

7. Manually change the lifetime of a cached entry (optional).

A **tools** command can be used to manually change the lifetime of an existing cached entry, with following syntax:

```
tool perform python-policy <name> cache {hex-key <hex-str>|string-key <str>} set-lifetime
<newlifetime>
```

However, manually changing the lifetime might cause issues with the Python script (for example, if reducing the lifetime causes the entry to expire) that needs the cached entry so it should be used with caution.

To demonstrate this recreate the cached entry by initiating a DHCPv4 discover from the PC (see Step 5, 5), then change the lifetime to 20 minutes.

```
A:BNG-1# tools dump python python-policy "python-policy-1" cache
=====
Python policy cache "python-policy-1" entries
=====
Key       : 00:20:fc:1e:cd:53
Value     : (hex) 03 07 08 09 43 6c 61 73 73 2d 31 43 6c 61 73 73 2d 32 32 43 6c 61 73 73 2d
33 33 33
Time Left : 0d 00:09:48
DDP Key   : N/A
=====
A:BNG-1# tools perform python-policy "python-policy-1" cache string-key "00:20:fc:1e:cd:53"
set-lifetime 1200
A:BNG-1# tools dump python python-policy "python-policy-1" cache
=====
Python policy cache "python-policy-1" entries
=====
Key       : 00:20:fc:1e:cd:53
Value     : (hex) 03 07 08 09 43 6c 61 73 73 2d 31 43 6c 61 73 73 2d 32 32 43 6c 61 73 73 2d
33 33 33
Time Left : 0d 00:19:57
DDP Key   : N/A
=====
```

8. Manually remove a cached entry (optional).

Manually removing an existing cached entry can be done using a clear command with following syntax:

```
clear python python-policy <name> cache {hex-key <hex-str>|string-key <str>}
```

Manually removing a cached entry can result in unexpected results (for example, if a script expects an entry to exist but it has been removed), so this should be used with caution.

The following command sequence demonstrates the effect of the clear command after initiating a DHCPv4 discover from the PC to recreate the cached entry (see Step 5, 5).

```
A:BNG-1# tools dump python python-policy "python-policy-1" cache

=====
Python policy cache "python-policy-1" entries
=====
Key       : 00:20:fc:1e:cd:53
Value     : (hex) 03 07 08 09 43 6c 61 73 73 2d 31 43 6c 61 73 73 2d 32 32 43 6c 61 73 73 2d
           33 33 33
Time Left : 0d 00:11:39
DDP Key   : N/A
=====

A:BNG-1# clear python python-policy "python-policy-1" cache string-key "00:20:fc:1e:cd:53"
A:BNG-1# tools dump python python-policy "python-policy-1" cache

=====
Python policy cache "python-policy-1" entries
=====
```

Configuration and Operational Guidelines

The following is a list of configuration and operational guidelines that a user should follow when using the Python cache:

- SR OS has a limit on the total amount of memory used for the python cache since the python cache can be demanding with respect to memory usage. The maximum memory allocated for the cache system wide is restricted to 256MB. However, it is good practice to configure per python-policy limits using the **entry-size** and **max-entries** commands; by doing this, one python-policy's cache will not impact another python-policy's cache memory usage.
- For applications needing a cache entry for the entire lifetime of an ESM host, lifetime management is essential. If the lifetime is too long then unneeded entries might reside in the system, wasting memory; if the lifetime is too short then entries might expire while they are still needed. One way to address this is by initially setting a relative short lifetime and then using the RADIUS interim-update message as a trigger to reset a new lifetime. This new lifetime should be larger than the interim-update interval. Then the entry should be removed by a script when an accounting-stop message is sent.
- With MCS enabled, each python cached entry will have a corresponding MCS record, resulting in each python cache entry consuming twice amount of memory. For example, 256MB cached entries would consume additional 256MB memory for MCS records.
- Choosing the right key is important, a network designer needs to choose the key meeting the application requirement in terms of their uniqueness. For example: if an application needs to store per-host information then the key for that host must be unique (for example, its MAC address or remote-id). The key must be derived from the trigger packet. For example: if an application needs to store information on DHCP discovery and retrieve it on receiving a RADIUS accounting-request message for the same host, then the script needs to be able derive the same key from both the DHCP discovery as well as from the RADIUS accounting-request message.
- Using tools or clear commands to manually change cached entries could cause problems if the entries are needed by a script. Only use these commands when it is absolutely necessary.

- The minimum-lifetimes exist to make the system more efficient in handling system memory before a cache entry could be synced or made persistent (e.g. when a new entry is created or when MCS is enabled). The cache entry's lifetime must be equal or larger than the configured minimal-lifetime listed below for that function to occur.
 1. high-availability — The minimum lifetime of a cache entry for it to be synchronized from the active CPM to the standby CPM.

The default is 0 seconds, resulting in all entries being synchronized.
 2. multi-chassis-redundancy — The minimum lifetime of a cache entry for it to be synchronized between chassis.

The default is 0 seconds, resulting in all entries being synchronized.
 3. persistence — The minimum lifetime of a cache entry for it to be written to the persistency file.

The default is 0 seconds, resulting in all entries will be synchronized.

Conclusion

The Python cache provides a very powerful and flexible way to share information across different Python scripts in SR OS.

RADIUS-Triggered Dynamic Data Service Provisioning

This chapter describes advanced RADIUS-triggered dynamic data service provisioning configurations.

Topics in this chapter include:

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

Applicability

The information and configuration are based on SR OS Release 11.0.R2.

Overview

RADIUS-triggered dynamic data services enables a zero touch, single-ended provisioning model for business services on the basis of Enhanced Subscriber Management functionality.

Triggered by the authentication of a single or dual stack PPPoE or IPoE session or a single stack IPv4 host as the "control channel" from the business CPE, parameters are passed in a RADIUS Access Accept or Change of Authorization (CoA) message to set up one or multiple Layer 2 or Layer 3 data services.

This concept removes the need to have an Operations Support System (OSS) responsible for the service provisioning and is particularly beneficial in a highly dynamic network environment, where physical network topologies – especially in the access – change frequently. With a regular service provisioning, frequent changes would be hard to keep track of. In the RADIUS-based model the service gets instantiated wherever it "pops-up" in the network. Even planned customer moves to a different office would not require advanced notifications and lead times but could be instantaneous, assuming the pure physical connectivity is given.

A variation of the current service offering will only require one or a few modified service parameters in the RADIUS user database and does not require timely and costly IT changes (for RADIUS those service parameters are just attributes; the RADIUS server does not check the logic). This speeds up the time-to-market for new service offerings, which is another big advantage.

Taking this logic to its full extent, it becomes immediately clear that the managed business CPE terminating the carrier service (and being responsible for the PPP or DHCP control channel) also needs to be provisioned in the most flexible way. Through the control channel or a dedicated management channel instantiated as the first dynamic data service, the business CPE should get its full configuration from a configuration server via a pre-populated configuration file. The details of the CPE provisioning are outside of the scope of this example and therefore not discussed further.

As the whole approach is centered around the principle of "highly flexible in a highly dynamic environment", it is naturally required to maintain as little state information about connections in the RADIUS parameter attributes as possible. For example, fixed remote peer IP-addresses for the SDPs used in a VPLS service

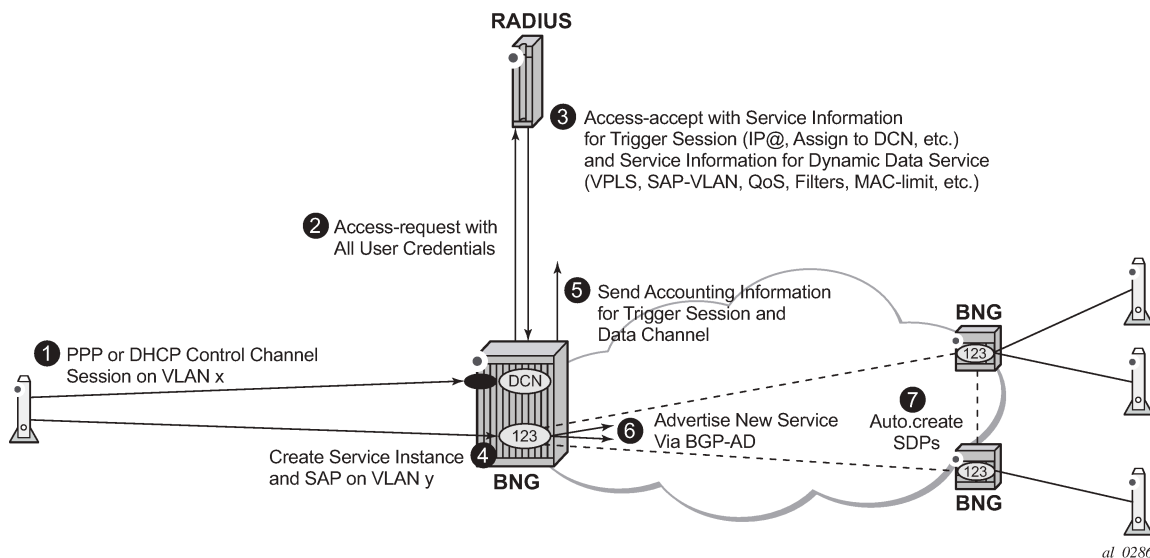
in the RADIUS parameter lists would remove all the flexibility and would not allow access services to be moved dynamically. As such the data services for this functionality focuses on those types where a control protocol like BGP is used to exchange VPN membership information. Dynamic data services supported include local Epipe VLL services, Epipe VLL services with dynamic Multi-Segment PseudoWires (MS-PWs) (FEC129) or spoke SDP, VPLS services with BGP-AD PWs or mesh/spoke SDPs, IES, and VPRN services.

To display the complete white list of dynamic data services CLI configuration commands, use the **tools dump service dynamic-services command-list** CLI command.

A Python script interface adds a flexible abstraction layer so that only the business user specific service parameters (service type, IP address, QoS and filter parameters, etc.) are required from RADIUS and are then used in a CLI template to set up the target service.

The setup sequence is shown in [Figure 146: Principle Model of Dynamic Data Services](#) with the example of a VPLS service.

Figure 146: Principle Model of Dynamic Data Services



1. Business CPE initiates a PPP or DHCP "control channel" session. This session is important to let the BNG and RADIUS understand the existence of a new access circuit and the location of the service endpoint.
2. BNG sends an Access-Request with all user credentials to RADIUS.
3. RADIUS replies using an Access-Accept with attributes for the PPP/DHCP control channel and attributes for the dynamic data service (Service-type, SAP-VLAN, QoS, Filter, etc.).
4. BNG creates a dynamic data service instance (if it is first access-circuit for this service) and a SAP, thus completing the service configuration.
5. BNG sends an Accounting Start for PPP/DHCP control channel session and also for the dynamic data service session (and subsequently interim accountings and accounting stop for both).
6. BNG advertises VPLS instance-ID (123) via BGP-AD to other PEs/BNGs.
7. PEs/BNGs with same service instance will auto-establish SDPs to the BNG.

The result is a fully functional service which is the same as a traditionally configured service.

The lifetime of the dynamic data services are bound to the existence of the control channel session. If, for whatever reason, the control channel session is torn down all associated dynamic data services will also be terminated.

Dynamic data service SAPs have to be located on dot1q or qinq encapsulated Ethernet ports and can be part of a LAG.

Both XML accounting and RADIUS accounting can be enabled on a dynamic data service SAP. The RADIUS accounting data can be sent to up to two different RADIUS servers.

There is a strict separation of services created by dynamic service provisioning and services created via the CLI or through other standard mechanisms (5620 SAM, SNMP). It is therefore not allowed to:

- create a dynamic services object in a local provisioned CLI/SNMP context (e.g. create a dynamic SAP in a local provisioned VPLS).
- create a local provisioned object in a dynamic service context (e.g. create a SAP via CLI/SNMP in a dynamic VPLS service).
- change parameters in a local provisioned CLI/SNMP context using the dynamic services model (change system name with dynamic services provisioning).
- change parameters with the CLI/SNMP in a dynamically created context.
- delete a local provisioned object using dynamic provisioning model.
- delete a dynamic provisioning object using the CLI/SNMP.
- create a reference to a dynamic services object in a local provisioned CLI/SNMP context (reference to dynamic interface in **router ospf**)


A special command exists to overcome some of the above rules. This command is designed to ease Python script creation and testing and not for normal operations. This is discussed in [Configuration](#).

Configuration

It is assumed that the reader is familiar with the regular Enhanced Subscriber Management (ESM) functionality as well as with general service related configurations. Furthermore certain knowledge about Python programming is also assumed.

The test topology is shown in [Figure 147: Test Topology](#).

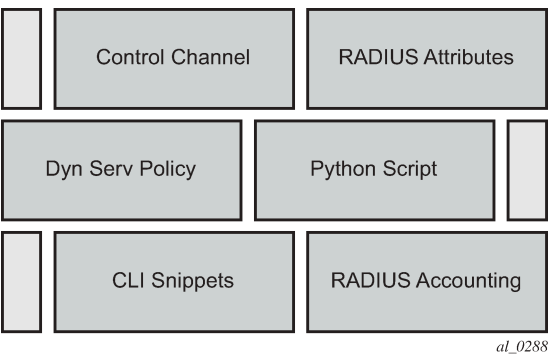
Figure 147: Test Topology



The pure service setup can be tested with a single node acting as BNG. However an Epipe service between two nodes will not normally become status “up” with only one endpoint in an up state. As such, for packets should be sent through the established dynamic data service, a remote PE could also be configured. The remote PE could have its data service configured in a regular fashion, meaning via CLI/ SNMP or 5620 SAM.

The required functionality on the BNG is divided into multiple building blocks. The following sections discuss each building block in detail.

Figure 148: Building Blocks of Dynamic Data Services



Based on a PPP or DHCP control session, RADIUS will return the required parameters for the dynamic data service via dedicated Vendor Specific Attributes (VSAs). The existence of those attributes in the RADIUS Accept message will trigger the relaying of the parameters relating to those attributes towards the Python script defined in the dynamic service policy, which will process them to generate the regular CLI output for the various service types (IES, VPRN, Epipe, VPLS).

For efficiency and flexibility the Python script needs to be structured into different parts per service which then reference each other internally. Those parts are called snippets.

Finally, as the services are initiated from RADIUS, RADIUS accounting messages per dynamic data service will be sent to the RADIUS server as a necessary feedback mechanism to inform the RADIUS server about a successful or failed service setup.

Building Block: Control Channel

The configuration to authenticate and instantiate a dynamic data service control channel is identical to a residential Enhanced Subscriber Management (ESM) configuration. Examples for this can be found in other chapters of the advanced configuration guide and will not be covered here in detail.

Building Block: Dynamic Services Policy

The dynamic services parameters are configured under the **configure service dynamic-services** CLI context. The following output shows two policy examples.

```
configure service dynamic-services
    dynamic-services-policy "dynamic-services-1" create
        accounting-1
            server-policy "radius-server-policy-1"
            update-interval min 5
        exit
        accounting-2
            server-policy "radius-server-policy-2"
            stats-type time
            update-interval min 5
            update-interval-jitter absolute 10
        exit
        cli-user "dynuser"
        description "Dynamic Service Policy #1"
        sap-limit 4000
        script-policy "script-policy-1"
    exit
    dynamic-services-policy "dynamic-services-2" create
        accounting-1
            server-policy "radius-server-policy-2"
        stats-type volume-time
            update-interval min 30
            update-interval-jitter absolute 20
        exit
        accounting-2
            server-policy "radius-server-policy-2"
            stats-type time
            update-interval min 5
            update-interval-jitter absolute 10
        exit
        cli-user "dynuser"
        description "Dynamic Service Policy #2"
        sap-limit 100
        script-policy "script-policy-2"
    exit
    service-range 1000 10000
    timers
        setup-timeout access-accept 3
    exit
```

Details of each command and the possible parameters can be found in the *7450 ESS, 7750 SR, and VSR Triple Play Service Delivery Architecture Guide* in the RADIUS Triggered Dynamic Data Services section.

On the top command level under the dynamic-services sub-tree there are three options:

- dynamic-services-policy
- service-range
- timers

The setup-timeout value under **timers** is used to limit the maximum delay allowed for a dynamic data service setup. In addition, it also protects the node during times where there is a high load on the CPU. If a requested dynamic data service cannot be established in the specified time the request will be dropped.

Dynamic data services are not preferred over regular ESM subscribers. As such, given a BNG with a mix of residential ESM subscribers and business customers with dynamic data services, all compete for the same CPU resources to establish the connections.

However, dynamic data services are expected to have a very long lifetime compared to potentially very dynamic lifetimes for residential subscribers. In a regular operating mode the amount of additional setup requests for dynamic data services should be relatively small. Only in the event of a node reboot will all users again compete to gain access, where longer setup-times are inevitable.

The service-range value reserves a certain amount of service IDs for the use of dynamic data services. The configured range is no longer available for regular provisioned services configured via the CLI/SNMP.

The dynamic-services-policy contains a CLI-user identifier, SAP-limits, accounting parameters and reference to a Python script policy which is used when creating a dynamic data service. Multiple dynamic services policies can be created to enable different profiles to be used for different users/customers or services (as an example, two different departments within the service provider, one responsible for Layer 2 services, one for Layer 3 services). The policy used for a dynamic data service is determined from the Alc-Dyn-Serv-Policy [26-6527-167] RADIUS attribute. If the attribute is not present and a policy named *default* exists, then the *default* policy is used, otherwise the dynamic data service creation fails.

Up to two accounting server policies can be defined. This allows the use of separate RADIUS accounting servers independent from the accounting servers used for residential services. The parameters defined in the accounting sections are the default values which are used if no specific values are sent via RADIUS VSAs.

As the service is established via RADIUS, a feedback mechanism towards RADIUS is most likely required which would be at least RADIUS start and stop messages per service/session. In addition performance counters (with a fixed set of parameters) can also be included in the RADIUS messages. It is also possible to use the standard service-accounting under the service instance and remove any counters from the RADIUS accounting messages.

The specification of a CLI user allows linking of the dynamic data service to a specific user-profile. In addition, this facilitates limiting of the scope of allowed service configurations even further, based on the specified context under the user profile.

The CLI-user needs to be configured locally on the node and needs to have a local user profile (remote authorization via TACACS/RADIUS is not possible).

The radius-script-policy is configured under the **configure aaa** CLI context.

```
configure aaa
radius-script-policy "script-policy-2"
  action-on-fail passthrough
  primary
    script-url "cf3:/scripts/dyn_services.py"
  no shutdown
```

```

        exit
        secondary
            script-url ftp://user*:pwd@10.255.137.80/scripts/dyn_services.py"
            no shutdown
        exit
    exit
exit

```

The parameters are no different to what have been defined generally for the use of Python scripting on the BNG.

When the very first session request arrives, the Python script is loaded into memory and executed. For all subsequent session requests the script is executed without the need for a reload. It is possible for both primary and secondary locations to be FTP sites (the small transfer delay for the first session is acceptable), however, it is recommended to have a compact-flash (cf1 or cf2) as the primary location and a remote location as backup.

Building Block: RADIUS Attributes

A series of vendor specific attributes (VSAs) have been defined to setup, teardown or modify dynamic data services from RADIUS.

The VSAs and their meaning are as follows:

- Alc-Dyn-Serv-SAP-Id [26-6527-164], type "string"

This attribute identifies the dynamic service SAP. The format can be any valid Ethernet SAP format (dot1q or qinq encapsulation), including LAGs. A wildcard ("#") can be specified for the port field and optionally for one of the tag fields of a qinq interface. To define the dynamic data service SAP-ID, the wildcard fields are replaced with the corresponding field from the Control Channel SAP-ID.

Examples: "1/2/7:10.100" or "#:#.100"

- Alc-Dyn-Serv-Script-Action [26-6527-166], type "integer"

A mandatory VSA in a COA to the control channel accounting session ID or the accounting session ID of the dynamic data service (only applicable for modify or teardown). Tells the system what script action is required: setup, modify or teardown of a dynamic data service.

Values: 1=setup, 2=modify, 3=teardown

- Alc-Dyn-Serv-Policy [26-6527-167], type "string"

Specifies the dynamic service policy to use for provisioning the dynamic service. The policy must be configured in the "configure service dynamic-services dynamic-services-policy < dynsrv-policy-name>" CLI context.

- Alc-Dyn-Serv-Script-Params [26-6527-165], type "string"

This VSA contains parameters that can be used by the Python script to setup or modify a dynamic data service. The parameters can be split into multiple instances of the same attribute, linked together by the same tag, that is, the parameters can cross an attribute boundary. The concatenation of all "Alc-Dyn-Serv-Script-Params" attributes with the same tag in a single message must be formatted as "function-key = {dictionary}" where function-key specifies which Python functions will be called and {dictionary} contains the actual parameters in a Python dictionary structure format.

Example: "business_1 = { 'as_id' : '100', 'comm_id' : '200', 'if_name' : 'itf1', 'ipv4_address' : '172.16.1.1', 'egr_ip_filter' : '100', 'routes' : [{ 'to' : '172.16.100.0/24', 'next-hop' : '172.16.1.2' }, { 'to' : '172.16.200.0/24', 'next-hop' : '172.16.1.2' }] }

The above example shows each parameter with a keyword and the associated value. Alternatively only the parameter values can be sent with a pre-defined (and always constant) sequence.

Example: "business_1 = {"t": '100', '200', 'itf1', '172.16.1.1', '100', '172.16.100.0/24', '172.16.1.2', '172.16.200.0/24', '172.16.1.2'}.

- Alc-Dyn-Serv-Acct-Interim-Ivl-1 [26-6527-168], type "integer"

This VSA defines the number of seconds between each accounting interim update message for the primary accounting server. It overrides the local configured "update-interval" value in the dynamic services policy "accounting-1" CLI context. A value of 0 (zero) corresponds to no accounting interim update messages. A value [1..299] seconds is rounded to 300s (min. CLI value) and a value above 15552000 seconds (180 days, maximum CLI value) is rounded to the maximum CLI value.

Range = 0 | [300 - 15552000].

- Alc-Dyn-Serv-Acct-Interim-Ivl-2 [26-6527-169], type "integer"

Same function and values as Alc-Dyn-Serv-Acct-Interim-Ivl-1 [26-6527-168], for the second accounting server. It overrides the locally configured "update-interval" value in the dynamic services policy "accounting-2" CLI context.

- Alc-Dyn-Serv-Acct-Stats-Type-1 [26-6527-170], type "integer"

Enable or disable dynamic data service accounting to the primary accounting server and specify the type of statistics that should be reported: volume and time or time only. It overrides the locally configured value in the dynamic services policy "accounting-1" CLI context.

Values: 1=off, 2=volume-time, 3=time

- Alc-Dyn-Serv-Acct-Stats-Type-2 [26-6527-171], type "integer"

Enable or disable dynamic data service accounting to the secondary accounting server and specify the type of statistics that should be reported: volume and time or time only. It overrides the locally configured "stats-type" value in the dynamic services policy "accounting-2" CLI context.

Values: 1=off, 2=volume-time, 3=time

All VSAs are tagged to enable manipulation of up to 32 (tag values 0..31) dynamic data services in a single RADIUS message. VSAs with an identical tag belong to the same dynamic data service.

The use of the VSAs in RADIUS Access-Accept, CoA and Disconnect Messages is detailed in [Table 31: Dynamic Service Attribute List for Setup, Modify and Teardown](#). An Access-Accept message can only contain dynamic data service setup requests. A CoA can be used to setup, modify or terminate a dynamic data service. A Disconnect Message can only be used to terminate a dynamic data service.

Table 31: Dynamic Service Attribute List for Setup, Modify and Teardown

Attribute Name	Access Accept	CoA			Disc. Message	Comment
	Setup	Setup	Modify	Teardown	Teardown	
Acct-Session-Id	N/A	M	M	M	M	Acct-Session-Id of the Control Channel or in case of a CoA: any other valid CoA key for ESM hosts/sessions.

Attribute Name	Access Accept	CoA			Disc. Message	Comment
	Setup	Setup	Modify	Teardown	Teardown	
Alc-Dyn-Serv-SAP-Id	M	M(*)	M(*)	M(*)	N/A	Identifies the dynamic data service
Alc-Dyn-Serv-Script-Params	O	M(*)	M(*)	N/A	N/A	For a Modify, the script parameters represent the new parameters required for the change.
Alc-Dyn-Serv-Script-Action	O	M(*)	M(*)	M(*)	N/A	Must be "setup" if specified in an access-accept.
Alc-Dyn-Serv-Policy	O	O	O	O	N/A	The <i>default</i> policy used when not specified for create. In CoA, must be same as used for Setup if Specified for Modify or Teardown.
Alc-Dyn-Serv-Acct-Interim-lvl-1	O	O	X(**)	X(**)	N/A	
Alc-Dyn-Serv-Acct-Interim-lvl-2	O	O	X(**)	X(**)	N/A	
Alc-Dyn-Serv-Acct-Stats-Type-1	O	O	X(**)	X(**)	N/A	
Alc-Dyn-Serv-Acct-Stats-Type-2	O	O	X(**)	X(**)	N/A	
M = Mandatory, O= Optional, X = May not, N/A = Not Applicable (ignored)						
(*) = CoA Nak'd, if not specified (Error Cause: 402 - Missing Attribute)						
(**) = CoA Nak'd if specified (Error Cause:405 - Unsupported Service)						

To summarize, [Table 32: Dynamic Service Actions on Control- and Data-Channel](#) shows resulting dynamic service script actions as function of the RADIUS message (Access-Accept, CoA or DM) and the target (Control Channel or Dynamic Service SAP).

Table 32: Dynamic Service Actions on Control- and Data-Channel

Target	RADIUS Message	Dynamic Service Script Action	Comments
Control Channel	Access-Accept	Setup	Up to 32 dynamic data services in single message.

Target	RADIUS Message	Dynamic Service Script Action	Comments
			Alc-Dyn-Serv-Script-Action VSA optional.
		Modify/Teardown	Not supported.
	CoA (acct-session-id or any other valid CoA key for ESM hosts/sessions)	Create/Modify/ Teardown	Cannot be mixed with session/post parameter changes in the same RADIUS message (results in CoA NAK). Up to 32 dynamic data services in single message. Alc-Dyn-Serv-Script-Action VSA mandatory.
	Disconnect	N/A	Teardown the Control Channel session and all associated dynamic data services.
Dynamic Service	CoA (acct-session-id of the dynamic data service sap)	Modify/Teardown	Only single dynamic data service per message (Acct- Session-Id). Alc-Dyn-Serv-Script-Action VSA mandatory.
		Setup	Not supported.
	Disconnect (acct-session-id of the dynamic data service sap)	N/A	Teardown the corresponding dynamic data service.

Building Block: Python Script

Dynamic data services scripts are built using a Python script engine. The following dedicated functions are available in the alc.dyn module:

- `dyn.reference(function-key, reference-id string, dictionary)`

This function creates a dynamic reference to another function in the script. This function eases the creation of N:1 relationships in the script. For more information about use cases, see [Building Block: CLI Snippets](#). The function-key specifies the key in the action dictionary to find the corresponding setup/modify/teardown function calls.

The reference-id (typically derived from a parameter specified from RADIUS, for example: service-name) specifies a unique instance string that identifies this reference.

The dictionary specifies a dictionary with parameters that can be used in the parent function to generate CLI script output.

- `dyn.action(d)`

When called, the `dyn.action` will take the "function-key" string specified in the Alc-Dyn-Serv-Script-Params attribute, and perform a lookup in the specified dictionary `d` to find the corresponding Python function to execute. The format of the dictionary is `d = {key-1 : (Setup-1, Modify-1, Revert-1, Teardown-1), ..., key-n : (Setup-n, Modify-n, Revert-n, Teardown-n) }`. If the function-key matches, for example, `key-1` and the corresponding Alc-Dyn-Script-Action is "setup", then the function specified as

"Setup-1" will be executed. Setup and teardown functions are mandatory. Modify and revert functions are optional. If a modify function is defined, a corresponding revert function must also be defined. If no modify/revert function is required, the keyword **None** should be used instead.

- `dyn.add_cli(string)`

This function is used to generate CLI output in the Python script. The use of `dyn.add_cli("""`) allows the specification of strings spanning multiple lines, which drastically improves the readability of the script.

A subset of all available CLI commands is currently enabled for dynamic data services. The command "tools dump service dynamic-services command-list" provides a complete overview of all available CLI nodes for dynamic data services. In the allowed nodes section, all CLI nodes are listed that can be navigated to and where attributes can be modified. The pass through nodes section shows CLI nodes that can be navigated to but no attribute changes are allowed. For example, it is not allowed to change the autonomous system of a router (configure router autonomous-system <autonomous-system>) because "configure router" is a "pass through node". However, you can navigate to configure router, because you can add a static route: "/configure router static-route 0.0.0.0/0 next-hop 192.168.1.1" is part of the "allowed nodes".

- `dyn.select_free_id("service-id")`

This function is used to select a free service ID within the service ID range defined under dynamic-services context. An automatic assignment of the service id is one option, but it is also possible to provide the service id as one of the parameters in the "Alc-Dyn-Serv-Script-Params" list from RADIUS.

The service-ID is a node-internal attribute. As such it is valid to let the node select the ID itself. However, in a network with multiple BNGs and a single customer service spanning two or more BNGs, a network administrator may actually prefer to use the same service-id for this customer service on all nodes for better visibility, which cannot be guaranteed if the automatic option is chosen. 5620 SAM is also using the service-ID as one attribute in addition to others to discover service-entities across the whole network. If SAM is in use for general management and service assurance, it is advised to manually specify the service-ID and not to use the automatic selection.

In any case, the administrator needs to make a choice between the automatic ID assignment and the specific assignment for all dynamic data services, as a mix between both is not recommended.

When the automatic assignment is chosen, there is no "binding/memory" of a service ID to a provisioned service, which means a service that may have service ID xyz initially may get another service ID the next time it comes up. In other words, as soon as a service is disconnected, the service ID is freed up for the next activated service.

- `dyn.get_sap()`

This function returns the value of the evaluation of the "Alc-Dyn-Serv-SAP-Id" attribute as a string. Wildcards ("#") in the Alc-Dyn-Serv-SAP-Id are replaced with the corresponding port/vlan information of the control channel SAP-ID. So if, for example, the "Alc-Dyn-Serv-SAP-Id" contains "#:#.1" and the control channel SAP ID is "1/1/5:100.100", the resulting SAP for the data service would be "1/1/5:100.1".

- `dyn.get_circuit_id()`

This function returns a string which is equal to the Control Channel Circuit-ID (from the DHCP relay agent option 82 or PPP tags). This function may be useful, for example, to use the circuit id in the SAP description.

- `dyn.get_remote_id()`

This function returns a string which is equal to the Control Channel Remote-ID (from the DHCP relay agent option 82 or PPP tags). This function may be useful, for example, to use the remote id in the SAP description.

In addition to the RADIUS dictionary, the node will also store service-related parameters in a service-specific dictionary. The information in the RADIUS messages or in the stored dictionary are used for the various functions as outlined in [Table 33: Function and Dictionary Relationship](#) :

Table 33: Function and Dictionary Relationship

Function Name	Input	Returns
setup_dynsvc(rd*)	rd : radius dictionary in the parameter list in Alc-Dyn-Serv-Script-Params. VSA Passed to setup function.	A dictionary that will be stored for the lifetime of the dynamic service (sd).
modify_dynsvc(rd,sd**)	rd : radius dictionary in the parameter list in Alc-Dyn-Serv-Script-Params. VSA passed to modify function. sd : previously stored dictionary of the setup/previous modify functions.	Updated stored dictionary (sd)
revert_dynsvc(rd, sd)	rd : radius dictionary in the parameter list in Alc-Dyn-Serv-Script-Params. VSA passed to revert function. sd : previously stored dictionary of the setup/previous modify function.	The function does not return (store) any information. The previously stored dictionary (sd) is kept.
teardown_dynsvc(sd)	sd : previously stored dictionary by the setup function or a previous modify function are passed to the teardown function.	The function does not return (store) any information. The stored dictionary (sd) is deleted.
(*) rd = radius dictionary		
(**) sd = stored dictionary. sd is required for modifies, reverts and teardowns.		

Building Block: CLI Snippets

The necessary functional parts of a service configuration cannot typically be put into one large script (one single actionable function). This is best described with a small and simple example:


Imagine a single script where the setup action creates both the service instance and the SAP, and the teardown action removes the service instance and the SAP. For a service with just one SAP per service instance this may work fine, however, in a multi SAP service like a VPLS this will cause problems, especially during the service teardown action. This is because if multiple SAPs have been instantiated in a single service, the disconnect of just one SAP would trigger the teardown action which would try to remove the SAP (still ok) but then would try to remove the service instance. This action would fail as other SAPs still exist in the service. As such the script execution would fail.

It is therefore necessary to structure the whole required configuration into individual actionable pieces which are referenced by each other with specific reference-IDs. Those actionable pieces are called "snippets".

Referenced snippets may or may not be executed depending on whether the functional instance exists already or not. As shown on the left of the picture below, the action to create a SAP references the creation of a service and then to the creation of a customer. For the very first business site to come up all three

snippets will be executed. For any further business site to come up in the same service the script to create the SAP will be executed, the referenced service script and subsequently the customer script will not be executed again as those instances already exist. The same logic applies during the teardown action. Only when the last SAP in a service is removed is the service-instance itself removed, and potentially also the customer (unless it too is associated with other services).

Figure 149: Hierarchy of Snippets



The implementation supports a three level hierarchy of snippets for high flexibility as shown in the picture. A reference to the fourth level as shown on the right side would result in an error.

Furthermore, snippets can be scaled "horizontally", so from one level multiple references to other snippets are possible. An example for that would be the creation of a SAP triggers the creation of a service as well as the creation of an Ethernet CFM association for that SAP.

Identifiers are needed for the referencing. The same identifier can be used on the "horizontal" level, but not on the vertical level between the same pair of snippets, also shown above.

Snippets are heavily used in the service examples in [Bringing it all together](#) where the logic and the referencing are described with real data.

Building Block: RADIUS Accounting

As dynamic data services are instantiated through RADIUS, it is also typically required to provide feedback to the RADIUS server for service establishment and teardown. This is achieved via RADIUS accounting records for the dynamic data channels in addition to the accounting messages for the PPP or DHCP control channel.

Up to two dedicated accounting destinations can be defined within the dynamic services policy. Thus, the accounting for the dynamic data services can be handled by an independent set of accounting servers (from the accounting for general ESM subscribers). But the same servers can also be used.

Each dynamic data service has its own accounting start/stop/interim messages based on a unique accounting session ID. In addition, the accounting packets contain a multi-session ID which is identical to the accounting session ID of the control channel and is therefore displayed in show commands as Acct-Session-ID-Ctrl as shown below.

```
A:BNG-1# show service dynamic-services saps summary
=====
Dynamic Services SAP's summary
=====
SAP                               Acct-Session-ID      Acct-Session-ID-Ctrl
-----
3/2/1:4.3                        D6E559000000B951668AEB D6E559000000B851668AEB
3/2/2:1.1                        D6E559000000C75166CFF4 D6E559000000C45166CFF4
3/2/2:1.2                        D6E559000000C85166CFF4 D6E559000000C45166CFF4
3/2/2:1.3                        D6E559000000C95166CFF4 D6E559000000C45166CFF4
3/2/2:1.4                        D6E559000000CA5166CFF4 D6E559000000C45166CFF4
-----
No. of SAP's: 5
=====
```

The Accounting Session ID (in the centre above) is the one for the dynamic data service itself, the one on the right is from the control-channel. The above example clearly shows that the last 4 dynamic services all belong to the same control channel, as they all have the same Acct-Session-ID-Ctrl.

If the accounting stats-type is set to "volume-time", the interim and stop accounting messages will also contain counters for the data traffic through the service. With the accounting stats-type "time", no counters are included, only session time is reported.

As a dynamic data service is functionally no different from a regular data service, traffic volumes can also be gathered by assigning accounting policies within the service for file-based XML accounting.

Bringing it all together

This section gives examples of all of the above parameters and will also cover show, log and debug information.

In the given example, a single user in the database has four different associated data services. Not only are the data service types all different, but also other aspects of the parameter set, this has an effect on how the data is entered in the RADIUS VSAs and how the Python script is constructed. More detail is given below. The different models for specifying parameters are presented to show the flexibility. An operator typically chooses a single model and uses that for all its services.

As all of the information for these four services will potentially be sent in one RADIUS message, the VSAs need to be tagged so that the BNG can link the appropriate VSAs to each other and differentiate the services. For better visibility, the different sections in the RADIUS users file are displayed with bold black and dark grey text.

The freeradius users file format is used for this example.

```
1.  "subscriber12@domain2.com" Cleartext-Password := "ALU"
2.  Alc-Subsc-ID-Str := "pppoe-user12",
3.  Framed-IP-Address = 10.2.1.200,
4.  Alc-Dyn-Serv-SAP-Id:1 = "#:#.1",
```



```

5.      Alc-Dyn-Serv-Script-Params:1 = "business_epipe={'t':('EPipe-
6.      CustomerName', 'CustomerName-Circuit-1', '3', '3', '64496',
7.      '192.0.2.5', '192.0.2.1', '3333'))}",
8.      Alc-Dyn-Serv-Policy:1 = "dynamic-services-2",
9.      Alc-Dyn-Serv-SAP-Id:2 += ":#.2",
10.     Alc-Dyn-Serv-Script-Params:2 += "business_vprn={'t':('9999',
11.     'VPRN-CustomerName', '64497', '100000', 'CustomerName-Circuit-
12.     1', '172.16.10.1/30', '3', '1', '3', '2', '172.16.100.0/24',
13.     '172.16.10.2', '100'))}",
14.     Alc-Dyn-Serv-Acct-Interim-Ivl-1:2 += "600",
15.     Alc-Dyn-Serv-Acct-Interim-Ivl-2:2 += "0",
16.     Alc-Dyn-Serv-Policy:2 += "dynamic-services-2",
17.     Alc-Dyn-Serv-SAP-Id:3 += ":#.3",
18.     Alc-Dyn-Serv-Script-Params:3 += "business_vpls={'inst':
19.     'VPLS-CustomerName', 'if_name': 'CustomerName-Circuit-1', 'ing_qos': '3',
20.     'egr_qos': '3', 'imp_comm_val': '10000', 'exp_comm_val': '10000',
21.     'rt': '64498', 'rd': '64498'})",
22.     Alc-Dyn-Serv-Policy:3 += "dynamic-services-2",
23.     Alc-Dyn-Serv-Acct-Interim-Ivl-1:3 += "0",
24.     Alc-Dyn-Serv-Acct-Interim-Ivl-2:3 += "0",
25.     Alc-Dyn-Serv-Acct-Stats-Type-1:3 += off,
26.     Alc-Dyn-Serv-Acct-Stats-Type-2:3 += off,
27.     Alc-Dyn-Serv-SAP-Id:4 += ":#.4",
28.     Alc-Dyn-Serv-Script-Params:4 += "business_ies={'t':
29.     ('IES-CustomerName', 'CustomerName-Circuit-1', '172.16.11.1/30',
30.     '2001:db8:5100:1000::1/64', '5', '1', '1', '6', '2', '2', '5', '25',
31.     'cfm-Mep-to-CPE', '100', ",
32.     Alc-Dyn-Serv-Script-Params:4 += "[{'to': '172.16.110.0/24',
33.     'n-h': '172.16.11.2'}, {'to': '2001:db8:bbbb::/56',
34.     'n-h': '2001:db8:5100:1000::2'})]",
35.     Alc-Dyn-Serv-Policy:4 += "dynamic-services-2",
36.     Alc-Dyn-Serv-Acct-Interim-Ivl-1:4 += "600",
37.     Alc-Dyn-Serv-Acct-Interim-Ivl-2:4 += "0",
38.     Alc-Dyn-Serv-Acct-Stats-Type-1:4 += "3",
39.     Alc-Dyn-Serv-Acct-Stats-Type-2:4 += "2",

```

The first section (lines 1 — 3) shows a minimal parameter set for the (PPP) control channel. As the focus of this example is on the dynamic data services, all default parameters will be used for the control-session which are defined under the msap.

The second section (lines 4 — 8, attributes with tag ":1") shows a possible parameter set for an Epipe service. Only the absolutely minimum set of VSAs is used (see [Table 31: Dynamic Service Attribute List for Setup, Modify and Teardown](#)). Furthermore, all service parameters are listed without keywords in a pre-defined order. No service ID number is specified in "Alc-Dyn-Serv-Script-Params", hence the Python script should dynamically select the next free ID.

The third section (lines 9 — 16, attributes with tag ":2") shows a possible parameter set for a VPRN service. A few more VSAs are defined, thus some of the default parameters in the dynamic service policy are overwritten for this service. The first entry in the "Alc-Dyn-Serv-Script-Params" attribute specifies the Service-ID number for this service, so the Python script should not select a service ID automatically. Furthermore, static-routing information towards the CPE is added as normal attributes at the end of the list.

The fourth section (line 17 — 26, attributes with tag ":3") shows a possible parameter set for a VPLS service. Notice the difference with the first two services in the "Alc-Dyn-Serv-Script-Params" part: now all parameters are given their specific keyword. As such, the sequence of those parameters is not important. The effect on the Python script is shown further down.

The fifth section (lines 27 — 39, attributes with tag ":4") finally shows a possible parameter set for an IES service. All of the required parameters for this service do not fit into a single "Alc-Dyn-Serv-Script-Params" attribute anymore (limited to 247 bytes). As is shown, multiple VSAs can be "concatenated" by simply splitting the attributes. It is important that the order in which the different "Alc-Dyn-Serv-Script-Params"

attributes with the same tag is received can be guaranteed. Furthermore the second appearance of this VSA shows a different way of provisioning static-routing information towards the CPE.

To better understand the details it is necessary to take a closer look into the active Python script. The first important part is the section with the dynamic actions.

```
---snip---

d = {
    "vprn": (setup_vprn, None, None, teardown_vprn),
    "ies": (setup_ies, None, None, teardown_ies),
    "vpls": (setup_vpls, None, None, teardown_vpls),
    "epipe": (setup_epipe, None, None, teardown_epipe),
    "ethcfm": (setup_ethcfm_domain, None, None, teardown_ethcfm_domain),
    "business_vprn" : (setup_business_vprn, None, None, teardown_business_vprn),
    "business_ies" : (setup_business_ies, None, None, teardown_business_ies),
    "business_vpls" : (setup_business_vpls, None, None, teardown_business_vpls),
    "business_epipe" : (setup_business_epipe, modify_business_epipe,
                       revert_business_epipe, teardown_business_epipe)}

dyn.action(d)
```

The function-key string specified at the start of the "Alc-Dyn-Serv-Script-Params" (for example Alc-Dyn-Serv-Script-Params:1 = "business_epipe={...}") has a 1:1 mapping with the keys of the dictionary "d" in the highlighted section of the above sample (for example d = { ..., "business_epipe" : (...)}). For services, different values for setup, modify, revert and teardown are given which point to other sections in the Python script (see below). Setup and teardown functions are mandatory, whereas modify and revert functions are optional.

In the unbolded text of the previous example, there are other actions defined that are not contained in the RADIUS attributes (for example d = { "vprn": (...), ... }). Those actions are referenced by the four main functions.

In the next part, there is more detail presented in each service example and maps it to the corresponding Python function.

It is advisable to read through all examples, as only the deltas between each service are explicitly explained.

Example 1 – Epipe service

```
# copy of the RADIUS attributes from above
---snip---
    Alc-Dyn-Serv-SAP-Id:1 = "#:#.1",
    Alc-Dyn-Serv-Script-Params:1 = "business_epipe={'t':
        ('Epipe-CustomerName', 'CustomerName-Circuit-1', '3', '3', '64496'
        , '192.0.2.5', '192.0.2.1', '3333')}",
    Alc-Dyn-Serv-Policy:1 = "dynamic-services-2",
---snip---

# Python-part
d = {
---snip---
    "business_epipe" : (setup_business_epipe, modify_business_epipe,
                       revert_business_epipe, teardown_business_epipe)

dyn.action(d)
---snip---
def setup_business_epipe(d):
    keys = ('inst', 'if_name', 'ing_qos', 'egr_qos', 'as', 'remote_ip',
            'local_ip', 'glb_svc_id')
    d = dict(zip(keys, d['t']))
    ref_d = dyn.reference("epipe", d['inst'], d)
```

```

        d['svc_id'] = ref_d['svc_id']
        d['sap_id'] = dyn.get_sap()
        dyn.add_cli("""
configure
service
    epipe %(svc_id)s
        sap %(sap_id)s create
            description "%(if_name)s"
            ingress
                qos %(ing_qos)s
            exit
            egress
                qos %(egr_qos)s
            exit
        exit
        spoke-sdp-fec %(svc_id)s fec 129 aii-type 2 create
        pw-template-bind 2
        saii-type2 %(as)s:%(local_ip)s:%(glb_svc_id)s
        taii-type2 %(as)s:%(remote_ip)s:%(glb_svc_id)s
        no shutdown
    exit
exit
exit
""")
    return d

def setup_epipe(d):
    d['svc_id'] = dyn.select_free_id("service-id")
    dyn.add_cli("""
configure
service
    epipe %(svc_id)s customer 1 create
        service-name "%(inst)s"
        description "%(inst)s"
        no shutdown
    exit
exit
""")
    return {'svc_id':d['svc_id']}

def teardown_epipe(d):
    dyn.add_cli("""
configure
service
    epipe %(svc_id)s
        shutdown
    exit
    no epipe %(svc_id)s
exit
""")

def teardown_business_epipe(d):
    dyn.add_cli("""
configure
service
    epipe %(svc_id)s
        sap %(sap_id)s
        shutdown
    exit
    spoke-sdp-fec %(svc_id)s
        shutdown

```

```

        exit
        no sap %(sap_id)s
        no spoke-sdp-fec %(svc_id)s
    exit
exit
exit
""" % d)
---snip---
```

Based on the dictionary specified in the `dyn.action(d)` call, the function definition `setup_business_epipe` in the Python script corresponds with the function that will be called if the function-key `"business-epipe"` is specified in the `"Alc-Dyn-Serv-Script-Params"` attribute as dictionary name and if a setup action is required. The dictionary containing the parameters in the RADIUS VSA `"Alc-Dyn-Serv-Script-Params"` has a single key-value pair, with the parameters stored in a tuple. The individual parameters cannot be identified with a keyword hence the order in which they are specified in the RADIUS VSA should match the order in which they are extracted in the Python script. The first two lines in this part of the script extract the parameters out of the array `"t"` and link them to unique keywords, which are used for the rest of the script.

The parameter `"inst"` is important in this logic, as it defines whether access circuits belong to the same service-instance or different instances (the RADIUS VSAs for two SAPs belonging to the same service therefore need to have the same `"inst"` value). If you look at the CLI of the `setup_business_epipe` function, you can see that it creates the SAP and all related attributes, but not the service itself. It is the `"ref_d = dyn.reference("epipe", d["inst"], d)"` that references a part in the script to create the actual service-instance. The referenced function is found by using the first parameter in the `dyn.reference` call (`"epipe"`) as a function-key lookup in the dictionary specified in the `dyn.action(d)` and finding the corresponding setup function: `d = { ..., "epipe": (setup_epipe, ...), ...}`. The second parameter (`"d["inst"]"`) is used as unique identification of the service instance. The last parameter (`"d"`) is a dictionary with parameters that can be used by the references function. When the first customer endpoint with a new `"inst"` name comes up, the service itself gets created.

By looking at `"def setup_epipe(d):"` the first line `"d['svc_id'] = dyn.select_free_id("service-id")"` of the script automatically picks a free service-id out of the range defined in the dynamic service policy, as no service ID was provided in the RADIUS parameters. The rest of this function creates the service instance. Service attributes that were provided by RADIUS and are placed in a service specific dictionary are available to this function via the third parameter in the `dyn.reference` call. The newly generated service ID is returned to the calling script by the `"return {'svc_id':d['svc_id']}"` command at the end of the function. The service specific dictionary (as explained in the Python Script Building Block) is updated with the appropriate information.

Back to `"def setup_business_epipe(d):"`, the service ID together with the SAP ID and the parameters from the `Alc-Dyn-Serv-Script-Params` VSA are used to create the appropriate CLI code for the SAP and the SDP within the service.

Similar to the setup, there is also a teardown part for both service and SAP. The teardown function is called either through the termination of the control-channel, through a COA with `Alc-Dyn-Script-Action = teardown` or through a disconnect message. The CLI for the teardown script must be written in the correct sequence as applied by the SR OS CLI logic so that SAP(s) and service(s) are removed in the correct order.

Example 2 – VPRN service

```

RADIUS-part from above
---snip---
Alc-Dyn-Serv-SAP-Id:2 += "#:#.2",
Alc-Dyn-Serv-Script-Params:2 += "business_vprn={'t':
('9999','VPRN-CustomerName','64497','100000',
'CustomerName-Circuit-1','172.16.10.1/30','3','1','3','2',
'172.16.100.0/24','172.16.10.2','100')}"",
Alc-Dyn-Serv-Acct-Interim-Ivl-1:2 += "600",
Alc-Dyn-Serv-Acct-Interim-Ivl-2:2 += "0",
```

```
Alc-Dyn-Serv-Policy:2 += "dynamic-services-2",
---snip---

Python-part
d = {
---snip---
"business_vprn" : (setup_business_vprn, None, None, teardown_business_vprn)
dyn.action(d)
---snip---
def setup_business_vprn(d):
    keys = ('svc_id', 'inst', 'as_id', 'comm_id', 'if_name', 'ipv4_address',
            'ing_qos', 'ing_ip_filter', 'egr_qos', 'egr_ip_filter', 'lan_pfx',
            'nxt_hop', 'metric')
    d = dict(zip(keys, d['t']))
    ref_d = dyn.reference("vprn", d['inst'], d)
    d['sap_id'] = dyn.get_sap()
    dyn.add_cli("""
configure
service
    vprn %(svc_id)s
        interface "%(if_name)s" create
            address %(ipv4_address)s
            urpf-check mode strict
            sap %(sap_id)s create
                ingress
                    qos %(ing_qos)s
                    filter ip %(ing_ip_filter)s
                exit
                egress
                    qos %(egr_qos)s
                    filter ip %(egr_ip_filter)s
                exit
            exit
        exit
    exit
router
    static-route %(lan_pfx)s next-hop %(nxt_hop)s metric %(metric)s
    exit
exit
""" % d)
    return d

def setup_vprn(d):
    dyn.add_cli("""
configure
service
    vprn %(svc_id)s customer 1 create
        service-name "%(inst)s"
        description "%(inst)s"
        autonomous-system %(as_id)s
        route-distinguisher %(as_id)s:%(comm_id)s
        auto-bind mpls
        vrf-target target:%(as_id)s:%(comm_id)s
        no shutdown
    exit
exit
exit
""" % d)
    return {'svc_id':d['svc_id']}

def teardown_vprn(d):
    dyn.add_cli("""
configure
```

```

service
    vprn %(svc_id)s
    shutdown
    exit
    no vprn %(svc_id)s
    exit
exit
""" % d)

def teardown_business_vprn(d):
    dyn.add_cli("""
configure
router
    no static-route %(lan_pfx)s next-hop %(nxt_hop)s
    exit
service
    vprn %(svc_id)s
    interface "%(if_name)s"
    sap %(sap_id)s
    shutdown
    exit
    no sap %(sap_id)s
    shutdown
    exit
    no interface "%(if_name)s"
    exit
exit
exit
""" % d)
---snip---

```

In this example of a VPRN service two additional RADIUS VSAs are used to overwrite the accounting interim update intervals for the two RADIUS Accounting servers that are specified in the dynamic services policy. The Stats-Type configuration (time or volume-time) is obtained from the dynamic services policy as no RADIUS VSA is provided for that.

The beginning of the "setup_business_vprn" definition is identical to the earlier Epipe service example. This time a service identifier is provided as part of the parameter list. The referenced function to create the VPRN service (def setup_vprn) does not need the line to auto-generate the service ID.

At the end of the setup-procedure there is a basic example to add static-route information in case they are needed for PE-CE communication. Later on, in the IES service example, a more flexible alternative is shown.

Example 3 – VPLS service

```

RADIUS-part from above
---snip---
    Alc-Dyn-Serv-SAP-Id:3 += "#:3",
    Alc-Dyn-Serv-Script-Params:3 += "business_vpls={'inst':
        'VPLS-CustomerName','if_name':'CustomerName-Circuit-1','ing_qos':'3',
        'egr_qos':'3','imp_comm_val':'10000','exp_comm_val':'10000',
        'rt':'64498','rd':'64498'}",
    Alc-Dyn-Serv-Policy:3 += "dynamic-services-2",
    Alc-Dyn-Serv-Acct-Interim-Ivl-1:3 += "0",
    Alc-Dyn-Serv-Acct-Interim-Ivl-2:3 += "0",
    Alc-Dyn-Serv-Acct-Stats-Type-1:3 += off,
    Alc-Dyn-Serv-Acct-Stats-Type-2:3 += off,
---snip---

Python-part
d = {

```

```
---snip---
"business_vpls" : (setup_business_vpls, None, None, teardown_business_vpls)
---snip---
def setup_business_vpls(d):
    ref_d = dyn.reference("vpls", d['inst'], d)
    d['svc_id'] = ref_d['svc_id']
    d['sap_id'] = dyn.get_sap()
    dyn.add_cli("""
configure
service
vpls %(svc_id)s
sap %(sap_id)s create
description "%(if_name)s"
ingress
qos %(ing_qos)s
exit
egress
qos %(egr_qos)s
exit
collect-stats
accounting-policy 10
exit
exit
exit
""")
    return d

def setup_vpls(d):
    d['svc_id'] = dyn.select_free_id("service-id")
    dyn.add_cli("""
configure
service
vpls %(svc_id)s customer 1 create
service-name "%(inst)s"
description "%(inst)s"
bgp
route-distinguisher %(rd)s:%(exp_comm_val)s
route-target export target:%(rt)s:%(exp_comm_val)s
import target:%(rt)s:%(imp_comm_val)s
pw-template-binding 1
exit
exit
bgp-ad
vpls-id %(rt)s:%(exp_comm_val)s
no shutdown
exit
no shutdown
exit
exit
""")
    return {'svc_id':d['svc_id']}

def teardown_vpls(d):
    dyn.add_cli("""
configure
service
vpls %(svc_id)s
shutdown
bgp-ad
shutdown
exit
no bgp-ad
""")
```

```

        bgp
        no pw-template-binding 1
        exit
    exit
    no vpls %(svc_id)s
    exit
exit
""" % d)

def teardown_business_vpls(d):
    dyn.add_cli("""
configure
service
    vpls %(svc_id)s
    sap %(sap_id)s
    shutdown
    exit
    no sap %(sap_id)s
    exit
exit
exit
""" % d)
---snip---

```

In the VPLS example the "Alc-Dyn-Serv-Acct-Stats-Type" is set to "off" for both RADIUS accounting destinations, meaning RADIUS accounting is switched off, even if it is enabled in the dynamic data services policy. In the script you can see that this service uses XML-accounting on the SAP instead ("collect-stats" and "accounting-policy 10").

The dictionary containing the parameters in the RADIUS VSA "Alc-Dyn-Serv-Script-Params" has a key-value pair for each parameter. In the Python script the individual parameters can be identified immediately with the dictionary key. The order in which they are specified in the RADIUS VSA does not have to be strictly defined. The drawback of this approach is that the length of the parameter VSA increases. A single parameter VSA is limited to a length of 246 bytes and the total length of all parameter VSAs for a single service is limited to 1000 bytes.

Example 4 – IES service

```

RADIUS-part from above
---snip---
    Alc-Dyn-Serv-SAP-Id:4 += "#:4",
    Alc-Dyn-Serv-Script-Params:4 += "business_ies={'t':
        ('IES-CustomerName', 'CustomerName-Circuit-1', '172.16.11.1/30',
        '2001:db8:5100:1000::1/64', '5', '1', '1', '6', '2', '2', '5', '25',
        'cfm-Mep-to-CPE', '100', ",
    Alc-Dyn-Serv-Script-Params:4 += "[{'to': '172.16.110.0/24',
        'n-h': '172.16.11.2'}, {'to': '2001:db8:bbbb::/56',
        'n-h': '2001:db8:5100:1000::2'})]",
    Alc-Dyn-Serv-Policy:4 += "dynamic-services-2",
    Alc-Dyn-Serv-Acct-Interim-Ivl-1:4 += "600",
    Alc-Dyn-Serv-Acct-Interim-Ivl-2:4 += "0",
    Alc-Dyn-Serv-Acct-Stats-Type-1:4 += "3",
    Alc-Dyn-Serv-Acct-Stats-Type-2:4 += "2",
---snip---

Python-part
d = {
---snip---
"business_ies" : (setup_business_ies, None, None, teardown_business_ies)
---snip---
def setup_business_ies(d):
    keys = ('inst', 'if_name', 'ipv4_address', 'ipv6_address', 'ing_qos',

```

3HE 20810 AAAC TQZZA


```
router
""" % d)
    for route in d['routes']:
        dyn.add_cli("""
            static-route %s next-hop %s metric %s tag 80
""") (route["to"], route["n-h"], d['metric'])
        dyn.add_cli("""
            exit
exit
""" % d)
    return d

def setup_ies(d):
    d['svc_id'] = dyn.select_free_id("service-id")
    dyn.add_cli("""
configure
service
    ies %(svc_id)s customer 1 create
        service-name "%(inst)s"
        description "%(inst)s"
        no shutdown
    exit
exit
""" % d)
    return {'svc_id':d['svc_id']}

def setup_ethcfm_domain(d):
    dyn.add_cli("""
configure
eth-cfm
    domain %(cfm_domain)s format none level 1
    exit
exit
""" % d)
    return {'cfm_domain':d['cfm_domain']}

def teardown_ethcfm_domain(d):
    dyn.add_cli("""
configure
eth-cfm
    no domain %(cfm_domain)s
    exit
exit
""" % d)

def teardown_ies(d):
    dyn.add_cli("""
configure
service
    ies %(svc_id)s
        shutdown
    exit
    no ies %(svc_id)s
    exit
""" % d)

def teardown_business_ies(d):
    dyn.add_cli("""
configure
router
""")
```

```

        for route in d['routes']:
            dyn.add_cli("""
                no static-route %s next-hop %s
            """)
            dyn.add_cli("""
                exit
            """)
            dyn.add_cli("""
configure
service
    ies %(svc_id)s
        interface "%(if_name)s"
            sap %(sap_id)s
            shutdown
            eth-cfm
                mep 1 domain %(cfm_domain)s association %(svc_id)s
                shutdown
            exit
            no mep 1 domain %(cfm_domain)s association %(svc_id)s
        exit
        no sap %(sap_id)s
        shutdown
    exit
    no interface "%(if_name)s"
    exit
exit
eth-cfm
    domain %(cfm_domain)s
    association %(svc_id)s
    no bridge-identifier %(svc_id)s
    exit
    no association %(svc_id)s
    exit
exit
exit
""")
    % d)
---snip---
```

The IES example has the most attributes. The maximum length of a tagged RADIUS VSA is 246 bytes. If the amount of data is too big to fit into one attribute, simply add a second or third one in the syntax shown above in the RADIUS part. There is no need to separate the attributes exactly at 246 bytes; it can be cut at any position in the list (preferably between two attributes for better readability). Note also that all the parameter VSAs that belong to the same service should have the same tag (":4" in this example).

In case of multiple parameter VSAs, the order in which they are specified is important and must be guaranteed as the concatenation of all the attributes must result in a Python dictionary in the form: "dictionary-name = {...}". The Python script is not aware that multiple attributes were used.

Another difference to the previous examples is that there is not only a reference to the function for the service creation, but also a similar reference to a function for Ethernet Connectivity Fault Management (CFM). Considering that you may want to put all of the Eth-CFM endpoints under the same domain within unique associations, the Eth-CFM domain needs to be created first and torn down as last.

Finally, a different way to provide static-route information is shown at the end of the "setup_business_ies" definition (starting with "for route in d['routes']:"). Also note the difference in how this information is implemented at the end of the "Alc-Dyn-Serv-Script-Params" list. The static routes themselves are defined as a dictionary and thus as many routes as required can be added with this method. Compare this to the VPRN example where a more basic mechanism was used.

As outlined before, dynamic data services can be triggered during the Access-Accept for the control channel but also through a CoA to the control channel Accounting Session ID.

Example 5 – modify an Epipe service using CoA

So far the focus was on service establishment and teardown. It is also possible to change a running dynamic data service using the "modify" function. This will be explained with the previously configured Epipe service.

```
RADIUS attributes in the COA message
Acct-Session-Id = D6E559000000BD5166BF34 #
Alc-Dyn-Serv-SAP-Id:1 = "#:#.1",
Alc-Dyn-Serv-Script-Params:1 = "business_epipe={'ing_qos':'4','egr_qos':'4'}",
Alc-Dyn-Serv-Script-Action:1 = modify,
Alc-Dyn-Serv-Policy:1 = "dynamic-services-2",

Python-part
d = {
---snip---
"business_epipe" : (setup_business_epipe, modify_business_epipe, revert_business_epipe,
teardown_business_epipe)}
dyn.action(d)
---snip---
def modify_business_epipe(d, sd):
    sd['ing_qos'] = d['ing_qos']
    sd['egr_qos'] = d['egr_qos']
    dyn.add_cli("""
configure
service
    epipe %(svc_id)s
        sap %(sap_id)s
            ingress
                qos %(ing_qos)s
            exit
            egress
                qos %(egr_qos)s
            exit
        exit
    exit
exit
"""% sd)
    return sd

def revert_business_epipe(d, sd):
    dyn.add_cli("""
configure
service
    epipe %(svc_id)s
        sap %(sap_id)s
            ingress
                qos %(ing_qos)s
            exit
            egress
                qos %(egr_qos)s
            exit
        exit
    exit
exit
"""% sd)
---snip---
```

Through the function-key in the parameter list (Alc-Dyn-Serv-Script-Params:1 = "business_epipe= ...") and the action attribute of "modify" (Alc-Dyn-Serv-Script-Action:1 = modify), the script will identify the relevant routine to be invoked for the modification (modify_business_epipe). If a modify function is defined, there must also be a definition for a revert function. A revert function cannot be initiated from RADIUS, but it is automatically executed to restore the initial configuration in case the modify script execution fails.

A modify action for an existing service is triggered with a CoA message. For this CoA, either the Accounting Session ID (ASID) of the control channel or the Accounting Session ID of the dynamic data channel can be used. In case the ASID of the control channel is used, the "Alc-Dyn-Serv-SAP-Id" can contain wildcards, as the appropriate port and VLAN information will be taken from the control channel. If the ASID of the dynamic data channel itself is used, the "Alc-Dyn-Serv-SAP-Id" needs to be fully specified, without wildcards. Otherwise the script execution will fail.

For a modify action, the "Alc-Dyn-Serv-Script-Params" only contains the parameters to be changed and does not need any further service identifying information. The service is identified based on the ASID and the "Alc-Dyn-Serv-SAP-Id". Parameters which have been previously received by the setup or an earlier modify function are available in the stored dictionary (sd). Those are combined with the dictionary in the RADIUS message (d). Service modifications which relate to subsequent modifications, or for the service teardown, need to be updated in the stored dictionary so that they can be used in those later actions. This is achieved by the "return sd" command.

As with "manual" provisioned services, the new QoS settings from our example take effect immediately.

A dynamic data service can also be disconnected using a RADIUS Disconnect Message containing the Accounting Session ID of the dynamic data service, or indirectly via a RADIUS Disconnect Message containing the Accounting Session ID of the control channel which would result in a teardown of all associated dynamic data services.

Debugging

It is obvious that the Python scripts need extensive testing in the lab before they are deployed in the field. This testing may require a number of iterations: write the script, testing, verification, improvement and testing again. Every time there is a change in the Python script the node needs to reload the script. This is achieved by a **shutdown** and **no shutdown** of the active script using the command:

configure aaa radius-script-policy <script-policy-name> <primary/secondary> shutdown

configure aaa radius-script-policy <script-policy-name> <primary/secondary> no shutdown

Testing the script may result in some problems if certain aspects may not work as expected (see also debug functions later in this section). It can be that a dynamically created service cannot be removed properly because the teardown script contains errors and the whole service, or fragments of that service, may still exist on the node.

Dynamic data services cannot be edited in normal CLI mode as it may potentially make a later removal of that service through the script impossible. For troubleshooting there is a procedure to manipulate those services during the testing phase, thus avoiding the need to reboot the box to clear the state. The **enable-dynamic-services-config** command allows for the editing dynamic services just like normal services. As this is an action that should only be executed by authorized personnel, the activation of this command is protected by the use of a password, defined under **configure system security password dynsvc-password**.

The **show users** command has been extended to visualize the respective mode ('D' indicates a user is in dynamic service edit mode). A user in dynamic services edit mode cannot modify regular services.

no enable-dynamic-services-config returns the user to normal mode.

To support the creation and the troubleshooting during the test phase the SR OS debug functions have been extended extensively to allow for a detailed review of what is happening in the script and on the CLI.

```
debug dynamic-services
debug dynamic-services scripts
debug dynamic-services scripts event
debug dynamic-services scripts event cli
debug dynamic-services scripts event errors
debug dynamic-services scripts event executed-cmd
debug dynamic-services scripts event state-change
debug dynamic-services scripts event warnings
debug dynamic-services scripts instance
debug dynamic-services scripts instance event
debug dynamic-services scripts instance event cli
debug dynamic-services scripts instance event errors
debug dynamic-services scripts instance event executed-cmd
debug dynamic-services scripts instance event state-change
debug dynamic-services scripts instance event warnings
debug dynamic-services scripts script
debug dynamic-services scripts script event
debug dynamic-services scripts script event cli
debug dynamic-services scripts script event errors
debug dynamic-services scripts script event executed-cmd
debug dynamic-services scripts script event state-change
debug dynamic-services scripts script event warnings
```

It is advised to enable all debug options when starting and then remove more and more debugs options as the script becomes more complete and stable. The debug output gives clear indications about errors in the script or its execution in case something goes wrong.

An additional aid is the use of "print" commands in the Python script itself for certain attributes during the execution of the script. The print output will appear in the debug log. "Print" commands in the Python script should only be used during the testing phase and not in the normal operations mode.

The following command allows the execution of a dynamic services Python script without the need for RADIUS interaction:

tools perform service dynamic-services evaluate-script sap <sap-id> control-session <acct-session-id> action <script-action> [dynsvc-policy <name>]

show service dynamic-services script statistics provides general statistics about script execution.

show service dynamic-services script snippets displays the individual service configuration parts and allows to check if all "snippets" are actually referenced (the counter will increment/decrement with every function call).

In the case of a failed script action a SAP may not be deleted properly and it remains in the configuration as an "orphaned" object.

An orphaned object no longer has any references, which can be seen using **show service dynamic-services root-objects** command where the snippet name and snippet instance is set to "N/A".

Complete setup flow example

To finalize the section about the interaction between RADIUS and the Python script, the complete setup flow for the Epipe example is shown using extracts from the debug output (any missing sequence numbers in the flow below are simple acknowledge messages from RADIUS and are left out to focus on the important information). The debug settings to be used for this output are the following.

```
*A:BNG-1# show debug
debug
```

```

router "Base"
  radius
    packet-type authentication accounting coa
    detail-level medium
  exit
exit
router "management"
  radius
    packet-type authentication accounting coa
    detail-level medium
  exit
exit
dynamic-services
  scripts
    event
      cli
    exit
    instance "dynamic-services-1"
      event
        cli
      exit
    exit
  exit
exit
exit
exit

```

The first sequence in the flow is the Access-Request to the RADIUS server for the control channel. The information provided is that configured as part of the regular ESM configuration.

```

9 2013/04/12 20:47:23.73 UTC MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
Access-Request(1) 172.31.1.2:1812 id 70 len 206 vrid 1 pol authentication-2
USER NAME [1] 24 subscriber12@domain2.com
NAS IP ADDRESS [4] 4 192.0.2.1
SERVICE TYPE [6] 4 Framed(2)
FRAMED PROTOCOL [7] 4 PPP(1)
CHAP PASSWORD [3] 17 1 0xd4b73e0a17c0ad7f03c19bc1db5c291d
CHAP CHALLENGE [60] 41
  0x620fa5f8be193d2066f6abad96c7de2df03986e3421f9733220d9520137b0bf40b30edc9c92bea30a2
VSA [26] 29 DSL(3561)
AGENT CIRCUIT ID [1] 13 circuit-id-12
AGENT REMOTE ID [2] 12 remote-id-12
NAS PORT ID [87] 11 3/2/2:1.100
CALLING STATION ID [31] 17 00:00:64:01:02:03
NAS IDENTIFIER [32] 5 BNG-1
NAS PORT TYPE [61] 4 PPPoEoQinQ(34)

```

If the subscriber can be authenticated and authorized, RADIUS responds with an Access-Accept containing attributes for both the control channel and the dynamic data service.

```

10 2013/04/12 20:47:23.73 UTC MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
Access-Accept(2) id 70 len 211 from 172.31.1.2:1812 vrid 1 pol authentication-2
VSA [26] 14 Alcatel(6527)
  SUBSC ID STR [11] 12 pppoe-user12
FRAMED IP ADDRESS [8] 4 10.2.1.200
VSA [26] 8 Alcatel(6527)
  DYN SERV SAP ID [164] 6 1 #:#.1
VSA [26] 118 Alcatel(6527)
  DYN SERV SCRIPT PARAMS [165] 116 1 business_epipe={'t':('EPipe-CustomerName', 'Customer
Name-Circuit-1', '3', '3', '64496', '192.0.2.5', '192.0.2.1', '3333')}
VSA [26] 21 Alcatel(6527)

```

```
DYN SERV POLICY [167] 19 1 dynamic-services-2
```

The existence of the Dyn Serv VSAs in the response triggers the BNG to start the execution of the Python script, but first the control channel session is completely established and an accounting start message is sent to RADIUS. This is a standard accounting message for ESM subscribers.

```
11 2013/04/12 20:47:23.75 UTC MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Accounting-Request(4) 172.31.1.2:1813 id 108 len 191 vrid 1 pol accounting-2
    STATUS TYPE [40] 4 Start(1)
    NAS IP ADDRESS [4] 4 192.0.2.1
    SERVICE TYPE [6] 4 Framed(2)
    FRAMED PROTOCOL [7] 4 PPP(1)
    FRAMED IP ADDRESS [8] 4 10.2.1.200
    FRAMED IP NETMASK [9] 4 255.255.255.255
    NAS IDENTIFIER [32] 5 BNG-1
    SESSION ID [44] 22 D6E559000000D2516872DB
    MULTI SESSION ID [50] 22 D6E559000000D3516872DB
    EVENT TIMESTAMP [55] 4 1365799643
    NAS PORT TYPE [61] 4 PPPoEoQinQ(34)
    NAS PORT ID [87] 11 3/2/2:1.100
    VSA [26] 29 DSL(3561)
      AGENT CIRCUIT ID [1] 13 circuit-id-12
      AGENT REMOTE ID [2] 12 remote-id-12
    VSA [26] 14 Alcatel(6527)
      SUBSC ID STR [11] 12 pppoe-user12
"
```

Next, the creation of the dynamic data service starts. As this is the first SAP for this service, the script which we reviewed above first creates the service instance.

```
12 2013/04/12 20:47:23.74 UTC MINOR: DEBUG #2001 Base dyn-script cli 1/1
"dyn-script cli 1/1: epipe:EPipe-CustomerName(cli 172 dict 0->31)

configure
  service
    epipe 1000 customer 1 create
      service-name "EPipe-CustomerName"
      description "EPipe-CustomerName"
      no shutdown
    exit
  exit
exit
"
```

Next, the SAP and the SDP are created within this service by the main function.

```
14 2013/04/12 20:47:23.74 UTC MINOR: DEBUG #2001 Base dyn-script cli 1/1
"dyn-script cli 1/1: business_epipe:3/2/2:1.1(cli 418 dict 0->308)

configure
  service
    epipe 1000
      sap 3/2/2:1.1 create
        description "CustomerName-Circuit-1"
        ingress
          qos 3
        exit
      egress
```

```
        qos 3
        exit
    exit
    spoke-sdp-fec 1000 fec 129 aii-type 2 create
        pw-template-bind 2
        saii-type2 64496:192.0.2.1:3333
        taii-type2 64496:192.0.2.5:3333
        no shutdown
    exit
exit
exit
exit
exit
"
```

The service is created and is now active. As two RADIUS accounting destinations are configured in the dynamic services policy a RADIUS Accounting-Start message is sent to each destination to indicate the service is up.

```
16 2013/04/12 20:47:23.76 UTC MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
Accounting-Request(4) 172.31.1.2:1813 id 252 len 294 vrid 1 pol radius-server-policy-2
  STATUS TYPE [40] 4 Start(1)
  NAS IP ADDRESS [4] 4 192.0.2.1
  SESSION ID [44] 22 D6E559000000D4516872DB
  NAS PORT ID [87] 9 3/2/2:1.1
  DELAY TIME [41] 4 0
  NAS IDENTIFIER [32] 5 BNG-1
  EVENT TIMESTAMP [55] 4 1365799643
  MULTI SESSION ID [50] 22 D6E559000000D1516872DB
  USER NAME [1] 24 subscriber12@domain2.com
  VSA [26] 29 DSL(3561)
    AGENT CIRCUIT ID [1] 13 circuit-id-12
    AGENT REMOTE ID [2] 12 remote-id-12
  VSA [26] 117 Alcatel(6527)
    DYN SERV SCRIPT PARAMS [165] 115 business_epipe={'t':('EPipe-CustomerName','CustomerName-
Circuit-1','3','3','64496','192.0.2.5','192.0.2.1','3333')}
"
```

```
15 2013/04/12 20:47:23.76 UTC MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
Accounting-Request(4) 172.31.1.2:1813 id 251 len 294 vrid 1 pol radius-server-policy-2
  STATUS TYPE [40] 4 Start(1)
  NAS IP ADDRESS [4] 4 192.0.2.1
  SESSION ID [44] 22 D6E559000000D4516872DB
  NAS PORT ID [87] 9 3/2/2:1.1
  DELAY TIME [41] 4 0
  NAS IDENTIFIER [32] 5 BNG-1
  EVENT TIMESTAMP [55] 4 1365799643
  MULTI SESSION ID [50] 22 D6E559000000D1516872DB
  USER NAME [1] 24 subscriber12@domain2.com
  VSA [26] 29 DSL(3561)
    AGENT CIRCUIT ID [1] 13 circuit-id-12
    AGENT REMOTE ID [2] 12 remote-id-12
  VSA [26] 117 Alcatel(6527)
    DYN SERV SCRIPT PARAMS [165] 115 business_epipe={'t':('EPipe-CustomerName','CustomerName-
Circuit-1','3','3','64496','192.0.2.5','192.0.2.1','3333')}
"
```

For both RADIUS accounting destinations the interim accounting updates are also configured.

```
21 2013/04/12 20:51:46.69 UTC MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
```



```
Accounting-Request(4) 172.31.1.2:1813 id 173 len 511 vrid 1 pol radius-server-policy-1
  STATUS TYPE [40] 4 Interim-Update(3)
  NAS IP ADDRESS [4] 4 192.0.2.1
  SESSION ID [44] 22 D6E559000000D4516872DB
  NAS PORT ID [87] 9 3/2/2:1.1
  DELAY TIME [41] 4 0
  NAS IDENTIFIER [32] 5 BNG-1
  EVENT TIMESTAMP [55] 4 1365799906
  SESSION TIME [46] 4 125174
  MULTI SESSION ID [50] 22 D6E559000000D1516872DB
  USER NAME [1] 23 subscriber12@domain2.com
  VSA [26] 27 DSL(3561)
    AGENT CIRCUIT ID [1] 12 circuit-id-12
    AGENT REMOTE ID [2] 11 remote-id-12
  VSA [26] 241 Alcatel(6527)
    DYN SERV SCRIPT PARAMS [165] 115 business_epipe={'t':('EPipe-CustomerName','CustomerName-
Circuit-1','3','3','64496','192.0.2.5','192.0.2.1','3333')}
    INPUT_INPROF_OCTETS_64 [19] 10 0x00010000000000000000
    INPUT_OUTPROF_OCTETS_64 [20] 10 0x00010000000000000000
    INPUT_INPROF_PACKETS_64 [23] 10 0x00010000000000000000
    INPUT_OUTPROF_PACKETS_64 [24] 10 0x00010000000000000000
    INPUT_HIGH_OCTETS_OFFER_64 [73] 10 0x00010000000000000000
    INPUT_LOW_PACK_OFFER_64 [76] 10 0x00010000000000000000
    INPUT_HIGH_PACK_OFFER_64 [75] 10 0x00010000000000000000
    INPUT_LOW_OCTETS_OFFER_64 [74] 10 0x00010000000000000000
    INPUT_UNC_PACK_OFFER_64 [78] 10 0x00010000000000000000
    INPUT_UNC_OCTETS_OFFER_64 [77] 10 0x00010000000000000000
    INPUT_HIGH_PACK_DROP_64 [71] 10 0x00010000000000000000
    INPUT_LOW_PACK_DROP_64 [72] 10 0x00010000000000000000
    INPUT_HIGH_OCTETS_DROP_64 [69] 10 0x00010000000000000000
    INPUT_LOW_OCTETS_DROP_64 [70] 10 0x00010000000000000000
    OUTPUT_INPROF_OCTETS_64 [21] 10 0x00010000000000000033c
  VSA [26] 84 Alcatel(6527)
    OUTPUT_OUTPROF_OCTETS_64 [22] 10 0x00010000000000000000
    OUTPUT_INPROF_PACKETS_64 [25] 10 0x00010000000000000000b
    OUTPUT_OUTPROF_PACKETS_64 [26] 10 0x00010000000000000000
    OUTPUT_INPROF_PACK_DROP_64 [81] 10 0x00010000000000000000
    OUTPUT_OUTPROF_PACK_DROP_64 [82] 10 0x00010000000000000000
    OUTPUT_INPROF_OCTS_DROP_64 [83] 10 0x00010000000000000000
    OUTPUT_OUTPROF_OCTS_DROP_64 [84] 10 0x00010000000000000000
"

19 2013/04/12 20:48:56.69 UTC MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
Accounting-Request(4) 172.31.1.2:1813 id 253 len 241 vrid 1 pol radius-server-policy-2
  STATUS TYPE [40] 4 Interim-Update(3)
  NAS IP ADDRESS [4] 4 192.0.2.1
  SESSION ID [44] 22 D6E559000000D4516872DB
  NAS PORT ID [87] 9 3/2/2:1.1
  DELAY TIME [41] 4 0
  NAS IDENTIFIER [32] 5 BNG-1
  EVENT TIMESTAMP [55] 4 1365799736
  SESSION TIME [46] 4 125004
  MULTI SESSION ID [50] 22 D6E559000000D1516872DB
  USER NAME [1] 23 subscriber12@domain2.com
  VSA [26] 27 DSL(3561)
    AGENT CIRCUIT ID [1] 12 circuit-id-12
    AGENT REMOTE ID [2] 11 remote-id-12
  VSA [26] 61 Alcatel(6527)
    DYN SERV SCRIPT PARAMS [165] 115 business_epipe={'t':('EPipe-CustomerName','CustomerName-
Circuit-1','3','3','64496','192.0.2.5','192.0.2.1','3333')}
"
```

The "Stats-Type" in the dynamic service policy (or obtained via RADIUS in a VSA) defines what information is sent back to the accounting server (per server). In this example one was set to Stats-Type "time" and the other to "volume-time". The first accounting message displays the content of "volume-time". A full set of statistics counters per service class are provided for the dynamic service. This is equivalent to the extended accounting statistics also provided in the ESM context. The second accounting message shows the content of "time". No volume statistics counters are provided in this case.

Once the dynamic data services are instantiated they can be displayed with the regular show commands.

```
A:BNG-1# show service service-using
=====
Services
=====
ServiceId   Type      Adm  Opr  CustomerId Service Name
-----
1           VPLS      Up   Up   1          VPLS_For_Capture_SAPs
2           VPRN      Up   Up   1          VPRN_Control_Channel
3           VPRN      Up   Up   1          VPRN_REsidential_Sub
4           IES       Up   Up   1
10          VPRN      Up   Up   1
99          Mirror    Up   Up   1
500         Mirror    Up   Up   1
[1000]      Epipe     Up   Up   1          EPipe-CustomerName
[1001]      VPLS      Up   Up   1          VPLS-CustomerName
[1002]      IES       Up   Up   1          IES-CustomerName
[5000]      IES       Up   Up   1          IES-5000
[9999]      VPRN      Up   Up   1          VPRN-CustomerName
10001       VPLS      Up   Up   1
10002       Epipe     Up   Up   1
---snip---
Matching Services : 20
Dynamic Services : 5, indicated by [<svc-id>]
=====
```

The dynamically created services are shown in the standard service list with their service IDs between brackets. It is possible to filter only the dynamic services using the **origin dyn-script** option.

```
A:BNG-1# show service service-using origin dyn-script
=====
Services
=====
ServiceId   Type      Adm  Opr  CustomerId Service Name
-----
[1000]      Epipe     Up   Up   1          EPipe-CustomerName
[1001]      VPLS      Up   Up   1          VPLS-CustomerName
[1002]      IES       Up   Up   1          IES-CustomerName
[5000]      IES       Up   Up   1          IES-5000
[9999]      VPRN      Up   Up   1          VPRN-CustomerName
Matching Services : 5
Dynamic Services : 5, indicated by [<svc-id>]
=====
```

Similarly, the active SAPs can also be shown with the regular command.

```
A:BNG-1# show service sap-using
```

```
=====
Service Access Points
=====
PortId                      SvcId      Ing.  Ing.  Egr.  Egr.  Adm  Opr
                        QoS    Fltr  QoS   Fltr
-----
3/2/1:*.100                 1          1    none  1     none  Up   Up
3/2/1:*.200                 1          1    none  1     none  Up   Up
3/2/2:*.100                 1          1    none  1     none  Up   Up
[3/2/1:4.100]               2          1    none  1     none  Up   Up
[3/2/2:1.100]               2          1    none  1     none  Up   Up
3/2/2:1000.1000            2          1    none  1     none  Up   Up
[3/2/1:2.200]               3          1    none  1     none  Up   Up
[3/2/1:3.200]               3          1    none  1     none  Up   Up
3/2/1:1001.1001            3          1    none  1     none  Up   Up
3/2/2:500.500              3          1    none  1     none  Up   Up
3/2/2:100.100              4          1    none  1     none  Up   Up
3/2/2:99.99                99          1    none  1     none  Up   Up
[3/2/2:1.1]                 [1000]      3     none  3     none  Up   Up
[3/2/2:1.3]                 [1001]      3     none  3     none  Up   Up
[3/2/2:1.4]                 [1002]      5    ip4+ip6 6    ip4+i*  Up   Up
[3/2/1:4.3]                 [5000]      1     none  1     none  Up   Up
[3/2/2:1.2]                 [9999]      3     ip4    3     ip4    Up   Up
3/2/1:99.99                10001       1     none  1     none  Up   Up
3/2/19:100                 10001       1     none  1     none  Up   Up
3/2/20:100                 10001       1     none  1     none  Up   Up
---snip---
-----
Number of SAPs : 31
-----
Number of Managed SAPs : 4, indicated by [<sap-id>]
-----
Number of Dynamic Service SAPs : 5, indicated by [<sap-id>] [<svc-id>]
-----
=====
* indicates that the corresponding row element may have been truncated.
=====
```

The description at the end of this show command explains how the dynamic services SAPs are displayed. Note that there are managed SAPs created for the control channel as well as dynamic data services SAPs.

If only the SAPs for dynamic data services should be displayed, the command **show service sap-using dyn-script** can be used.

```
A:BNG-1# show service sap-using dyn-script
=====
Service Access Points
=====
PortId                      SvcId      Ing.  Ing.  Egr.  Egr.  Adm  Opr
                        QoS    Fltr  QoS   Fltr
-----
[3/2/2:1.1]                 [1000]      3     none  3     none  Up   Up
[3/2/2:1.3]                 [1001]      3     none  3     none  Up   Up
[3/2/2:1.4]                 [1002]      5    ip4+ip6 6    ip4+i*  Up   Up
[3/2/1:4.3]                 [5000]      1     none  1     none  Up   Up
[3/2/2:1.2]                 [9999]      3     ip4    3     ip4    Up   Up
-----
Number of SAPs : 5
-----
Number of Dynamic Service SAPs : 5, indicated by [<sap-id>] [<svc-id>]
-----
=====
* indicates that the corresponding row element may have been truncated.
=====
```

Dynamic data services CLI is not saved as part of the configuration file. The active dynamic data services configuration is hidden in the output of the "admin display-config" CLI command. To display the dynamic services configuration, use:

- "info" in a configuration CLI context for SR OS Releases prior to 14.0.R1
- "info include-dynamic" in a configuration context for SR OS Release 14.0.R1 or later

Conclusion

RADIUS-based dynamic data services provide an innovative way for business service provisioning. They are created both automatically and instantaneously.

Raw Formatting of DHCPv4/v6 Options in ESM

This chapter provides information about raw formatting of DHCPv4/v6 options in ESM.

Topics in this chapter include:

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

Applicability

This example is applicable to DHCPv4/v6 subscriber-hosts using the routed central office ESM model.

A local DHCPv4/v6 server is used for address/prefix assignment, which implies a DHCP *relay* scenario (as opposed to a DHCP *proxy* scenario where the IP address/prefix is assigned via a RADIUS server or an LUDB).

The information and configuration in this chapter is based on a single homed environment using SR OS Release 12.0.R4.

Overview

An SR OS node supports IP address assignment to its DHCP clients via two IP address assignment authorities:

- DHCP server — In this model the SR OS node behaves as a DHCP relay between the DHCP client and the DHCP server.
- RADIUS/LUDB — In this model the IP address/prefix is assigned via a RADIUS server or an LUDB and the SR OS node internal or external behaves as a proxy between the DHCP client and the non-DHCP aware RADIUS/LUDB.



Note: The term proxy can also refer to the functionality where the DHCP server is used for address assignment. In this case, the SR OS node would hide the DHCP server from the client and pretend to be the DHCP server to the client, passing the DHCP parameters between the client and the server (lease times, etc).

Within these two fundamental address assignment models, there are several mechanisms available on the SR OS node by which DHCP parameters (DHCP options and various parameters within the options) can be passed to the DHCP client during the address assignment phase.

For example, in the RADIUS/LUDB address assignment model, the DHCP parameters can be supplied via RADIUS, LUDB and Python, while in the DHCP server model, the DHCP parameters can also be supplied via the DHCP server itself (in addition to RADIUS, LUDB and Python).

Some of the more commonly used DHCP parameters have their own RADIUS and CLI constructs. For example, a default router has its own RADIUS attributes:

```
Alc-Default-Router (26-6527-18)
```

or even its own CLI keyword:

```
config>router>dhcp>server>pool>subnet>options# default-router
config>service>router>dhcp>server>pool>subnet>options# default-router
config>subscr-mgmt>ludb>ipoe>host>options# default-router
```

Other less common DHCP options can be defined and inserted by the DHCP relay agent using the pre-formatted (IP address, domain, or string) or the non-formatted (hex) custom-option CLI command:

```
config>router>dhcp>server>pool>options# custom-option
config>router>dhcp>server>pool>subnet>options# custom-option
config>router>dhcp6>server>pool>options# custom-option
config>router>dhcp6>server>pool>prefix>options# custom-option
config>service>vprn>dhcp>server>pool>options# custom-option
config>service>vprn>dhcp>server>pool>subnet>options# custom-option
config>service>vprn>dhcp6>server>pool>options# custom-option
config>service>vprn>dhcp6>server>pool>prefix>options# custom-option
config>subscriber-mgmt>ludb>ipoe>host>options# custom-option
```

The most flexible way of configuring DHCP parameters is by means of 'raw' (or hexadecimal) formatting. Any DHCP option can be hexadecimally (raw) formatted via the following RADIUS attributes:

```
Alc-ToClient-Dhcp-Options
Alc-ToClient-Dhcp6-Options
```

and/or via the custom-options CLI commands as outlined above. These options are then passed on to the DHCP client via the DHCP relay agent in the SR OS node.

In addition to raw formatting via RADIUS or CLI, Python scripting can be used to intercept DHCP messages and modify their content.

The focus of this example is to demonstrate how the **raw** DHCP options are formatted via RADIUS. The messages can be optionally pre/post-processed by a Python script in the SR OS node before they are passed on to the DHCP client.

In this example, the following DHCP parameters are passed to the DHCP client using the **Alc-ToClient-Dhcp-Options** and the **Alc-ToClient-Dhcp6-Options** RADIUS attributes:

Table 34: RADIUS inserted raw options

RADIUS	
DHCPv4 ToClient-Dhcp-Options	DHCPv6 ToClient-Dhcp6-Options
(default-)router [3] = 10.10.10.254	DNS server [23] = 2001:db8:1:1:1:1:1:1 2001:db8:1:1:1:1:1:2
DNS server [6] = 172.22.250.250	Domain search list [24] =

RADIUS	
DHCPv4 ToClient-Dhcp-Options	DHCPv6 ToClient-Dhcp6-Options
172.22.250/251	'Nokia.com' 'test.com'
Domain name [15] = 'alcatel.com'	Vendor specific-option [17] = 'custom-test-option'
Custom option [230] = 'custom test option'	
Renew time [58] = 5 min (300sec)	
Rebind time [59] = 6min 40sec (400sec)	

The DHCP parameters in the following DHCP messages are altered by a Python script:

Table 35: Python Modified DHCP Fields

Python	
DHCPv4 (DHCP Request)	DHCPv6 (LDRA DHCP Request)
Lease-time [51] = 8min 20sec (500sec)	IA-NA Preferred-Lifetime = 66min 40sec (4000sec)
	IA-NA Valid-Lifetime = 66min 40sec (4000sec)
	IA-NA Renew-Time (T1) = 33min 20sec (2000sec)
	IA-NA Rebind-Time (T2) = 50min (3000sec)
	IA-PD Preferred-Lifetime = 66min 40sec (4000sec)
	IA-PD Valid-Lifetime = 66min 40sec (4000sec)
	IA-PD Renew-Time (T1) = 33min 20sec (2000sec)
	IA-PD Rebind-Time (T2) = 50min (3000sec)

The following DHCP parameters are configured via CLI in the SR OS node DHCPv4/v6 server:

Table 36: CLI Inserted DHCP Options

CLI DHCP Server Pool/prefix Options	
DHCPv4	DHCPv6
DNS server [6] = 172.22.250.253	DNS server [23] = 2001:db8:1:1:1:1:1:3
Custom option [231] = 'dhcp injected custom option 231'	Custom option [232]= 'v6 custom option 232'

CLI DHCP Server Pool/prefix Options	
DHCPv4	DHCPv6
	IA-NA Preferred-Lifetime = 20min (1200sec)
	IA-NA Valid-Lifetime = 20min (1200sec)
	IA-NA Renew-Time (T1) = 10 min (600 sec)
	IA-NA Rebind-Time (T2) = 15min (900 sec)
	IA-PD Preferred-Lifetime = 20min (1200sec)
	IA-PD Valid-Lifetime = 20min (1200sec)
	IA-PD Renew-Time (T1) = 10min (600 sec)

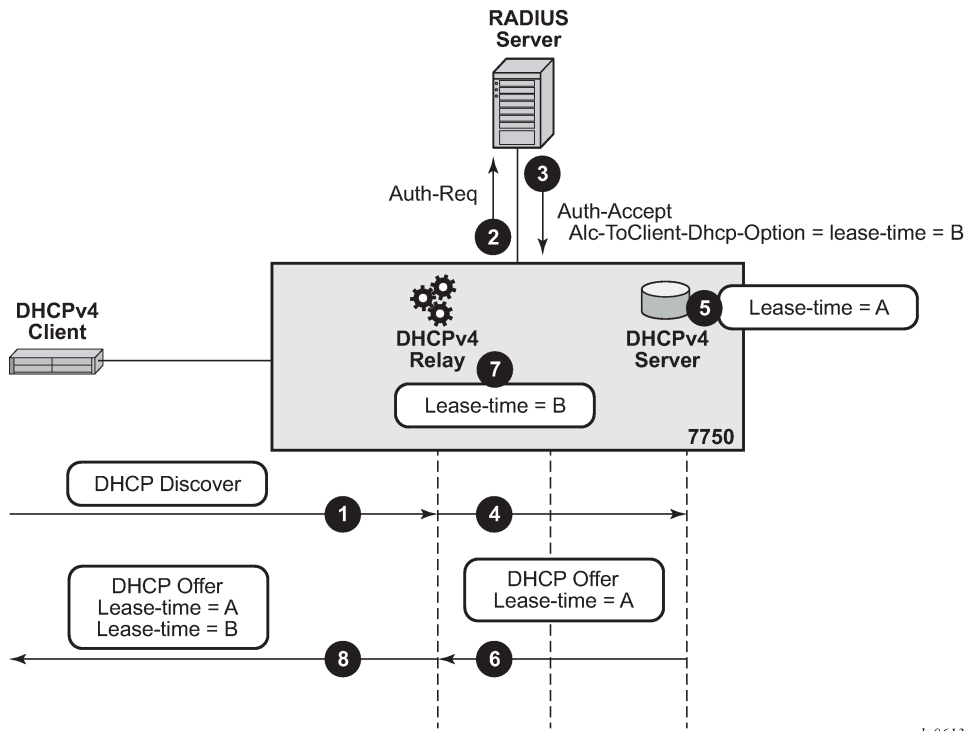
The RADIUS injected raw options are prepended by the DHCP relay agent in the SR OS node to any existing DHCP options already present in the DHCP message before being sent to the client. The existing options could be generated by the DHCP server (internal or external) or by the LUDB. No check is performed on the outgoing DHCP message towards the client in order to verify whether any of the RADIUS inserted options are already present in the DHCP message. This could potentially lead to duplication of DHCP options in the outgoing DHCP messages in case that the same option is inserted via the DHCP server and via RADIUS. To prove the point, this example supplies the same DHCP option (with different values) via multiple sources (RADIUS and CLI).

Configuration of DHCP lease related times requires closer examination. In DHCPv4, the DHCP lease-time option (51) is always supplied by the DHCPv4 server (this cannot be disabled). In case the lease-time is also supplied via RADIUS in an Alc-ToClient-Dhcp-Options VSA, the client would receive two lease-times for the same IP address. This can lead to unpredictable behavior not only on the client side but also on the SR OS node DHCPv4 server side since the DHCPv4 server (and the SR OS node DHCPv4 relay agent) creates the lease state only for the lease-time supplied by the DHCP server, and ignores the one supplied via RADIUS or LUDB. This scenario is shown in [Figure 150: DHCPv4 Lease-Time Inserted by RADIUS and DHCPv4 Server](#):

1. DHCP Discover arrives.
2. Radius authentication is triggered.
3. RADIUS returns lease-time value 'B' (Alc-ToClient-DHCP-Option) in Authentication-Accept message.
4. DHCP Discover is forwarded by the DHCP relay agent to the DHCP server.
5. DHCP server offers an IP lease with the configured lease-time of 'A'.
6. The DHCP offer is sent to the DHCP relay agent.
7. The DHCP relay agent appends the lease-time 'B' supplied by RADIUS to the DHCP message.
8. The DHCP relay forwards the message to the DHCP client with both lease-times 'A' and 'B'.

Note that the example in [Figure 150: DHCPv4 Lease-Time Inserted by RADIUS and DHCPv4 Server](#) does not represent a typical deployment case. This example is solely chosen to clarify the behavior in SR OS nodes.

Figure 150: DHCPv4 Lease-Time Inserted by RADIUS and DHCPv4 Server



al_0613

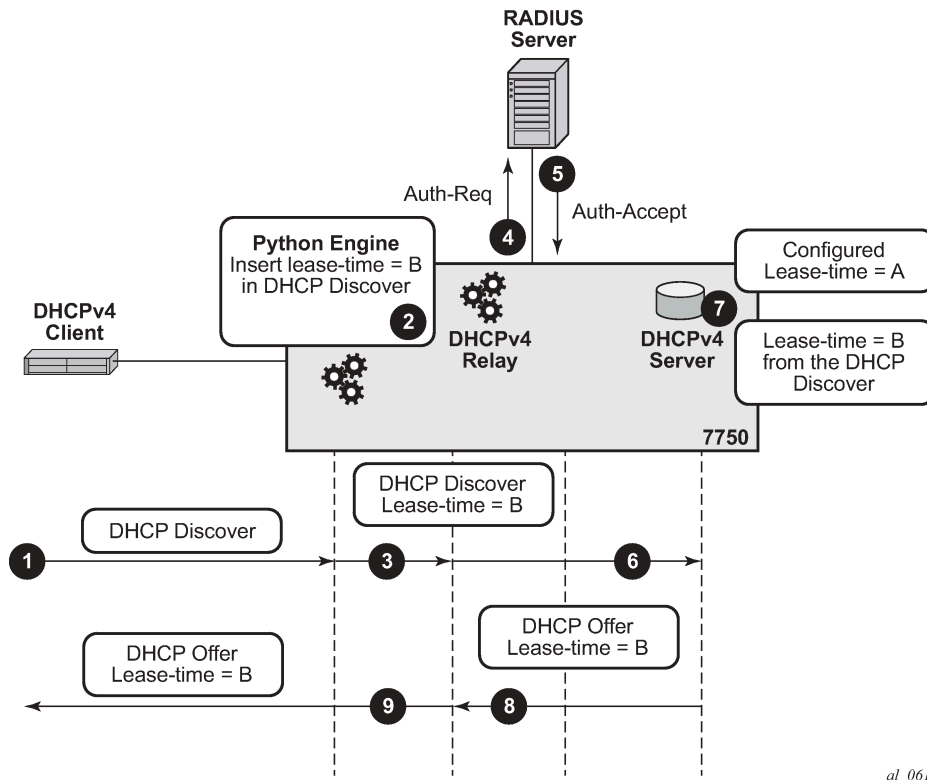
To ensure DHCPv4 lease time unambiguity, the lease-time should be supplied by a single source, in this case by the DHCPv4 server.

Since this eliminates RADIUS as a source of the DHCPv4 lease-time, an alternate method operating on the *raw level* is used to influence the automatic selection of the lease-time in the DHCPv4 server. This alternate method relies on the fact that the DHCPv4 server accepts hints received from the client as to what the desired lease-time should be. In other words, if the client sends the option 51 (lease-time) with a specific value, the SR OS node DHCPv4 server will honor this hint, as long as this value is within the configured range of values specified in the DHCP server. To demonstrate this behavior, a Python script is invoked upon receipt of a DHCPv4 Request message during the IP address assignment process (DORA – Discover-Offer-Request-Ack). The Python script inserts a new option 51 with the desired value for the lease-time. The DHCPv4 server honors this hint from the client and it returns the requested lease-time back to the client. This scenario is shown in [Figure 151: Python Injected Hint for Lease-Time](#).

1. DHCP Discover arrives.
2. DHCP Discover is intercepted by the Python processing engine and the lease-time 'B' is inserted in DHCP Discover message. This is then used as a hint to the DHCP server.
3. DHCP Discover message is sent to the DHCP relay agent.
4. RADIUS authentication is triggered.
5. User is authenticated. This time lease-time is not returned via RADIUS.
6. DHCP Discover is forwarded to the DHCP server.
7. The DHCP server honors the hint from the DHCP Discover and offers lease-time 'B', even though the server is configured with lease-time 'A'.

8. The DHCP server replies with a DHCP Offer message.
9. DHCP Offer is forwarded by the DHCP relay agent to the client.

Figure 151: Python Injected Hint for Lease-Time



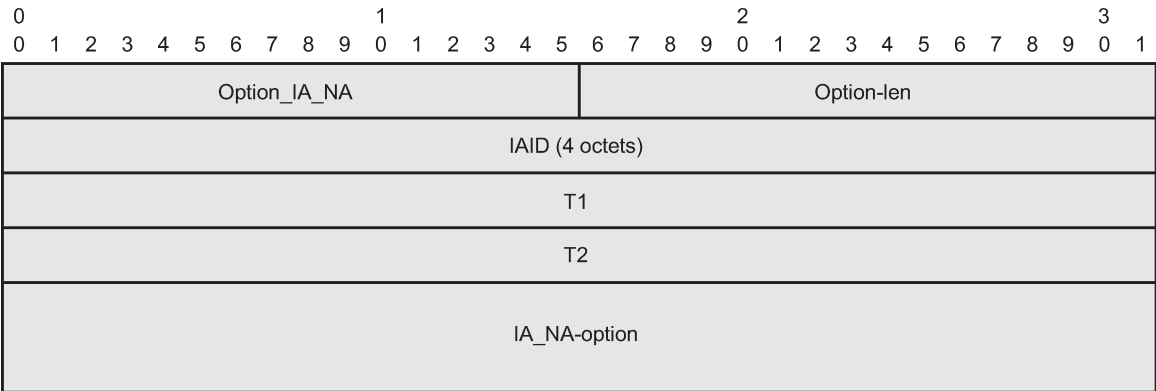
al_0614

By default the local DHCP server does **not** inject Renew (T1) and Rebind (T2) times so these two timers can still be supplied via RADIUS without duplication by the local DHCP server.

When it comes to lease-time related parameters, the behavior of the DHCPv6 server is different from the behavior of the DHCPv4 server.

DHCPv6 lease related timers are **not** DHCP options. Instead, they are parameters within the IPv6 addressing option. An IPv6 address or prefix is assigned to the client via the IA-NA or IA-PD option, which contains additional parameters (which are not considered options) such as the IP address/prefix and the lease related timers. [Figure 152: Format of the IA-NA Option](#) shows the IA-NA option that carries the T1/T2 parameters.

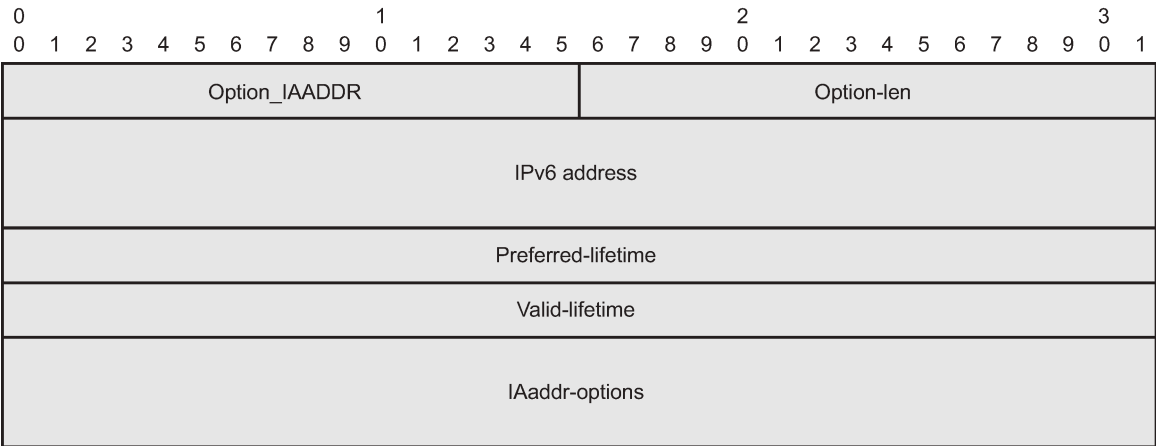
Figure 152: Format of the IA-NA Option



al_0615

The format of the IA address option is shown in [Figure 153: Format of the IA Address Option](#). This option carries preferred and valid lifetimes.

Figure 153: Format of the IA Address Option



al_0616

In this example, the IPv6 address/prefix is provided by the local DHCPv6 server and as such, RADIUS cannot modify the parameters within the DHCPv6 options supplied by the DHCP server. Therefore, the desired IPv6 lease timers (preferred-life time, valid-lifetime, renew-time[T1], rebind-time[T2]) are part of the IPv6 pool configuration in the DHCPv6 server.

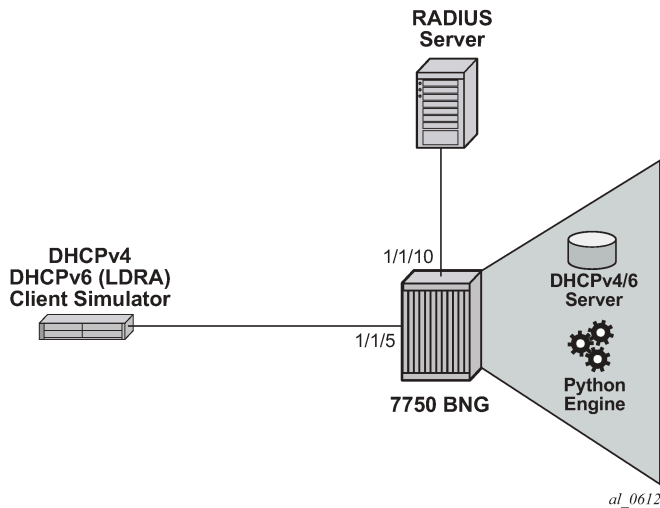
Alternatively Python can be used to intercept the outgoing DHCPv6 message and then change the timers within the IA-NA and IA-PD options. Although this would configure the lease timers for the client, the action of modifying the outgoing DHCP6 message occurs after the DHCPv6 server processing. This would result in different lease times in the client and the DHCPv6 server, without any intermediary between them (such as a DHCPv6 Proxy) to deal with the differences.

For consistency purposes with the DHCPv4 example, a Python script processes the incoming DHCPv6 message (DHCPv6 Request) altering the lease timers (preferred/valid/renew/rebind) as a hint to the DHCPv6 server to request those values. However, the SR OS node DHCPv6 server does not honor those hints and uses its own values (default or configured) instead.

Configuration

The topology is shown in [Figure 154: Topology](#).

Figure 154: Topology



Access Ethernet Port with QinQ Encapsulation

```
configure port 1/1/5
  ethernet
    mode access
    encap-type qinq
  exit
  no shutdown
```

Capture SAP

A capture SAP is used to dynamically detect VLAN ID(s) in incoming DHCP (trigger) packets. This example uses RADIUS authentication along with Python scripting for DHCP message processing and therefore the authentication and Python policies must be configured under the capture SAP.

```
configure service vpls 10
  sap 1/1/5:1.* capture-sap create
    description "circuit-id authentication"
    trigger-packet dhcp dhcp6
    dhcp-python-policy "acg"
    dhcp6-python-policy "acg"
    authentication-policy "rad"
```

MSAP-Policy Configuration

The MSAP-policy defines the anti-spoofing mode which is set to next-hop MAC (nh-mac) in this example. It also defines the default subscriber management parameters in case that they are not supplied via LUDB or RADIUS.

MSAP-policy configuration is mandatory when a capture-SAP is deployed. In this example, the MSAP-policy name is supplied via RADIUS:

```
Alc-MSAP-Policy = "msaps"

configure subscriber-mgmt msap-policy "msaps"
  sub-sla-mgmt
    sub-ident-policy "sub_ident_pol"
    multi-sub-sap limit 500
  exit
  ies-vprn-only-sap-parameters
    anti-spoof nh-mac
  exit
```

Subscriber-Interface and Group-Interface Configuration

In this example the subscriber-interface is a 'numbered interface' in which the interface IPv4 address and the interface IPv6 prefixes are explicitly configured. The IPv4 address is used as the default-gateway by the IPoE attached clients. The IPv4 subnet to which this address belongs and the configured IPv6 prefixes are used for routing aggregation and are treated as local subnets/prefixes in the SR OS node routing table.

The managed (dynamic) SAPs are created under the group-interface which contains the reference to the authentication-policy name, the Python script, the v4/6 policy names and the DHCPv4/v6 relay related configuration settings (for example, a reference to DHCP servers). Both the authentication-policy name and the Python policy name referenced under the group-interface must match those configured under the capture-SAP.

```
configure service vprn 1
  subscriber-interface "intl-1" create
  address 10.10.10.254/24 # Numbered IPv4 subscriber interface.
  ipv6
    delegated-prefix-len 54
    subscriber-prefixes
      prefix 2001:db8:3::/48 pd # Numbered IPv6 subscriber interface.
      prefix 2001:db8:4::/48 wan-host # Numbered IPv6 subscriber interface.
    exit
  exit
  group-interface "g1-1" create
  ipv6
    router-advertisements
      no shutdown
    exit
    dhcp6
      python-policy "acg" # Python script for DHCPv6 messages.
      relay
        server 2001:db8::1001 # IPv6 address of the DHCPv6 server.
        client-applications dhcp
        no shutdown
      exit
    exit
  exit
```

```

dhcp
python-policy "acg"          # Python script for DHCPv4 messages.
option
    action keep              # Keep option82 in the received DHCP packet.
    vendor-specific-option
        pool-name # Pool-name obtained via RADIUS (or LUDB) will be passed
                  # via DHCP relay to the local DHCP server. This name
                  # will be used for pool selection in DHCPv4 server.
    exit
exit
server 192.168.100.1 # IPv4 address of the DHCPv4 server.
lease-populate 100   # Maximum number of DHCPv4 lease under each
                    # SAP of the group-interface.
client-applications dhcp
no shutdown
exit
authentication-policy "rad" # RADIUS authentication policy.
exit
exit

```

Loopback (DHCP) Interface Configuration

The loopback interface is used for the DHCPv4/v6 server binding. It is configured with the IPv4/IPv6 addresses which are referenced from the DHCP relay configuration under the group-interface.

```

configure service vprn 1
interface "loopback1-1" create
    address 192.168.100.1/32      # IPv4 address of the DHCPv4 server.
    ipv6
        address 2001:db8::1001/128 # IPv6 address of the DHCPv6 server.
        local-dhcp-server "v6"     # Binding of the DHCPv6 server
                                   # to this interface.
    exit
    local-dhcp-server "v4"        # Binding of the DHCPv4 server
                                   # to this interface.
    loopback
exit

```

DHCPv4/6 Server Configuration

The local DHCP server configuration contains the pool selection method, pool information and DHCP options which are passed to the DHCP client at IP address/prefix assignment time.

```

configure service vprn 1
dhcp
    local-dhcp-server "v4"
        use-pool-from-client # Pool-name received in the DHCP messages
                             # sent by the DHCP relay. The pool-name
                             # is used in pool selection.
        pool "non-shared-left"
            options
                dns-server 172.22.250.253 # DHCPv4 option passed on to the client.
                custom-option 231 string "dhcp server injected custom option 231"
            exit
            subnet 10.10.10.0/24 create
                address-range 10.10.10.5 10.10.10.100 # IPv4 address range available
                                                         # for address allocation.

```

```
        exit
    exit
exit
dhcp6
    local-dhcp-server "v6"
        use-pool-from-client
        pool "pd-left" create
        options
            dns-server 2001:db8:1:1:1:1:3
            custom-option 232 string "v6 custom option 232"
        exit
        prefix 2001:db8:4::/48 pd      # IPv6 prefix range available for delegated
                                      # prefix allocation by this DHCPv6 server.
            preferred-lifetime min 20 # Preferred lifetime of the allocated
                                      # delegated prefix.
            rebind-timer min 15        # Rebind (T2) time of the allocated
                                      # delegated prefix.
            renew-timer min 10         # Renew (T1) time of the allocated
                                      # delegated prefix.
            valid-lifetime min 20      # Valid lifetime of the allocated
                                      # delegated prefix.
        exit
    exit
    pool "wan-left" create
        options
            dns-server 2001:db8:1:1:1:1:3
            custom-option 232 string "v6 custom option 232"
        exit
        prefix 2001:db8:3::/56 wan-host
            preferred-lifetime min 20 # Preferred lifetime of the
                                      # allocated IPv6 address.
            rebind-timer min 15        # Rebind (T2) time of the
                                      # allocated IPv6 address.
            renew-timer min 10         # Renew (T1) time of the
                                      # allocated IPv6 address.
            valid-lifetime min 20      # Valid lifetime of the
                                      # allocated IPv6 address.
        exit
    exit
    no shutdown
exit
exit
```

RADIUS Authentication-Policy Configuration

The RADIUS authentication-policy is referenced under the capture-sap and under the group-interface configuration.

```
authentication-policy "rad" create
    password "ALU" hash2
    radius-authentication-server
    router "Base"
        server 1 address 192.168.114.1 secret "ALU" hash2
    exit
    user-name-format circuit-id
    include-radius-attribute
        circuit-id
        remote-id
        nas-port-id
        nas-identifier
```

```
    exit
exit
```

Subscriber-Identification Policy

The subscriber-identification policy in this example defines a mapping method between the subscriber strings and the predefined subscriber profiles (*sub* and *sla*) locally configured on the SR OS node. In our example the subscriber strings (*sub* and *sla*) are provided via RADIUS and are directly mapped to the preconfigured sub-profiles and sla-profiles with the matching names.

The subscriber-identification policy can be configured with default subscriber profiles in case the strings are not explicitly obtained via other means (RADIUS, LUSB, Python or statically provisioned). Subscriber-identification policy configuration is mandatory.

```
sub-ident-policy "sub_ident_pol" create
  sub-profile-map
    use-direct-map-as-default
  exit
  sla-profile-map
    use-direct-map-as-default
  exit
```

Sla-Profile and Sub-Profile Configuration

The following is the configuration of the sub-profile and the sla-profile which are used to setup the subscriber-host. The sla and sub profiles are mandatory when creating subscriber-hosts in SR OS node.

```
sla-profile "sla-profile-1" create
  ingress
    qos 2
  exit
exit
  egress
    qos 2
  exit
exit
exit

sub-profile "sub-profile-1" create
exit
```

Python-Policy Configuration

The python-policy defined below is applied under the capture-sap and under the group-interface. It references the python-script command which defines the location of the script. A python-policy specifies the DHCP messages along with the direction to which the script processing applies.

The DHCPv4 script in this example is applied to incoming DHCPv4 Request messages. The python script inserts the lease-time option in the DHCPv4 Request message as a hint to the DHCPv4 server.

Similar logic is applied to the incoming Lightweight DHCPv6 Relay Agent

(LDRA) DHCPv6 messages where IA-NA and IA-PD related lease times are altered. Note that in the DHCPv6 case the local DHCPv6 server does not honor the hint and therefore the lease related times are explicitly configured in the DHCPv6 server.

```
python-script "acg" create
    action-on-fail passthrough #In case of python script failure, do not drop the
                               message but instead continue with message processing
                               in 7750.
    primary-url "ftp://a.b.c.d/pub/configs/alu/SIMS/acg.py"
    no shutdown
exit
python-script "acg6" create
    action-on-fail passthrough
    primary-url "ftp://a.b.c.d/pub/configs/alu/SIMS/acg6.py"
    no shutdown
exit
python-policy "acg" create #Python policy that is applied under the capture-sap and
                           under the group-interface.
    dhcp request direction ingress script "acg"
    dhcp6 relay-forward direction ingress script "acg6"
exit
```

Python Script Configuration

In this example the Python script is located in an external location and downloaded to the SR OS node once the python-script CLI node is enabled (**no shutdown**).

The DHCPv4 Python script has exception code included (try/except statements). This makes script debugging easier in case one of the commands in the script fails.

For simplicity reasons, the exception code is removed from the DHCPv6 Python script. Note that in real deployments it is recommended for the exception code to be included in all Python scripts.

DHCPv4 Python Script:

```
from alc import dhcpv4
try:
    myopt = dhcpv4.getOptionList()
    if myopt != []:
        print "option-list ", repr(myopt)
        print "\n"
except Exception:
    print "Can't retrieve DHCP options"
#lease 500s 8min 20sec
try:
    dhcpv4.set(51,('\x00\x00\x01\x04', #Insert the lease-time (opt51) in the incoming
                                     DHCPv4 request as a hint to the DHCPv4 server.
except Exception:
    print "Can't set time lease"
```

DHCPv6 Python Script:

```
from alc import dhcpv6
import struct
packet = dhcpv6.get_relaymsg()# Extract the original DHCPv6 packet within LDRA.

msgType = ord(packet.msg_type) # Get the message type.
ia_na = packet.get_iana()      # Store the IA-NA option for further processing later on.
ia_pd = packet.get_iapd()      # Store the IA-PD option for further processing later on.
```

```
if msgType == 3: # If the message in the LDRA packet is DHCPv6 Request, insert the lease
    related times in address/prefic options.

    ia_na[0][1] = '\x00\x00\x07\xd0' # Set the renew time (T1) in IA-NA to 2000sec.
    ia_na[0][2] = '\x00\x00\x0b\xb8' # Set the rebind time (T2) in IA-NA to 3000sec.
    ia_na[0][3][5][0][1] = '\x00\x00\x0f\xa0' # Set the preferred time in IA-NA to
                                                # 4000sec.
    ia_na[0][3][5][0][2] = '\x00\x00\x0f\xa0' # Set the valid time in IA-NA to 4000sec.
    packet.set_iana(ia_na) # Update the stored packet with the new values for IA-NA.

    ia_pd[0][1] = '\x00\x00\x07\xd0' # Set the renew time (T1) in IA-PD to 2000sec.
    ia_pd[0][2] = '\x00\x00\x0b\xb8' # Set the rebind time (T2) in IA-PD to 3000sec.
    ia_pd[0][3][26][0][0] = '\x00\x00\x0f\xa0' # Set the preferred time in IA-PD to
                                                # 4000sec.
    ia_pd[0][3][26][0][1] = '\x00\x00\x0f\xa0' # Set the valid time in IA-PD to 4000sec.
    packet.set_iapd(ia_pd) # Update the stored packet with the new values for IA-PD.
    dhcpv6.set_relaymsg(packet) # Insert the packet in the LDRA message.
```

RADIUS Access-Accept

Upon authentication, RADIUS returns the Access-Accept message with the following attributes:

```
Sending Access-Accept of id 66 to 192.168.114.2 port 64384
Alc-Subsc-Prof-Str = "sub-profile-1"
Alc-SLA-Prof-Str = "sla-profile-2"
Alc-MSAP-Interface = "gl-1"
Alc-MSAP-Policy = "msaps"
Alc-MSAP-Serv-Id = 1
Framed-Pool = "non-shared-left"
Framed-IPv6-Pool = "wan-left"
Alc-Delegated-IPv6-Pool = "pd-left"
Alc-ToClient-Dhcp-Options += 0x03040a0a0afe
Alc-ToClient-Dhcp-Options += 0x0608ac16fafaac16fafb
Alc-ToClient-Dhcp-Options += 0x0f0b616c636174656c2e636f6d
Alc-ToClient-Dhcp-Options += 0xe612637573746f6d2074657374206f7074696f6e
Alc-ToClient-Dhcp-Options += 0x3a040000012c
Alc-ToClient-Dhcp-Options += 0x3b0400000190
Alc-ToClient-Dhcp6-Options +=
0x0011001a0000197f00e60012637573746f6d2074657374206f7074696f6e
Alc-ToClient-Dhcp6-Options +=
0x0017002020010db80001000100010001000100010001000100010001000100010002
Alc-ToClient-Dhcp6-Options +=
0x0018001e0e616c636174656c2d6c7563656e7403636f6d00047465737403636f6d
```

It is possible to concatenate multiple DHCP options in a single RADIUS Alc-ToClient-DHCP6-Option but for clarity each option is in a separate attribute in this example.

The following table contains the explanation of the DHCP options inserted via RADIUS:

Table 37: DHCP options inserted via RADIUS

Alc-ToClient-Dhcp-Options += 0x03040a0a0afe (default) router (3) = 10.10.10.254
Alc-ToClient-Dhcp-Options += 0x0608ac16fafaac16fafb

dns (6) = 172.16.250.250 172.16.250.251
Alc-ToClient-Dhcp-Options += 0x0f0b616c636174656c2e636f6d domain-name (15) = alcatel.com
Alc-ToClient-Dhcp-Options += 0xe612637573746f6d2074657374206f7074696f6e custom -option (230) = "custom test option"
Alc-ToClient-Dhcp-Options += 0x3a040000012c renewal time T1 (58) = 300s (5min)
Alc-ToClient-Dhcp-Options += 0x3b0400000190 rebind time T2 (59) = 400s (6min 40sec)
Alc-ToClient-Dhcp6-Options += 0x0011001a0000197f00e60012637573746f6d2074657374206f7074696f6e v6 vendor option (17) [opt-id(2) len(2) entp-id(4) vopt-code(2) vlen(2) vdata] = 17 26 6527 230 18 "custom test option"
Alc-ToClient-Dhcp6-Options += 0x0017002020010db800010001000100010001000120010db8000100010001000100010002 dns servers (23) [opt-id(2) len(2) servers-v6@] = 23 32 2001:0db8:0001:0001:0001:0001:0001:0001 2001:0db8:0001:0001:0001:0001:0001:0002
Alc-ToClient-Dhcp6-Options += 0x0018001e0e616c636174656c2d6c7563656e7403636f6d00047465737403636f6d0 domain list (24) = Nokia.com test.com [formatting as described in section 3.1 of RFC 1035 (as referenced by RFC 4704 and RFC 3315)].

Results and Verification

The results are verified via debug output and show commands on the SR OS node, and also via pcap (Wireshark® packet capture) files on the DHCP client side.

Debug output on the SR OS node is enabled for DHCPv4/6 messages and for the Python script.

The DHCP debug output shows the options sent to the client in the DHCPv4/6 Ack/Reply messages.

The following commands enables debugging information to be sent to the current telnet/ssh session:

```
*A:BNGL# configure log
*A:BNGL>config>log# info
-----
log-id 50 # Capturing and displaying debug output is configured via log.
from debug-trace # Capture debug output.
to session      # Output the debug to the current tcp/ssh session.
exit
-----
```

The following commands enable DHCP related debugging:

```
*A:BNGL>config>log# show debug
```

```
debug
router "1"
ip
    dhcp
        detail-level high
        mode egr-ingr-and-dropped
    exit
    dhcp6
        mode egr-ingr-and-dropped
        detail-level high
    exit
exit
local-dhcp-server "v4"
    detail-level high
    mode egr-ingr-and-dropped
exit
local-dhcp-server "v6"
    detail-level high
    mode egr-ingr-and-dropped
exit
exit
```

DHCPv4 Results

The following output displays the DHCPv4 Request message as it was received by the SR OS node DHCP server.

This message has been modified by the Python script on ingress and the lease-time option [51] has been inserted as a hint to the DHCPv4 server.

Option [82] is partially added by the **access-node** (relay-agent —> circuit-id and remote-id) and partially by the internal SR OS node DHCP-relay (pool name).

```
32830 2014/07/24 03:02:46.44 WEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: v4
Rx DHCP Request

ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
siaddr: 0.0.0.0          giaddr: 10.10.10.254
chaddr: 00:00:65:01:03:01  xid: 0x159dd536

DHCP options:
[82] Relay agent information: len = 42
    [1] Circuit-id: ds-left
    [2] Remote-id: remote0
    [9] Vendor-Specific info: len = 22
        Enterprise [6527] : len = 17
        [13] dhcpPool: non-shared-left
[53] Message type: Request
[54] DHCP server addr: 192.168.100.1
[50] Requested IP addr: 10.10.10.34
[51] Lease time: 500
[255] End
```

The next output captures the DHCPv4 ACK message (within the SR OS node) that is on its way to the client.

It can be observed that the DHCPv4 server inserted options are listed first:

- Opt[82] is echoed back by SR OS node DHCPv4 server
- Opt[53], [54], [51] and [1] are by default inserted by the local DHCPv4 server and they cannot be disabled. The value for the lease-time [51] is set by the Python script.
- The next two options ([6] and [231]) are the options configured explicitly in the DHCPv4 server ([Table 36: CLI Inserted DHCP Options](#)).

The remaining options (with the exception of the **end** [255] option) are provided by RADIUS and they appear in the exact same order as they appear in the RADIUS Alc-ToClient-Dhcp-Options attributes ([Table 34: RADIUS inserted raw options](#)).

There are two options [6] since they are inserted by both DHCP and RADIUS server.

Custom options [231] and [230] are decoded in [Table 34: RADIUS inserted raw options](#) and [Table 36: CLI Inserted DHCP Options](#).

```
32834 2014/07/24 03:02:46.44 WEST MINOR: DEBUG #2001 vprn1 PIP
"PIP: DHCP
instance 2 (1), interface index 11 (g1-1),
  transmitted DHCP Boot Reply to Interface g1-1 (1/1/5:1.3) Port 68

H/W Type: Ethernet(10Mb) H/W Address Length: 6
ciaddr: 0.0.0.0 yiaddr: 10.10.10.34
siaddr: 192.168.100.1 giaddr: 10.10.10.254
chaddr: 00:00:65:01:03:01 xid: 0x159dd536

DHCP options:
[82] Relay agent information: len = 18
    [1] Circuit-id: ds-left
    [2] Remote-id: remote0
[53] Message type: Ack
[54] DHCP server addr: 192.168.100.1
[51] Lease time: 500
[1] Subnet mask: 255.255.255.0
[6] Domain name server: 172.22.250.253
[231] Unknown option: len = 38, value = 64 68 63 70 20 73 65 72 76 65 72
20 69 6e 6a 65 63 74 65 64 20 63 75 73 74 6f 6d 20 6f 70 74 69 6f 6e 20 32
33 31
[3] Router: 10.10.10.254
[6] Domain name server: length = 8
    172.22.250.250
    172.22.250.251
[15] Domain name: alcatel.com
[230] Unknown option: len = 18, value = 63 75 73 74 6f 6d 20 74 65 73 74
20 6f 70 74 69 6f 6e
[58] Renew timeout: 300
[59] Rebind timeout: 400
[255] End
```

The Wireshark® output shown on the next page is captured at the client side (N2X Ixia) and it effectively mirrors what is shown in the debug output.

```

❏ Bootstrap Protocol
  Message type: Boot Reply (2)
  Hardware type: Ethernet (0x01)
  Hardware address length: 6
  Hops: 0
  Transaction ID: 0x159dd536
  Seconds elapsed: 0
  ❏ Bootp flags: 0x8000 (Broadcast)
  Client IP address: 0.0.0.0 (0.0.0.0)
  Your (client) IP address: 10.10.10.34 (10.10.10.34)
  Next server IP address: 192.168.100.1 (192.168.100.1)
  Relay agent IP address: 10.10.10.254 (10.10.10.254)
  Client MAC address: NetworkG_01:03:01 (00:00:65:01:03:01)
  Client hardware address padding: 00000000000000000000
  Server host name not given
  Boot file name not given
  Magic cookie: DHCP
  ❏ Option: (53) DHCP Message Type
    Length: 1
    DHCP: ACK (5)
  ❏ Option: (54) DHCP Server Identifier
  ❏ Option: (51) IP Address Lease Time
    Length: 4
    IP Address Lease Time: (500s) 8 minutes, 20 seconds
  ❏ Option: (1) Subnet Mask
  ❏ Option: (6) Domain Name Server
    Length: 4
    Domain Name Server: 172.22.250.253 (172.22.250.253)
  ❏ Option: (231) Private
    Length: 38
    Value: 646863702073657276657220696e6a656374656420637573...
  ❏ Option: (82) Agent Information Option
  ❏ Option: (3) Router
    Length: 4
    Router: 10.10.10.254 (10.10.10.254)
  ❏ Option: (6) Domain Name Server
    Length: 8
    Domain Name Server: 172.22.250.250 (172.22.250.250)
    Domain Name Server: 172.22.250.251 (172.22.250.251)
  ❏ Option: (15) Domain Name
    Length: 11
    Domain Name: alcatel.com
  ❏ Option: (230) Private
    Length: 18
    Value: 637573746f6d2074657374206f7074696f6e
  ❏ Option: (58) Renewal Time Value
    Length: 4
    Renewal Time value: (300s) 5 minutes
  ❏ Option: (59) Rebinding Time value
    Length: 4
    Rebinding Time value: (400s) 6 minutes, 40 seconds
  ❏ Option: (255) End

```

The show command for the DHCP-relay lease state only displays the well known options inserted by the DHCPv4 server. The custom option inserted by the DHCPv4 server and any of the RADIUS supplied options are not kept as part of the DHCP-relay lease state.

```

*A:BNGL# show service id 1 dhcp lease-state detail
=====
DHCP lease states for service 1
=====
Service ID      : 1
IP Address      : 10.10.10.34
Client HW Address : 00:00:65:01:03:01

```

```
Subscriber-interface : int1-1
Group-interface      : g1-1
SAP                  : [1/1/5:1.3]
Up Time              : 0d 00:10:46
Remaining Lease Time : 0d 00:07:35
Remaining SessionTime: N/A
Persistence Key      : N/A

Sub-Ident            : "ds-left"
Sub-Profile-String   : "sub-profile-1"
SLA-Profile-String   : "sla-profile-2"
App-Profile-String   : ""
Lease ANCP-String    : ""
Lease Int Dest Id    : ""
Category-Map-Name    : ""

Lease Info origin    : DHCP

Ip-Netmask          : 255.255.255.0
Broadcast-Ip-Addr    : N/A
Default-Router        : N/A
Primary-Dns         : 172.22.250.253
Secondary-Dns         : N/A
Primary-Nbns          : N/A
Secondary-Nbns        : N/A

ServerLeaseStart      : 07/24/2014 03:02:46
ServerLastRenew       : 07/24/2014 03:12:46
ServerLeaseEnd        : 07/24/2014 03:21:06
Session-Timeout       : N/A
Lease-Time          : 0d 00:08:20
DHCP Server Addr    : 192.168.100.1

Relay Agent Information
  Circuit Id          : ds-left
  Remote Id           : remote0
  Radius User-Name    : "ds-left"
-----
Number of lease states : 1
=====
```

DHCPv6 Results

The DHCPv6 server receives the DHCPv6 Request message with Python modified lease times (preferred, valid, renew and rebind) for IA-NA and IA-PD.

```
32877 2014/07/24 03:15:28.32 WEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server: v6
Rx DHCPv6 RELAY_FORW
  Hop Count : 1
  Link Addr : 2001:db8:4::
  Peer Addr : fe80::200:65ff:fe01:301
  Option : RELAY_MSG (9), Length : 184
    Msg Type : RELAY_FORW (12)
    Hop Count : 0
    Link Addr : ::
    Peer Addr : fe80::200:65ff:fe01:301
    Option : INTERFACE_ID (18), Length : 7
      Interface Id : 64732d6c656674 (ds-left)
    Option : RELAY_MSG (9), Length : 135
      Msg Type : REQUEST (3)
```

```

Trans Id : 0x060000
Option : ELAPSED_TIME (8), Length : 2
    Time : 0 seconds
Option : CLIENTID (1), Length : 10
    LL : HwTyp=0001,LL=000065010301
    00030001000065010301
Option : SERVERID (2), Length : 10
    LL : HwTyp=0001,LL=d896ff000000
    00030001d896ff000000
Option : ORO (6), Length : 4
    Requested Option : IA_NA (3)
    Requested Option : IA_PD (25)
Option : IA_NA (3), Length : 40
    IAID : 0
    Time1: 2000 seconds
    Time2: 3000 seconds
Option : IAADDR (5), Length : 24
    Address : 2001:db8:3:1::1
    Preferred Lifetime : 4000 seconds
    Valid Lifetime : 4000 seconds
Option : IA_PD (25), Length : 41
    IAID : 0
    Time1: 2000 seconds
    Time2: 3000 seconds
Option : IAPREFIX (26), Length : 25
    Prefix : 2001:db8:4:400::/54
    Preferred Lifetime : 4000 seconds
    Valid Lifetime : 4000 seconds
Option : VENDOR_OPTS (17), Length : 37
    Enterprise : 0000197f
Option : WAN_POOL (1), Length : 8
    wan-left
Option : PFX_POOL (2), Length : 7
    pd-left
Option : PFX_LEN (3), Length : 1

```

The **hinted** DHCPv6 lease-times are not honored by the SR OS node DHCPv6 server and instead the SR OS node DHCPv6 server default values are inserted in the outgoing DHCPv6 Reply message towards the client as shown in the output below.

The explicitly configured DHCPv6 options are inserted by the DHCPv6 server first ([Table 36: CLI Inserted DHCP Options](#)) followed by the RADIUS supplied options inserted by the DHCPv6 relay ([Table 34: RADIUS inserted raw options](#)).

There are two DNS options [23] since they are supplied via two sources (DHCPv6 server and RADIUS Alc-ToClient-DHCP-Option VSA).

```

32885 2014/07/24 03:15:28.32 WEST MINOR: DEBUG #2001 vprn1 TIP
"TIP: DHCP6_PKT
  Outgoing DHCP6 Msg : RELAY_REPLY (13)
  to itf gl-1
  Hop Count : 0
  Link Addr : ::
  Peer Addr : fe80::200:65ff:fe01:301
  Option : RELAY_MSG (9), Length : 265
    Msg Type : REPLY (7)
    Trans Id : 0x060000
    Option : SERVERID (2), Length : 10
      LL : HwTyp=0001,LL=d896ff000000
      00030001d896ff000000
    Option : CLIENTID (1), Length : 10
      LL : HwTyp=0001,LL=000065010301
      00030001000065010301

```



```
Option : IA_NA (3), Length : 40
  IAID : 0
  Time1: 600 seconds
  Time2: 900 seconds
Option : IAADDR (5), Length : 24
  Address : 2001:db8:3:1::1
  Preferred Lifetime : 1200 seconds
  Valid Lifetime : 1200 seconds
Option : IA_PD (25), Length : 41
  IAID : 0
  Time1: 600 seconds
  Time2: 900 seconds
Option : IAPREFIX (26), Length : 25
  Prefix : 2001:db8:4:400::/54
  Preferred Lifetime : 1200 seconds
  Valid Lifetime : 1200 seconds
Option : DNS_NAME_SRVR (23), Length : 16
  Server : 2001:db8:1:1:1:1:1:3
Option : UNKNOWN (232), Length : 20
  763620637573746f6d206f7074696f6e20323332
Option : VENDOR_OPTS (17), Length : 26
  Enterprise : 0000197f
Option : UNKNOWN (230), Length : 18
  637573746f6d2074657374206f7074696f6e
Option : DNS_NAME_SRVR (23), Length : 32
  Server : 2001:db8:1:1:1:1:1:1
  Server : 2001:db8:1:1:1:1:1:2
Option : DOM_SRCH_LIST (24), Length : 30
  SearchList : .Nokia.com..test.com.
Option : INTERFACE_ID (18), Length : 7
  Interface Id : 64732d6c656674 (ds-left)
```

The Wireshark® capture of the DHCPv6 Reply message on the client side mirrors the debug information captured by the SR OS node:

```

❏ DHCPv6
  Message type: Relay-reply (13)
  Hopcount: 0
  Link address: :: (::)
  Peer address: fe80::200:65ff:fe01:301 (fe80::200:65ff:fe01:301)
❏ Relay Message
  Option: Relay Message (9)
  Length: 265
  Value: 070600000002000a00030001d896ff0000000001000a0003...
❏ DHCPv6
  Message type: Reply (7)
  Transaction ID: 0x060000
❏ Server Identifier
  Option: Server Identifier (2)
  Length: 10
  Value: 00030001d896ff000000
  DUID: 00030001d896ff000000
  DUID Type: link-layer address (3)
  Hardware type: Ethernet (1)
  Link-layer address: d8:96:ff:00:00:00
❏ Client Identifier
  Option: Client Identifier (1)
  Length: 10
  Value: 00030001000065010301
  DUID: 00030001000065010301
  DUID Type: link-layer address (3)
  Hardware type: Ethernet (1)
  Link-layer address: 00:00:65:01:03:01
❏ Identity Association for Non-temporary Address
  Option: Identity Association for Non-temporary Address (3)
  Length: 40
  Value: 00000000000000258000003840005001820010db800030001...
  IAID: 00000000
  T1: 600
  T2: 900
❏ IA Address
  Option: IA Address (5)
  Length: 24
  Value: 20010db8000300010000000000000001000004b00000004b0
  IPv6 address: 2001:db8:3:1::1 (2001:db8:3:1::1)
  Preferred lifetime: 1200
  valid lifetime: 1200
```

```

    Identity Association for Prefix Delegation
      Option: Identity Association for Prefix Delegation (25)
      Length: 41
      Value: 000000000000025800000384001a0019000004b0000004b0...
      IAID: 00000000
      T1: 600
      T2: 900
    IA Prefix
      Option: IA Prefix (26)
      Length: 25
      Value: 000004b0000004b03620010db8000404000000000000000...
      Preferred lifetime: 1200
      Valid lifetime: 1200
      Prefix length: 54
      Prefix address: 2001:db8:4:400:: (2001:db8:4:400::)
    DNS recursive name server
      Option: DNS recursive name server (23)
      Length: 16
      Value: 20010db8000100010001000100010003
      DNS server address: 2001:db8:1:1:1:1:1:3 (2001:db8:1:1:1:1:1:3)
    DHCP option 232
      Option: Unknown (232)
      Length: 20
      Value: 763620637573746f6d206f7074696f6e20323332
    Vendor-specific Information
      Option: Vendor-specific Information (17)
      Length: 26
      Value: 0000197f00e60012637573746f6d2074657374206f707469...
      Enterprise ID: Panthera Networks, Inc. (6527)
    option
    DNS recursive name server
      Option: DNS recursive name server (23)
      Length: 32
      Value: 20010db800010001000100010001000120010db800010001...
      DNS server address: 2001:db8:1:1:1:1:1:1 (2001:db8:1:1:1:1:1:1)
      DNS server address: 2001:db8:1:1:1:1:1:2 (2001:db8:1:1:1:1:1:2)
    Domain Search List
      Option: Domain Search List (24)
      Length: 30
      Value: 0e616c636174656c2d6c7563656e7403636f6d0004746573...
      DNS Domain Search List
      Domain: alcatel-lucent.com
      Domain: test.com
    Interface-Id
      Option: Interface-Id (18)
      Length: 7
      Value: 64732d6c656674
      Interface-ID: ds-left
  
```

The following command captures the information kept in the SR OS node DHCPv6 relay lease state:

```
*A:BNG1# show service id 1 dhcp6 lease-state detail
```

```
=====
DHCP lease states for service 1
=====
```

```

Service ID      : 1
IP Address      : 2001:db8:3:1::1/128
Client HW Address : 00:00:65:01:03:01
Subscriber-interface : int1-1
Group-interface  : gl-1
SAP             : [1/1/5:1.3]
Up Time         : 0d 00:02:41
Remaining Lease Time : 0d 00:17:18
Remaining SessionTime: N/A
  
```

```
Persistence Key      : N/A

Sub-Ident            : "ds-left"
Sub-Profile-String   : "sub-profile-1"
SLA-Profile-String   : "sla-profile-2"
App-Profile-String   : ""
Lease ANCP-String    : ""
Lease Int Dest Id    : ""
Category-Map-Name    : ""
Dhcp6 ClientId (DUID): 00030001000065010301
Dhcp6 IAID           : 0
Dhcp6 IAID Type      : non-temporary
Dhcp6 Client Ip      : fe80::200:65ff:fe01:301
Primary-Dns          : N/A
Secondary-Dns        : N/A
Pool Name            : "wan-left"
Dhcp6 Server Addr    : 2001:db8::1001
Dhcp6 ServerId (DUID): 00030001d896ff000000
Dhcp6 InterfaceId    : ds-left
Dhcp6 RemoteId       : N/A

Lease Info origin    : DHCP

ServerLeaseStart     : 07/24/2014 03:15:28
ServerLastRenew      : 07/24/2014 03:15:28
ServerLeaseEnd       : 07/24/2014 03:35:27
Session-Timeout      : N/A
Radius User-Name     : "ds-left"
-----
Service ID           : 1
IP Address           : 2001:db8:4:400::/54
Client HW Address    : 00:00:65:01:03:01
Subscriber-interface : int1-1
Group-interface      : gl-1
SAP                  : [1/1/5:1.3]
Up Time              : 0d 00:02:41
Remaining Lease Time : 0d 00:17:18
Remaining SessionTime: N/A
Persistence Key      : N/A

Sub-Ident            : "ds-left"
Sub-Profile-String   : "sub-profile-1"
SLA-Profile-String   : "sla-profile-2"
App-Profile-String   : ""
Lease ANCP-String    : ""
Lease Int Dest Id    : ""
Category-Map-Name    : ""
Dhcp6 ClientId (DUID): 00030001000065010301
Dhcp6 IAID           : 0
Dhcp6 IAID Type      : prefix
Dhcp6 Client Ip      : fe80::200:65ff:fe01:301
Primary-Dns          : N/A
Secondary-Dns        : N/A
Pool Name            : "pd-left"
Dhcp6 Server Addr    : 2001:db8::1001
Dhcp6 ServerId (DUID): 00030001d896ff000000
Dhcp6 InterfaceId    : ds-left
Dhcp6 RemoteId       : N/A

Lease Info origin    : DHCP

ServerLeaseStart     : 07/24/2014 03:15:28
ServerLastRenew      : 07/24/2014 03:15:28
ServerLeaseEnd       : 07/24/2014 03:35:27
```

```
Session-Timeout      : N/A
Radius User-Name     : "ds-left"
-----
Number of lease states : 2
=====
```

Python Debug Output

DHCPv4

For debugging purpose a line is added to the Python script printing all DHCP option numbers present in the incoming DHCP packets.

It can also be observed that all Python induced modifications to the original DHCP message are also displayed in the debugging output (inserting option [51] in this case).

Python script:

```
from alc import dhcpv4
myopt = dhcpv4.getOptionList()
print "option-list =", repr(myopt)
#lease 500s 8min 20sec
dhcpv4.set(51, ('\\x00\\x00\\x01\\xf4',))
```

Debug Output:

```
32826 2014/07/24 03:02:46.44 WEST MINOR: DEBUG #2001 Base Python Output
"Python Output: acg
option-list (53, 54, 50, 82, 255)
"
32827 2014/07/24 03:02:46.44 WEST MINOR: DEBUG #2001 Base Python Result
"Python Result: acg
DHCPv4 Option 51, SET
      '\\x00\\x00\\x01\\xf4'
"
```

DHCPv6

Also the DHCPv6 Python script has some lines added to demonstrate Python debugging capabilities. The new lines print assigned values to the debugging output.

DHCPv6 script

```
from alc import dhcpv6
import struct
packet = dhcpv6.get_relaymsg()
msgTop = ord(dhcpv6.msg_type)
msgBot = ord(packet.msg_type)
ia_na = packet.get_iana()
ia_pd = packet.get_iapd()
print 'ia-na = ', ia_na
print '\\n'
print 'ia-pd = ', ia_pd
print '\\n'
print 'msg type Top = ', msgTop
```

```
print 'msg type Bot = ', msgBot

msgType = struct.unpack('B',packet.msg_type)[0]
print "relay packet: ", msgType

# in relay request insert DHCPv6 lease times
if msgBot == 3:

    ia_na[0][1] = '\x00\x00\x07\xd0'
    ia_na[0][2] = '\x00\x00\x0b\xb8'
    ia_na[0][3][5][0][1] = '\x00\x00\x0f\xa0'
    ia_na[0][3][5][0][2] = '\x00\x00\x0f\xa0'
    packet.set_iana(ia_na)
    ia_pd[0][1] = '\x00\x00\x07\xd0'
    ia_pd[0][2] = '\x00\x00\x0b\xb8'
    ia_pd[0][3][26][0][0] = '\x00\x00\x0f\xa0'
    ia_pd[0][3][26][0][1] = '\x00\x00\x0f\xa0'
    packet.set_iapd(ia_pd)
    dhcpv6.set_relaymsg(packet)
```

Python debugging output

```
32873 2014/07/24 03:15:28.32 WEST MINOR: DEBUG #2001 Base Python Output
"Python Output: acg6

ia-na = [['\x00\x00\x00\x00', '\x00\x00\x02X', '\x00\x00\x03\x84', {5: [['\x01
\r\xb8\x00\x03\x00\x01\x00\x00\x00\x00\x00\x00\x00\x01', '\x00\x00\x04\xb0', '\x
00\x00\x04\xb0', {}]]}]]

ia-pd = [['\x00\x00\x00\x00', '\x00\x00\x02X', '\x00\x00\x03\x84', {26: [['\x00
\x00\x04\xb0', '\x00\x00\x04\xb0', '6', '\x01\r\xb8\x00\x04\x04\x00\x00\x00\x00
\x00\x00\x00\x00\x00', {}]]}]]

msg type Top = 12
msg type Bot = 3
relay packet: 3
"

32874 2014/07/24 03:15:28.32 WEST MINOR: DEBUG #2001 Base Python Result
"Python Result: acg6
DHCPv6 Option 9, SET
'\x03\x06\x00\x00\x00\x08\x00\x02\x00\x00\x00\x01\x00\n\x00\x03\x00\x01\x00\x00
e\x01\x03\x01\x00\x02\x00\n\x00\x03\x00\x01\xd8\x96\xff\x00\x00\x00\x00\x06\x00\
x04\x00\x03\x00\x19\x00\x03\x00(\x00\x00\x00\x00\x00\x00\x07\xd0\x00\x0b\xb8
\x00\x05\x00\x18 \x01\r\xb8\x00\x03\x00\x01\x00\x00\x00\x00\x00\x00\x01\x00\
x00\x0f\xa0\x00\x00\x0f\xa0\x00\x19\x00)\x00\x00\x00\x00\x00\x00\x07\xd0\x00\x00
\x0b\xb8\x00\x1a\x00\x19\x00\x00\x0f\xa0\x00\x00\x0f\xa06 \x01\r\xb8\x00\x04\x04
\x00\x00\x00\x00\x00\x00\x00\x00'
"
```

Conclusion

The most common DHCP options that need to be passed by the SR OS node to the clients can be directly configured in CLI with a DHCP option specific command (such as DNS or a router option in IPv4). The DHCP option specific commands hide the complexity of the option encoding from the operator.

Less common options can be configured via a custom-option command in CLI. This scenario requires the operator to be familiar with the encoding of the option.

Similarly, RADIUS provides the means to pass the DHCP options destined to the client in the form of option specific RADIUS attributes (lease-time, etc). For less common options, two RADIUS attributes are provided: **Alc-ToClient-Dhcp-Options** and **Alc-ToClient-Dhcp6-Options**. These two attributes allow the operator to encode client destined DHCP options using hexadecimal notation. Although this process requires manual encoding it provides a very flexible way of providing options to the client.

The custom options supplied via LUDB or RADIUS are appended by the SR OS node DHCP-relay agent to any existing options that may have been already inserted by the DHCP server in the DHCP packet.

Python processing can additionally assist in DHCP message processing where the options or the parameters within the existing options can be added, removed or modified.

Routed CO

This chapter provides information about Routed Central Office (Routed CO) configurations.

Topics in this chapter include:

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

Applicability

The information and configuration in this chapter are based on SR OS 11.0.R4.

Overview

In the Routed Central Office (Routed CO) model, subscriber management features are implemented on a Layer 3 subscriber interface, available in a VPRN or an IES service. Compared to regular Layer 3 interfaces, a subscriber interface supports multiple SAPs, see later.

Customer originated traffic enters an Access Node (AN) and can be aggregated via either a Layer 2 or a Layer 3 aggregation network before being handled by a Broadband Network Gateway (BNG). Alternatively, an AN can be directly connected to the BNG.

Routed CO supports numbered, unnumbered, and hybrid (combined numbered/unnumbered) subscriber interface configurations.

Enhanced Subscriber Management (ESM) is not mandatory for IPoEv4 in Routed CO, but is mandatory for PPPoE and all IPoEv6 scenarios.

The numbered and unnumbered scenarios in this example use an IES service with:

- Dual Stack IPoEv4 + IPoEv6
- Single stack PPPoEv4

General knowledge of Triple Play Service Delivery Architecture is assumed throughout this chapter. Refer to the *7450 ESS*, *7750 SR*, and *VSR Triple Play Service Delivery Architecture Guide*.

The Routed CO model offers through the subscriber and group interface construct:

- Flexible subnet management
 - Subnets can be shared across multiple access nodes.
- Support for different deployment models, for example:
 - VLAN/service model.
 - VLAN/subscriber model.

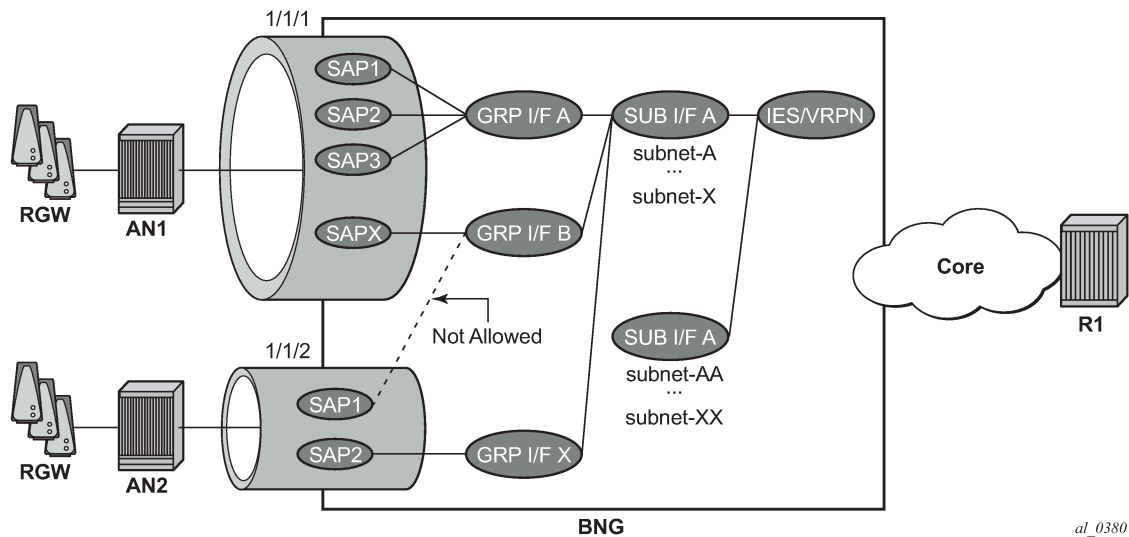
- VLAN/service/subscriber model.
- VLAN/access node model.
- Per group-interface load balancing in multi-chassis redundancy configurations. Redundancy is out of the scope of this example.

The components needed in the Routed CO model are depicted in [Figure 155: Components of the Routed CO Model](#).

For the Routed CO model two interface types are needed:

- First, one or more subscriber interfaces must be created.
- Second, one or more group interfaces must be created within the subscriber interface context.

Figure 155: Components of the Routed CO Model



Subscriber Interface

A subscriber interface is a set of one or more group interfaces and identified by name.

A subscriber interface is created under an IES or VPRN service context, and supports up to 256 subnets (sum of IPv4 subnets and IPv6 prefixes).

Three types of subscriber interface configurations are available:

- Numbered subscriber interface.
- Unnumbered subscriber interface.
- Hybrid subscriber interface (numbered and unnumbered combined).

Subnet/Prefix Assignment

For the numbered scenario, the subscriber interface is configured with

- One or more IPv4 subnets.

- One or more IPv6 subscriber prefixes:
 - For WAN-hosts, using the DHCPv6 Identity Association for Non-temporary Addresses (IA_NA) option or Stateless Address Auto Configuration (SLAAC) and the prefix length is /64.
 - For Prefix Delegation-hosts (PD-hosts), using the DHCPv6 Identity Association for Prefix Delegation (IA_PD) option and the prefix length is defined by the Delegated Prefix Length (DPL).

This allows for subscriber-host address assignment in these subnets/prefixes only.

For the unnumbered scenario, the subscriber interface is configured with:

- IPv4:
 - No IPv4 subnets.
 - The keyword **unnumbered** plus an interface in the same routing instance (for example the system interface). The IP address of the interface referenced in the unnumbered command is used in IPCP negotiation.
- IPv6:
 - No IPv6 prefixes.
 - **allow-unmatching-prefixes**.

This allows for subscriber-host address assignment in any subnet/prefix. For IPv4, the keywords **unnumbered** and **allow-unmatching-subnets** are mutually exclusive.

For the hybrid scenario the subscriber interface is configured with:

- One or more IPv4 subnets and/or IPv6 subscriber prefixes.
- For IPv4: the keyword **allow-unmatching-subnets**.
- For IPv6: the keyword **allow-unmatching-prefixes**.

This allows for both subscriber-host address assignment within and outside of these subnets/prefixes.

Host IP Reachability

For the numbered scenario, host IP reachability requires:

- Adding the subscriber interfaces to the Interior Gateway Protocol (IGP).
- Or an export policy matching the subscriber interface subnets/prefixes.

For the unnumbered scenario, host IP reachability requires:

- An export policy matching the addresses of all individual subscriber hosts (from protocol sub-mgmt).

For the hybrid scenario, host IP reachability requires:

- An export policy matching both the subscriber interface subnets/prefixes as well as all individual subscriber hosts addresses.

Detailed examples of numbered/unnumbered/hybrid scenarios, including host IP reachability are included below.

Group Interface

A group interface is a set of one or more SAPs belonging to the same port and identified by name.

Configuration

This section covers:

- The definition of subscriber and group interfaces.
- A description of the numbered, unnumbered and hybrid scenarios.
- Options ensuring host IP reachability throughout the network.

Subscriber Interface

The configuration of the subscriber interface appears as follows.

```
configure
service
  ies 1
    subscriber-interface "sub-int-1" create
      address 10.1.1.254/24
      address 10.1.2.254/24
      ipv6
        delegated-prefix-len 56
        link-local-address FE80::EA:48:FF
        subscriber-prefixes
          prefix 2001:DB8:101::/48 wan-host
          prefix 2001:DB8:102::/48 pd
          prefix 2001:DB8:F101::/48 wan-host
          prefix 2001:DB8:F102::/48 pd
        exit
      exit
    exit
  subscriber-interface "sub-int-2" create
    address 10.2.1.254/24
    address 10.2.2.254/24
    ipv6
      delegated-prefix-len 56
      link-local-address FE80::EA:48:FF
      subscriber-prefixes
        prefix 2001:DB8:201::/48 wan-host
        prefix 2001:DB8:202::/48 pd
        prefix 2001:DB8:F201::/48 wan-host
        prefix 2001:DB8:F202::/48 pd
      exit
    exit
  exit
```

Notice that once a subnet/prefix is assigned to a subscriber interface, the subnet/prefix is tied to that interface, meaning that the same subnet/prefix cannot be used on another subscriber interface or regular interface in the same routing instance. When using VPRN for the Routed CO model, overlapping subnets/prefixes are allowed when on different VPRN services.

As long as no group interfaces are configured within the subscriber interface context, the subscriber interfaces are in the operationally down state as shown in the following output. The subscriber-interfaces, sub-int-1 and sub-int-2, are operational down since no group-interfaces have been assigned at this stage.

```
*A:BNG# show router "Base" interface
=====
Interface Table (Router: Base)
```

Interface-Name IP-Address	Adm	Opr(v4/v6)	Mode	Port/SapId PfxState
sub-int-1	Up	Down/Down	IES Sub	subscriber
10.1.1.254/24				n/a
10.1.2.254/24				n/a
2001:DB8:101::/48				INACCESSIBLE
2001:DB8:102::/48				INACCESSIBLE
2001:DB8:F101::/48				INACCESSIBLE
2001:DB8:F102::/48				INACCESSIBLE
FE80::EA:48:FF/64				INACCESSIBLE
sub-int-2	Up	Down/Down	IES Sub	subscriber
10.2.1.254/24				n/a
10.2.2.254/24				n/a
2001:DB8:201::/48				INACCESSIBLE
2001:DB8:202::/48				INACCESSIBLE
2001:DB8:F201::/48				INACCESSIBLE
2001:DB8:F202::/48				INACCESSIBLE
FE80::EA:48:FF/64				INACCESSIBLE
system	Up	Up/Up	Network	system
192.0.2.75/32				n/a
2001:DB8::75/128				PREFERRED
toDHCP-1	Up	Up/Up	Network	loopback
10.11.11.1/32				n/a
2001:DB8::11/128				PREFERRED
FE80::E84B:FFFF:FE00:0/64				PREFERRED
toRADIUS-1	Up	Up/Down	Network	1/1/10
192.168.202.75/24				n/a

Interfaces : 5				
=====				
*A:BNG#				

The corresponding IPv4 routing table looks as follows.

*A:BNG# show router "Base" route-table ipv4					
=====					
Route Table (Router: Base)					
=====					
Dest Prefix[Flags]		Type	Proto	Age	Pref
Next Hop[Interface Name]				Metric	

10.11.11.1/32		Local	Local	00h30m12s	0
toDHCP-1				0	
192.0.2.75/32		Local	Local	00h30m12s	0
system				0	
192.168.202.0/24		Local	Local	00h29m54s	0
toRADIUS-1				0	

No. of Routes: 3					
Flags: L = LFA nexthop available B = BGP backup route available					
n = Number of times nexthop is repeated					
=====					
*A:BNG#					

The corresponding IPv6 routing table looks as follows.

*A:BNG# show router "Base" route-table ipv6				
=====				
IPv6 Route Table (Router: Base)				
=====				
Dest Prefix[Flags]	Type	Proto	Age	Pref

Next Hop[Interface Name]	Metric			
2001:DB8::11/128	Local	Local	00h30m19s	0
toDHCP-1			0	
2001:DB8::75/128	Local	Local	00h30m21s	0
system			0	

No. of Routes: 2				
Flags: L = LFA nexthop available B = BGP backup route available				
n = Number of times nexthop is repeated				
=====				
*A:BNG#				

No subscriber interface subnets/prefixes are present in the IPv4 and the IPv6 routing table as the subscriber interfaces are operational down.

Group Interface

A group interface is created under the subscriber-interface hierarchy.

```
configure
service
  ies 1
    subscriber-interface "sub-int-1" create
      group-interface "grp-int-1-1" create
        ipv6
        exit
        sap 1/1/1:111 create
        exit
        sap 1/1/1:112 create
        exit
      exit
      group-interface "grp-int-1-2" create
        ipv6
        exit
        sap 1/1/1:121 create
        exit
      exit
    exit
  subscriber-interface "sub-int-2" create
    group-interface "grp-int-2-1" create
      ipv6
      exit
      sap 1/1/2:211 create
      exit
    exit
    group-interface "grp-int-2-2" create
      ipv6
      exit
      sap 1/1/3:221 create
      exit
      sap 1/1/3:222 create
      exit
    exit
  exit
exit
```

Static SAPs are created manually under the group-interface context. Managed SAPs (MSAPs) are dynamically created when a trigger packet (DHCP, DHCPv6, ARP, PPPoE) is successfully authenticated,

which eliminates the provisioning of static SAPs. The creation and use of capture and managed SAPs (MSAPs) is explained in the example on [Managed SAPs with Routed CO](#).

A group interface is operationally up when at least one of its statically configured SAPs is operationally up or when no static SAPs are configured while the parameter **oper-up-while-empty** under the group-interface context is enabled. The following output shows all group interfaces are operationally up.

```
*A:BNB# show router "Base" interface ipv4
=====
Interface Table (Router: Base)
=====
Interface-Name      Adm      Opr(v4/v6)  Mode      Port/SapId
IP-Address          PfxState
-----
grp-int-1-1         Up       Up/Up       IES Grp   1/1/1
grp-int-1-2         Up       Up/Up       IES Grp   1/1/1
grp-int-2-1         Up       Up/Up       IES Grp   1/1/2
grp-int-2-2         Up       Up/Up       IES Grp   1/1/3
sub-int-1           Up       Up/Up       IES Sub   subscriber
10.1.1.254/24       n/a
10.1.2.254/24       n/a
sub-int-2           Up       Up/Up       IES Sub   subscriber
10.2.1.254/24       n/a
10.2.2.254/24       n/a
system              Up       Up/Up       Network  system
192.0.2.75/32       n/a
toDHCP-1            Up       Up/Up       Network  loopback
10.11.11.1/32       n/a
toRADIUS-1          Up       Up/Down     Network  1/1/10
192.168.202.75/24   n/a
-----
Interfaces : 9
=====
*A:BNB#
```

The IPv4 routing table includes the subnets configured on the subscriber-interfaces.

```
*A:BNB# show router "Base" route-table ipv4
=====
Route Table (Router: Base)
=====
Dest Prefix[Flags]      Type  Proto  Age      Pref
Next Hop[Interface Name] Metric
-----
10.1.1.0/24             Local  Local  00h25m32s 0
sub-int-1               0
10.1.2.0/24             Local  Local  00h25m32s 0
sub-int-1               0
10.2.1.0/24             Local  Local  00h25m32s 0
sub-int-2               0
10.2.2.0/24             Local  Local  00h25m32s 0
sub-int-2               0
10.11.11.1/32           Local  Local  01h01m53s 0
toDHCP-1                0
192.0.2.75/32           Local  Local  01h01m53s 0
system                  0
192.168.202.0/24        Local  Local  01h01m36s 0
toRADIUS-1              0
-----
No. of Routes: 7
Flags: L = LFA nexthop available    B = BGP backup route available
n = Number of times nexthop is repeated
```

```
=====
*A:BN#
```

For IPv6, the interface table looks as follows.

```
*A:BN# show router "Base" interface ipv6
```

```
=====
Interface Table (Router: Base)
=====
```

Interface-Name IP-Address	Adm	Opr(v4/v6)	Mode	Port/SapId PfxState
grp-int-1-1	Up	Up/Up	IES Grp	1/1/1
grp-int-1-2	Up	Up/Up	IES Grp	1/1/1
grp-int-2-1	Up	Up/Up	IES Grp	1/1/2
grp-int-2-2	Up	Up/Up	IES Grp	1/1/3
sub-int-1	Up	Up/Up	IES Sub	subscriber
2001:DB8:101::/48				PREFERRED
2001:DB8:102::/48				PREFERRED
2001:DB8:F101::/48				PREFERRED
2001:DB8:F102::/48				PREFERRED
FE80::EA:48:FF/64				PREFERRED
sub-int-2	Up	Up/Up	IES Sub	subscriber
2001:DB8:201::/48				PREFERRED
2001:DB8:202::/48				PREFERRED
2001:DB8:F201::/48				PREFERRED
2001:DB8:F202::/48				PREFERRED
FE80::EA:48:FF/64				PREFERRED
system	Up	Up/Up	Network	system
2001:DB8::75/128				PREFERRED
toDHCP-1	Up	Up/Up	Network	loopback
2001:DB8::11/128				PREFERRED
FE80::E84B:FFFF:FE00:0/64				PREFERRED
toRADIUS-1	Up	Up/Down	Network	1/1/10
-				-

```
-----
Interfaces : 9
=====
```

```
*A:BN# #
```

The IPv6 routing table includes the prefixes configured on the subscriber interfaces.

```
*A:BN# show router "Base" route-table ipv6
```

```
=====
IPv6 Route Table (Router: Base)
=====
```

Dest Prefix[Flags] Next Hop[Interface Name]	Type	Proto	Age Metric	Pref
2001:DB8::11/128	Local	Local	14d04h08m 0	0
toDHCP-1				
2001:DB8::75/128	Local	Local	14d04h08m 0	0
system				
2001:DB8:101::/48	Local	Local	14d03h10m 0	0
sub-int-1				
2001:DB8:102::/48	Local	Local	14d03h10m 0	0
sub-int-1				
2001:DB8:201::/48	Local	Local	14d03h08m 0	0
sub-int-2				
2001:DB8:202::/48	Local	Local	14d03h08m 0	0
sub-int-2				
2001:DB8:F101::/48	Local	Local	14d03h10m 0	0
sub-int-1				

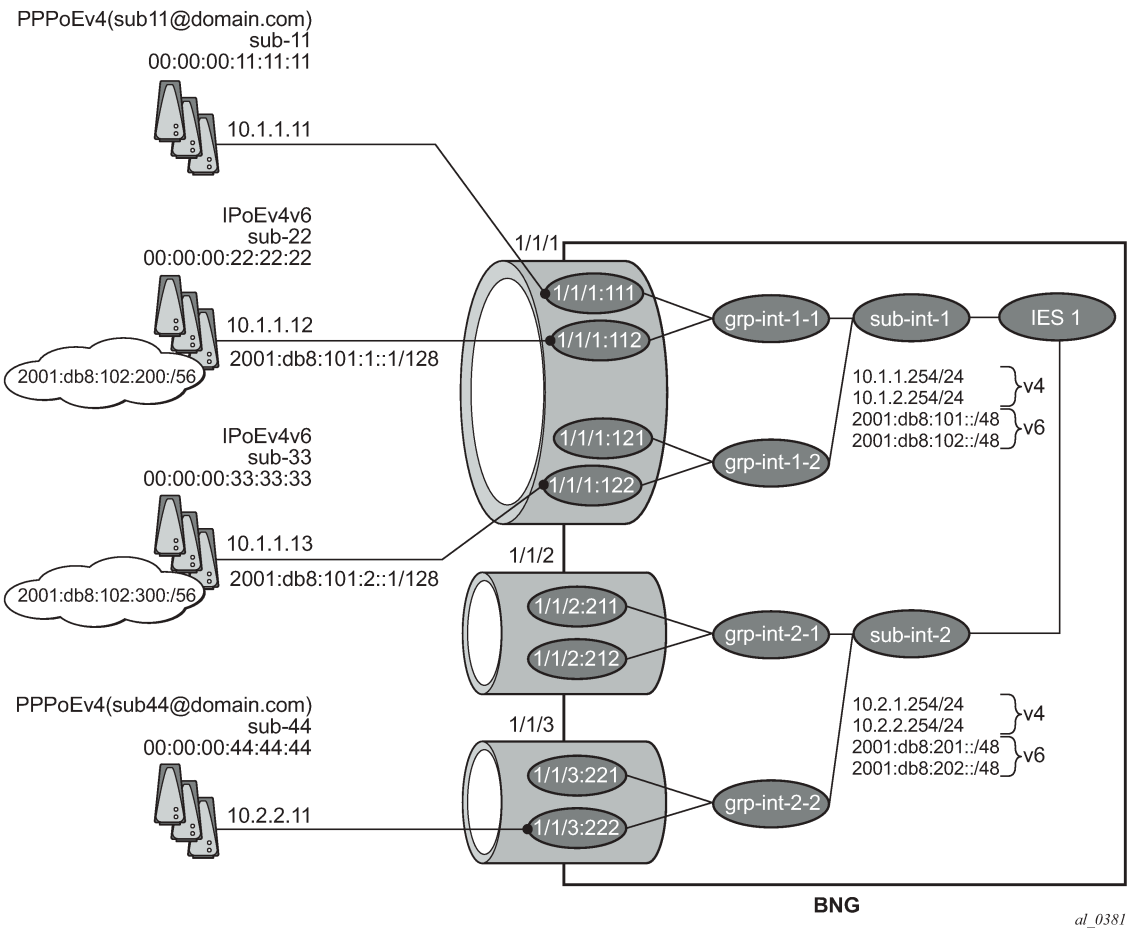
2001:DB8:F102::/48	Local	Local	14d03h10m	0
sub-int-1			0	
2001:DB8:F201::/48	Local	Local	14d03h08m	0
sub-int-2			0	
2001:DB8:F202::/48	Local	Local	14d03h08m	0
sub-int-2			0	

No. of Routes: 10				
Flags: L = LFA nexthop available B = BGP backup route available				
n = Number of times nexthop is repeated				
=====				
*A:BN#				

Numbered Scenario

Figure 156: Numbered Scenario For IES 1 depicts the numbered scenario outlined below, including the connecting subscribers and subscriber hosts. Subscribers sub-11 and sub-44 are using PPPv4 hosts, and subscribers sub-22 and sub-33 are using dual stack DHCP hosts. Their VLANs and the MAC addresses are shown, as are the IP addresses and prefixes assigned once they are connected.

Figure 156: Numbered Scenario For IES 1



The configuration for the numbered scenario is shown below. Only the configuration items specific to the numbered scenario are shown.

In the numbered scenario the subscriber interfaces have following configuration:

- IPv4
 - Subnets.
 - **no allow-unmatching-subnets.**
 - **no unnumbered.**
- IPv6
 - A delegated prefix length.
 - subscriber prefixes.
 - **no allow-unmatching-prefixes.**

```
configure
service
  ies 1
    subscriber-interface "sub-int-1" create
      address 10.1.1.254/24
      address 10.1.2.254/24
      ipv6
        delegated-prefix-len 56
        link-local-address FE80::EA:4B:FF
        subscriber-prefixes
          prefix 2001:DB8:101::/48 wan-host
          prefix 2001:DB8:102::/48 pd
        exit
      exit
    group-interface "grp-int-1-1" create
      ipv6
        ---snip---
      exit
      arp-populate
      dhcp
        ---snip---
        lease-populate 100
        no shutdown
      exit
      authentication-policy "auth-pol-1"
      local-proxy-arp
      sap 1/1/1:111 create
        anti-spoof ip-mac
        sub-sla-mgmt
        ---snip---
      exit
      exit
      sap 1/1/1:112 create
        anti-spoof ip-mac
        sub-sla-mgmt
        ---snip---
      exit
      exit
      pppoe
        ---snip---
        no shutdown
      exit
    exit
  group-interface "grp-int-1-2" create
    ipv6
```

```
        ---snip---
        exit
        arp-populate
        dhcp
        ---snip---
        lease-populate 100
        no shutdown
        exit
        authentication-policy "auth-pol-1"
        local-proxy-arp
        sap 1/1/1:121 create
        anti-spoof ip-mac
        sub-sla-mgmt
        ---snip---
        exit
        exit
        sap 1/1/1:122 create
        anti-spoof ip-mac
        sub-sla-mgmt
        ---snip---
        exit
        exit
        pppoe
        ---snip---
        no shutdown
        exit
        exit
    exit
subscriber-interface "sub-int-2" create
    address 10.2.1.254/24
    address 10.2.2.254/24
    ipv6
        delegated-prefix-len 56
        link-local-address FE80::EA:4B:FF
        subscriber-prefixes
            prefix 2001:DB8:201::/48 wan-host
            prefix 2001:DB8:202::/48 pd
        exit
    exit
group-interface "grp-int-2-1" create
    ipv6
        ---snip---
        exit
        arp-populate
        dhcp
        ---snip---
        lease-populate 100
        no shutdown
        exit
        authentication-policy "auth-pol-1"
        local-proxy-arp
        sap 1/1/2:211 create
        anti-spoof ip-mac
        sub-sla-mgmt
        ---snip---
        exit
        exit
        sap 1/1/2:212 create
        anti-spoof ip-mac
        sub-sla-mgmt
        ---snip---
        exit
        exit
        pppoe
```

```

        ---snip---
        no shutdown
    exit
exit
group-interface "grp-int-2-2" create
    ipv6
        ---snip---
    exit
    arp-populate
    dhcp
        ---snip---
        lease-populate 100
        no shutdown
    exit
    authentication-policy "auth-pol-1"
    local-proxy-arp
    sap 1/1/3:221 create
        anti-spoof ip-mac
        sub-sla-mgmt
        ---snip---
    exit
    exit
    sap 1/1/3:222 create
        anti-spoof ip-mac
        sub-sla-mgmt
        ---snip---
    exit
    exit
    pppoe
        ---snip---
        no shutdown
    exit
exit
exit
no shutdown

```

The following parameters are mandatory for the routed CO model:

- **lease-populate** — DHCPv4 lease state population is enabled by default on a group-interface with DHCPv4 configured as **no shutdown**. The number of leases allowed on each SAP of the group-interface must be configured. By default one single DHCPv4 host is allowed on each SAP. This parameter enables the creation of an ESM host table entry for each DHCPv4 lease. For DHCPv6 the ESM host table entry creation is implicit: no CLI parameter is required.
- **arp-populate** — The ARP table is populated with dynamically learned entries from the DHCP lease state table or static entries from the static host table. The BNG does not send downstream ARPs for those managed ARP table entries.
- **local-proxy-arp** — Enables user to user traffic in a split-horizon environment. The BNG responds with its own MAC address to ARP requests targeting subnets configured on the subscriber interface. If the ARP request is targeting a host of the same subscriber on the same SAP, the ARP request is silently discarded. This prevents traffic within a single bridged home to be attracted to the BNG. Local-proxy-arp is enabled by default.
- **anti-spoof** — Checks the source MAC and/or source IP of the upstream subscriber traffic. This parameter is configured at the SAP level with values **ip-mac** (default), **ip** or **nh-mac**. With ESM enabled, anti-spoof must include the source mac (values **ip-mac** or **nh-mac**).

Optional settings are:

- **description** — Can be used to assign a descriptive text to the item and used for administrative reasons.

- **delayed-enable** — To be used in redundant configurations. It is expressed in seconds and defines the additional time the BNG waits before the interface is enabled.

Verification

The interfaces on the BNG are listed using following command. Notice that all subscriber and group interfaces are operational up for IPv4 and IPv6.

```
A:BNG# show router "Base" interface
=====
Interface Table (Router: Base)
=====
Interface-Name      Adm      Opr(v4/v6)  Mode      Port/SapId  PfxState
IP-Address
-----
grp-int-1-1         Up        Up/Up        IES Grp    1/1/1
grp-int-1-2         Up        Up/Up        IES Grp    1/1/1
grp-int-2-1         Up        Up/Up        IES Grp    1/1/2
grp-int-2-2         Up        Up/Up        IES Grp    1/1/3
sub-int-1           Up        Up/Up        IES Sub    subscriber
10.1.1.254/24       n/a
10.1.2.254/24       n/a
2001:DB8:101::/48   PREFERRED
2001:DB8:102::/48   PREFERRED
FE80::EA:48:FF/64   PREFERRED
sub-int-2           Up        Up/Up        IES Sub    subscriber
10.2.1.254/24       n/a
10.2.2.254/24       n/a
2001:DB8:201::/48   PREFERRED
2001:DB8:202::/48   PREFERRED
FE80::EA:48:FF/64   PREFERRED
system              Up        Up/Up        Network    system
192.0.2.75/32       n/a
2001:DB8::75/128    PREFERRED
toDHCP-1            Up        Up/Up        Network    loopback
10.11.11.1/32       n/a
2001:DB8::11/128    PREFERRED
FE80::E84B:FFFF:FE00:0/64 PREFERRED
toR1                Up        Up/Up        Network    1/1/12
192.168.12.1/24     n/a
2001:DEAD::1/64     PREFERRED
FE80::E84B:FFFF:FE00:0/64 PREFERRED
toRADIUS-1          Up        Up/Down      Network    1/1/10
192.168.202.75/24   n/a
-----
Interfaces : 10
=====
A:BNG#
```

Successfully created hosts have forwarding state Fwding. Hosts not in the Fwding state cannot forward any data.

```
A:BNG# show service id 1 subscriber-hosts
=====
Subscriber Host table
=====
Sap      Subscriber
IP Address
MAC Address  PPPoE-SID Origin  Fwding State
-----
```

```
1/1/1:111          sub-11
 10.1.1.11
 00:00:00:11:11:11    1      IPCP      Fwding
1/1/1:112          sub-22
 10.1.1.12
 00:00:00:22:22:22    N/A     DHCP      Fwding
1/1/1:112          sub-22
2001:DB8:101:1::1/128
 00:00:00:22:22:22    N/A     IPoE-DHCP6  Fwding
1/1/1:112          sub-22
2001:DB8:102:200::/56
 00:00:00:22:22:22    N/A     IPoE-DHCP6  Fwding
1/1/1:122          sub-33
 10.1.1.13
 00:00:00:33:33:33    N/A     DHCP      Fwding
1/1/1:122          sub-33
2001:DB8:101:2::1/128
 00:00:00:33:33:33    N/A     IPoE-DHCP6  Fwding
1/1/1:122          sub-33
2001:DB8:102:300::/56
 00:00:00:33:33:33    N/A     IPoE-DHCP6  Fwding
1/1/3:222          sub-44
 10.2.2.11
 00:00:00:44:44:44    1      IPCP      Fwding
-----
Number of subscriber hosts : 8
=====
A:BNG#
```

The list of active subscribers can be displayed as follows.

```
A:BNG# show service active-subscribers
=====
Active Subscribers
=====
-----
Subscriber sub-11 (sub-prof-1)
-----
(1) SLA Profile Instance sap:1/1/1:111 - sla:sla-prof-1
-----
IP Address          MAC Address          PPPoE-SID Origin
-----
10.1.1.11           00:00:00:11:11:11    1      IPCP
-----
Subscriber sub-22 (sub-prof-1)
-----
(1) SLA Profile Instance sap:1/1/1:112 - sla:sla-prof-1
-----
IP Address          MAC Address          PPPoE-SID Origin
-----
10.1.1.12           00:00:00:22:22:22    N/A     DHCP
2001:DB8:101:1::1/128
                    00:00:00:22:22:22    N/A     IPoE-DHCP6
2001:DB8:102:200::/56
                    00:00:00:22:22:22    N/A     IPoE-DHCP6
-----
Subscriber sub-33 (sub-prof-1)
```

```
-----
(1) SLA Profile Instance sap:1/1/1:122 - sla:sla-prof-1
-----
```

IP Address	MAC Address	PPPoE-SID	Origin
10.1.1.13	00:00:00:33:33:33	N/A	DHCP
2001:DB8:101:2::1/128	00:00:00:33:33:33	N/A	IPoE-DHCP6
2001:DB8:102:300::/56	00:00:00:33:33:33	N/A	IPoE-DHCP6

```
-----
Subscriber sub-44 (sub-prof-1)
-----
```

```
-----
(1) SLA Profile Instance sap:1/1/3:222 - sla:sla-prof-1
-----
```

IP Address	MAC Address	PPPoE-SID	Origin
10.2.2.11	00:00:00:44:44:44	1	IPCP

```
-----
Number of active subscribers : 4
-----
```

```
A:BNG#
```

Manually cross-referencing the SAPs from this output with the actual configuration shows the following for IPv4, and is depicted in [Figure 156: Numbered Scenario For IES 1](#).

- Sub-11 and sub-22 are connected to the same subscriber and group interface (sub-int-1 and grp-int-1-1) but via different SAPs (1/1/1:111 and 1/1/1:112) and are sharing the same IPv4 subnet.
- Sub-33 is also connected to the same subscriber interface (sub-int-1) but via a different group-interface (grp-int-1-2). Sub-33 shares the same IPv4 subnet as sub-11 and sub-12, showing that the same subnet is shared across multiple group-interfaces.
- Sub-44 is connected to a different subscriber and group interface, and does not share a subnet with the other subscribers.

An alternative way to find where, for example, subscriber sub-33 is connected is shown below.

```
*A:BNG# show service active-subscribers subscriber "sub-33" detail
```

```
=====
Active Subscribers
=====
```

```
-----
Subscriber sub-11 (sub-prof-1)
-----
```

```
I. Sched. Policy : N/A
```

```
---snip---
```

```
Oper-Rate-Limit : Maximum
```

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

```
-----
(1) SLA Profile Instance
  - sap:1/1/1:112 (IES 1 - grp-int-1-2)
  - sla:sla-prof-1
-----
```

```
Description : (Not Specified)
```

---snip---

An alternative to find where, for example, IP address 10.1.1.13 is connected is shown below.

```
*A:BNG# show service id 1 dhcp lease-state ip-address 10.1.1.13 detail
=====
DHCP lease states for service 1
=====
Service ID      : 1
IP Address      : 10.1.1.13
Client HW Address : 00:00:00:33:33:33
Subscriber-interface : sub-int-1
Group-interface  : grp-int-1-2
SAP             : 1/1/1:122
---snip---

Sub-Ident       : "sub-33"
Sub-Profile-String : "sub-prof-1"
SLA-Profile-String : "sla-prof-1"
App-Profile-String : ""
---snip---

DHCP Server Addr : 10.11.11.1
Radius User-Name  : "00:00:00:33:33:33"
-----
Number of lease states : 1
=====
*A:BNG#
```

For IPv6, the situation is as follows:

- Sub-22 and sub-33 are connected to the same subscriber interface (sub-int-1) but to different group interfaces. Both subscribers share the same IPv6 prefix for prefix-delegation (PD) and wan-host.

With these subscriber hosts connected, the IPv4 routing table (RIB) for the base router looks as follows.

```
*A:BNG# show router "Base" route-table ipv4
=====
Route Table (Router: Base)
=====
```

Dest Prefix[Flags] Next Hop[Interface Name]	Type	Proto	Age Metric	Pref
10.1.1.0/24 sub-int-1	Local	Local	02h25m15s 0	0
10.1.1.11/32 [grp-int-1-1]	Remote	Sub Mgmt	02h25m10s 0	0
10.1.1.12/32 [grp-int-1-1]	Remote	Sub Mgmt	00h49m52s 0	0
10.1.1.13/32 [grp-int-1-2]	Remote	Sub Mgmt	00h47m40s 0	0
10.1.2.0/24 sub-int-1	Local	Local	02h25m15s 0	0
10.2.1.0/24 sub-int-2	Local	Local	02h25m15s 0	0
10.2.2.0/24 sub-int-2	Local	Local	02h25m15s 0	0
10.2.2.11/32 [grp-int-2-2]	Remote	Sub Mgmt	02h25m10s 0	0
10.11.11.1/32 toDHCP-1	Local	Local	02h25m33s 0	0

```

192.0.2.75/32          Local  Local  02h25m33s  0
    system
192.0.2.76/32          Remote ISIS    02h24m43s  15
    192.168.12.2
192.168.12.0/24        Local  Local  02h25m15s  0
    toR1
192.168.202.0/24       Local  Local  02h25m15s  0
    toRADIUS-1
-----
No. of Routes: 13
Flags: L = LFA nexthop available    B = BGP backup route available
      n = Number of times nexthop is repeated
=====
*A:BN#

```

The IPv6 routing table (RIB) for the base router displays as follows.

```

*A:BN# show router "Base" route-table ipv6
=====
IPv6 Route Table (Router: Base)
=====
Dest Prefix[Flags]      Type  Proto  Age      Pref
Next Hop[Interface Name] Metric
-----
2001:DB8::11/128        Local  Local  02h27m23s  0
    toDHCP-1
2001:DB8::75/128        Local  Local  02h27m24s  0
    system
2001:DB8::76/128        Remote ISIS    02h26m32s  15
    FE80::E84C:FFFF:FE00:0-"toR1"
2001:DB8:101::/48       Local  Local  02h27m06s  0
    sub-int-1
2001:DB8:101:1::1/128   Remote Sub Mgmt 01h13m43s  0
    [grp-int-1-1]
2001:DB8:101:2::1/128   Remote Sub Mgmt 01h13m25s  0
    [grp-int-1-2]
2001:DB8:102::/48       Local  Local  02h27m06s  0
    sub-int-1
2001:DB8:102:200::/56   Remote Sub Mgmt 01h13m43s  0
    [grp-int-1-1]
2001:DB8:102:300::/56   Remote Sub Mgmt 01h13m25s  0
    [grp-int-1-2]
2001:DB8:201::/48       Local  Local  02h27m06s  0
    sub-int-2
2001:DB8:202::/48       Local  Local  02h27m06s  0
    sub-int-2
2001:DEAD::/64          Local  Local  02h27m05s  0
    toR1
-----
No. of Routes: 12
Flags: L = LFA nexthop available    B = BGP backup route available
      n = Number of times nexthop is repeated
=====
*A:BN#

```

The corresponding IPv4 FIB on card 1 looks as follows.

```

*A:BN# show router "Base" fib 1 ipv4
=====
FIB Display
=====
Prefix                                Protocol
NextHop

```



```
-----
10.1.1.0/24                                LOCAL
   10.1.1.0 (sub-int-1)
10.1.2.0/24                                LOCAL
   10.1.2.0 (sub-int-1)
10.2.1.0/24                                LOCAL
   10.2.1.0 (sub-int-2)
10.2.2.0/24                                LOCAL
   10.2.2.0 (sub-int-2)
10.11.11.1/32                             LOCAL
   10.11.11.1 (toDHCP-1)
192.0.2.75/32                             LOCAL
   192.0.2.75 (system)
192.0.2.76/32                             ISIS
   192.168.12.2 (toR1)
192.168.12.0/24                             LOCAL
   192.168.12.0 (toR1)
192.168.202.0/24                            LOCAL
   192.168.202.0 (toRADIUS-1)
-----
Total Entries : 9
-----
=====
*A:BNG#
```

The corresponding IPv6 FIB on card 1 is as follows.

```
*A:BNG# show router "Base" fib 1 ipv6
=====
FIB Display
=====
Prefix                                     Protocol
  NextHop
-----
2001:DB8::11/128                          LOCAL
   2001:DB8::11 (toDHCP-1)
2001:DB8::75/128                          LOCAL
   2001:DB8::75 (system)
2001:DB8::76/128                          ISIS
   FE80::E84C:FFFF:FE00:0 (toR1)
2001:DB8:101::/48                         LOCAL
   2001:DB8:101:: (sub-int-1)
2001:DB8:102::/48                         LOCAL
   2001:DB8:102:: (sub-int-1)
2001:DB8:201::/48                         LOCAL
   2001:DB8:201:: (sub-int-2)
2001:DB8:202::/48                         LOCAL
   2001:DB8:202:: (sub-int-2)
2001:DEAD::/64                            LOCAL
   2001:DEAD:: (toR1)
-----
Total Entries : 8
-----
=====
*A:BNG#
```

The addresses of the individual subscriber hosts show up in the RIB but they do not show up in the FIB.

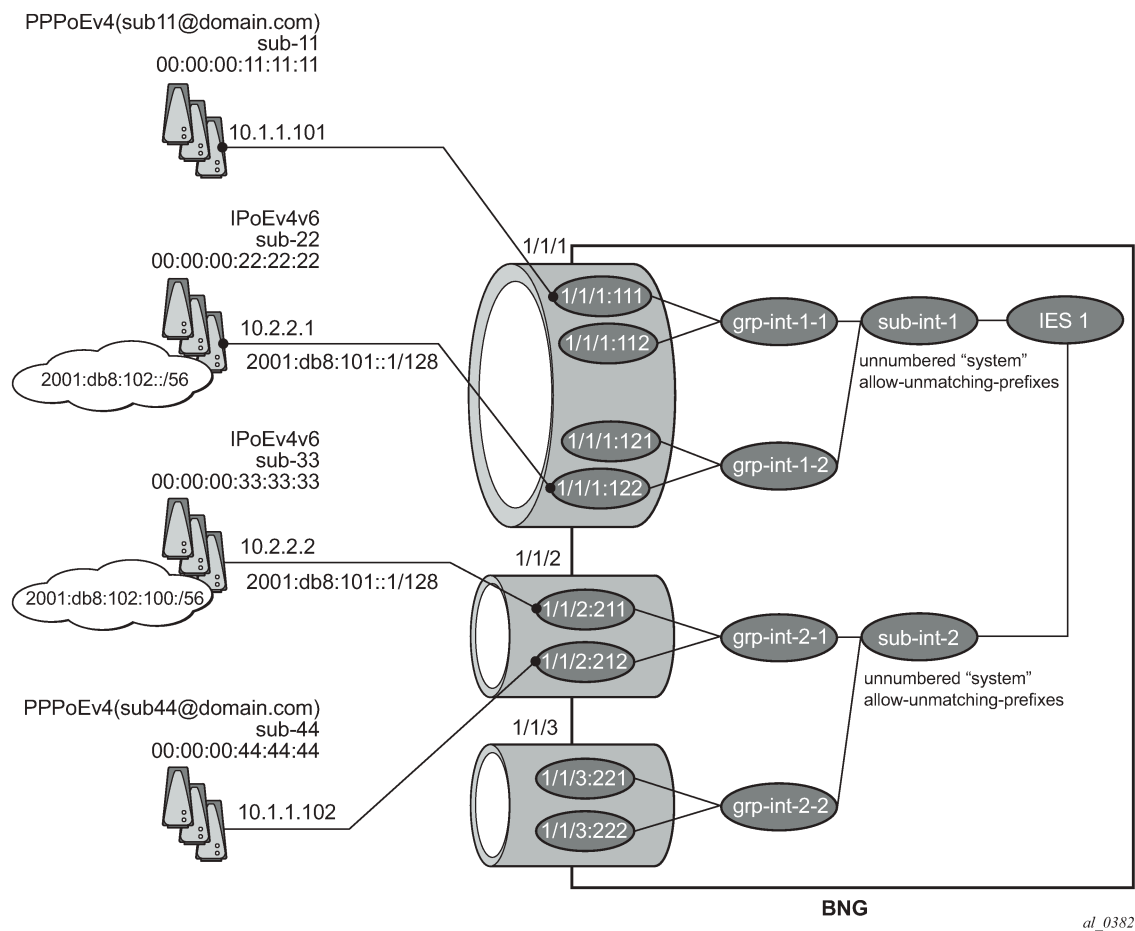
- /32 for IPv4-hosts.
- /DPL (Delegated Prefix Length) for IPv6 DP hosts, /56 in this example.
- /128 or /64 for IPv6 wan host.

Downstream traffic is forwarded based on a subscriber host table lookup. For specific network designs, subscriber host IPv4 addresses can optionally be included in the FIB with the `populate-host-routes` statement added to the subnet configuration. This is out of scope of this example.

Unnumbered Scenario

Figure 157: Unnumbered Scenario for IES 1 depicts the unnumbered scenario outlined below, including the connecting subscribers and subscriber hosts. Sub-11 and sub-44 are using single stack PPPoEv4 hosts, and sub-22 and sub-33 are using dual stack DHCP hosts. Their VLANs and the MAC addresses are shown, as are the IP addresses and prefixes assigned once they are connected.

Figure 157: Unnumbered Scenario for IES 1



The configuration for the unnumbered scenario is show below. Only the configuration items specific to the unnumbered scenario are shown.

In the unnumbered scenario the subscriber interfaces have following properties:

- IPv4:
 - No subnets configured.

- **unnumbered**, with an IPv4 interface or an IPv4 address used for IPCP negotiation.
- **no allow-unmatching-subnets**.
- IPv6:
 - No subscriber prefixes configured.
 - **allow-unmatching-prefixes**.

```
configure
  service
    ies 1
      subscriber-interface "sub-int-1" create
        unnumbered "system"
        ipv6
          delegated-prefix-len 56
          allow-unmatching-prefixes
          link-local-address FE80::EA:4B:FF
        exit
      group-interface "grp-int-1-1" create
        ipv6
          ---snip---
        exit
        arp-populate
        dhcp
          ---snip---
          lease-populate 100
          no shutdown
        exit
        authentication-policy "auth-pol-1"
        sap 1/1/1:111 create
          anti-spoof ip-mac
          sub-sla-mgmt
          ---snip---
        exit
      exit
      sap 1/1/1:112 create
        anti-spoof ip-mac
        sub-sla-mgmt
        ---snip---
      exit
      pppoe
        ---snip---
        no shutdown
      exit
    exit
  group-interface "grp-int-1-2" create
    ipv6
      ---snip---
    exit
    arp-populate
    dhcp
      ---snip---
      lease-populate 100
      no shutdown
    exit
    authentication-policy "auth-pol-1"
    sap 1/1/1:121 create
      anti-spoof ip-mac
      sub-sla-mgmt
      ---snip---
```

```
        exit
    exit
    sap 1/1/1:122 create
        anti-spoof ip-mac
        sub-sla-mgmt
        ---snip---
    exit
    exit
    pppoe
        ---snip---
        no shutdown
    exit
exit
subscriber-interface "sub-int-2" create
    unnumbered "system"
    ipv6
        delegated-prefix-len 56
        allow-unmatching-prefixes
        link-local-address FE80::EA:4B:FF
    exit
    group-interface "grp-int-2-1" create
        ipv6
            ---snip---
        exit
        arp-populate
        dhcp
            ---snip---
            lease-populate 100
            no shutdown
        exit
        authentication-policy "auth-pol-1"
        sap 1/1/2:211 create
            anti-spoof ip-mac
            sub-sla-mgmt
            ---snip---
        exit
    exit
    sap 1/1/2:212 create
        anti-spoof ip-mac
        sub-sla-mgmt
        ---snip---
    exit
    exit
    pppoe
        ---snip---
        no shutdown
    exit
exit
group-interface "grp-int-2-2" create
    ipv6
        ---snip---
    exit
    arp-populate
    dhcp
        ---snip---
        lease-populate 100
        no shutdown
    exit
    authentication-policy "auth-pol-1"
    sap 1/1/3:221 create
        anti-spoof ip-mac
        sub-sla-mgmt
        ---snip---
```

```

        exit
    exit
    sap 1/1/3:222 create
        sub-sla-mgmt
        anti-spoof ip-mac
        sub-sla-mgmt
        ---snip---
    exit
    exit
    exit
    pppoe
        ---snip---
        no shutdown
    exit
    exit
    exit
    no shutdown

```

The same mandatory and optional settings as for the numbered scenario apply.

Verification

The interfaces on the BNG are listed using following command. Notice that all subscriber and group interfaces are operational up for IPv4 and IPv6.

```
A:BNG# show router "Base" interface
```

```
=====
```

```
Interface Table (Router: Base)
```

```
=====
```

Interface-Name IP-Address	Adm	Opr (v4/v6)	Mode	Port/SapId PfxState

grp-int-1-1	Up	Up/Up	IES Grp	1/1/1
grp-int-1-2	Up	Up/Up	IES Grp	1/1/1
grp-int-2-1	Up	Up/Up	IES Grp	1/1/2
grp-int-2-2	Up	Up/Up	IES Grp	1/1/3
lb-pool4-1 10.1.1.254/24	Up	Up/Down	Network	loopback n/a
lb-pool4-2 10.1.2.254/24	Up	Up/Down	Network	loopback n/a
lb-pool4-3 10.2.1.254/24	Up	Up/Down	Network	loopback n/a
lb-pool4-4 10.2.2.254/24	Up	Up/Down	Network	loopback n/a
sub-int-1 Unnumbered If[system] FE80::EA:4B:FF/64	Up	Up/Up	IES Sub	subscriber n/a PREFERRED
sub-int-2 Unnumbered If[system] FE80::EA:4B:FF/64	Up	Up/Up	IES Sub	subscriber n/a PREFERRED
system 192.0.2.75/32 2001:DB8::75/128	Up	Up/Up	Network	system n/a PREFERRED
toDHCP-1 10.11.11.1/32 2001:DB8::11/128 FE80::E84B:FFFF:FE00:0/64	Up	Up/Up	Network	loopback n/a PREFERRED PREFERRED
toR1 192.168.12.1/24	Up	Up/Down	Network	1/1/12 n/a
toRADIUS-1 192.168.202.75/24	Up	Up/Down	Network	1/1/10 n/a

```
-----
Interfaces : 14
=====
```

```
A:BNG#
```

Successfully created hosts have forwarding state Fwding. Hosts not in the Fwding state cannot forward any data.

```
*A:BNG# show service id 1 subscriber-hosts
```

```
=====
Subscriber Host table
=====
```

Sap	IP Address	MAC Address	Subscriber	PPPoE-SID Origin	Fwding State
1/1/1:111	10.1.1.101	00:00:00:11:11:11	sub-11	1 IPCP	Fwding
1/1/1:122	10.2.2.1	00:00:00:22:22:22	sub-22	N/A DHCP	Fwding
1/1/1:122	2001:DB8:101::1/128	00:00:00:22:22:22	sub-22	N/A IPoE-DHCP6	Fwding
1/1/1:122	2001:DB8:102::/56	00:00:00:22:22:22	sub-22	N/A IPoE-DHCP6	Fwding
1/1/2:211	10.2.2.2	00:00:00:33:33:33	sub-33	N/A DHCP	Fwding
1/1/2:211	2001:DB8:101:1::1/128	00:00:00:33:33:33	sub-33	N/A IPoE-DHCP6	Fwding
1/1/2:211	2001:DB8:102:100::/56	00:00:00:33:33:33	sub-33	N/A IPoE-DHCP6	Fwding
1/1/2:212	10.1.1.102	00:00:00:44:44:44	sub-44	1 IPCP	Fwding

```
-----
Number of subscriber hosts : 8
=====
```

```
*A:BNG#
```

A variant of the show service active-subscribers command shows the subscriber hierarchy.

```
*A:BNG# show service active-subscribers hierarchy
```

```
=====
Active Subscriber hierarchy
=====
```

```
-- sub-11 (sub-prof-1)
|
|-- sap:1/1/1:111 - sla:sla-prof-1
|   |
|   |-- 10.1.1.101
|   |   00:00:00:11:11:11 - 1 (IPCP)
|   |
|
-- sub-22 (sub-prof-1)
|
|-- sap:1/1/1:122 - sla:sla-prof-1
```

*A: BNG#

- Sub-11 and sub-44 share the same IPv4 subnet even though they are connected to different subscriber interfaces.
- Sub-22 and sub-33 share the same subnet even though they are connected to different subscriber interfaces.

- Sub-22 and sub-33 are in different subscriber interfaces and do not share IPv6 prefixes in this example.

```
A:BNG# show router "Base" route-table ipv4
```

Route Table (Router: Base)					
Dest Prefix[Flags]	Type	Proto	Age	Metric	Pref
Next Hop[Interface Name]					
10.1.1.0/24	Local	Local	00h49m42s	0	
lb-pool4-1			0		
10.1.1.101/32	Remote	Sub Mgmt	00h23m24s	0	
[grp-int-1-1]			0		
10.1.1.102/32	Remote	Sub Mgmt	00h02m32s	0	
[grp-int-2-1]			0		
10.1.2.0/24	Local	Local	00h49m42s	0	

```

      lb-pool4-2
10.2.1.0/24                Local   Local   00h49m42s  0
      lb-pool4-3
10.2.2.0/24                Local   Local   00h49m42s  0
      lb-pool4-4
10.2.2.1/32               Remote  Sub Mgmt 00h27m18s  0
      [grp-int-1-2]
10.2.2.2/32               Remote  Sub Mgmt 00h27m10s  0
      [grp-int-2-1]
10.11.11.1/32             Local   Local   00h49m42s  0
      toDHCP-1
192.0.2.75/32             Local   Local   00h49m42s  0
      system
192.0.2.76/32             Remote  ISIS    00h41m48s  15
      192.168.12.2
192.168.12.0/24           Local   Local   00h49m24s  0
      toR1
192.168.202.0/24          Local   Local   00h49m24s  0
      toRADIUS-1
                                0
-----
No. of Routes: 13
Flags: L = LFA nexthop available    B = BGP backup route available
      n = Number of times nexthop is repeated
=====
A:BNG#
```

The IPv6 RIB for the base router looks as follows.

```

.A:BNG# show router "Base" route-table ipv6
=====
IPv6 Route Table (Router: Base)
=====
Dest Prefix[Flags]          Type  Proto   Age           Pref
Next Hop[Interface Name]    Metric
-----
2001:DB8::11/128           Local  Local   01h03m27s    0
      toDHCP-1
      0
2001:DB8::75/128           Local  Local   00h06m34s    0
      system
      0
2001:DB8:101::1/128        Remote  Sub Mgmt 00h36m12s    0
      [grp-int-1-2]
      0
2001:DB8:101:1::1/128      Remote  Sub Mgmt 00h35m58s    0
      [grp-int-2-1]
      0
2001:DB8:102::/56          Remote  Sub Mgmt 00h36m12s    0
      [grp-int-1-2]
      0
2001:DB8:102:100::/56      Remote  Sub Mgmt 00h35m58s    0
      [grp-int-2-1]
      0
-----
No. of Routes: 6
Flags: L = LFA nexthop available    B = BGP backup route available
      n = Number of times nexthop is repeated
=====
A:BNG# #
```

The corresponding IPv4 FIB on card 1 looks as follows.

```

A:BNG# show router "Base" fib 1 ipv4
=====
FIB Display
=====
Prefix                      Protocol
NextHop
-----
```



```

10.1.1.0/24 LOCAL
  10.1.1.0 (lb-pool4-1)
10.1.1.101/32 LOCAL
  10.1.1.101 (sub-int-1)
10.1.1.102/32 LOCAL
  10.1.1.102 (sub-int-2)
10.1.2.0/24 LOCAL
  10.1.2.0 (lb-pool4-2)
10.2.1.0/24 LOCAL
  10.2.1.0 (lb-pool4-3)
10.2.2.0/24 LOCAL
  10.2.2.0 (lb-pool4-4)
10.2.2.1/32 LOCAL
  10.2.2.1 (sub-int-1)
10.2.2.2/32 LOCAL
  10.2.2.2 (sub-int-2)
10.11.11.1/32 LOCAL
  10.11.11.1 (toDHCP-1)
192.0.2.75/32 LOCAL
  192.0.2.75 (system)
192.0.2.76/32 ISIS
  192.168.12.2 (toR1)
192.168.12.0/24 LOCAL
  192.168.12.0 (toR1)
192.168.202.0/24 LOCAL
  192.168.202.0 (toRADIUS-1)

```

Total Entries : 13

=====

The corresponding IPv6 FIB on card 1 looks as follows:

```
A:BNG# show router "Base" fib 1 ipv6
```

=====

FIB Display

```

=====
Prefix                                     Protocol
  NextHop
-----
2001:DB8::11/128                          LOCAL
  2001:DB8::11 (toDHCP-1)
2001:DB8::75/128                          LOCAL
  2001:DB8::75 (system)
2001:DB8:101::1/128                       LOCAL
  2001:DB8:101::1 (sub-int-1)
2001:DB8:101:1::1/128                     LOCAL
  2001:DB8:101:1::1 (sub-int-2)
2001:DB8:102::/56                         LOCAL
  2001:DB8:102:: (sub-int-1)
2001:DB8:102:100::/56                     LOCAL
  2001:DB8:102:100:: (sub-int-2)

```

Total Entries : 6

=====

The addresses of the individual subscriber hosts appear in the RIB and the FIB, which is the main difference with the numbered model. The forwarding plane here needs the individual addresses to forward the traffic towards the individual subscriber hosts.

Hybrid Scenario

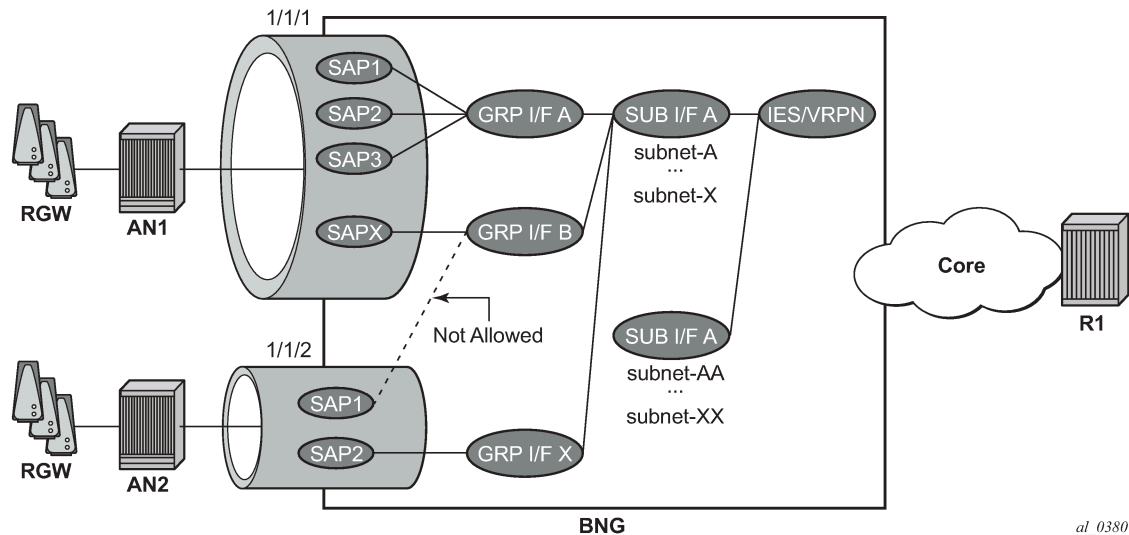
An alternative to the scenarios described above does exist in the form of a mixed numbered/unnumbered (hybrid) scenario as depicted in [Figure 158: Hybrid Configuration](#).

The subscriber interface is configured with

- One or more IPv4 subnets and/or IPv6 subscriber prefixes.
- For IPv4: the keyword **allow-unmatching-subnets**.
- For IPv6: the keyword **allow-unmatching-prefixes**.

No explicit configuration is shown as it is a mix of the numbered and the unnumbered scenario described above, and as such the behavior is mixed.

Figure 158: Hybrid Configuration



Host IP Reachability

To ensure reachability to the individual subscriber hosts, the subnets and prefixes of the subscriber interfaces/subscriber hosts need to be distributed to other routers in the network.

Three options are available:

- Without an export policy.
- With an export policy using, for example, from protocol direct.
- With an export policy using, for example, from protocol sub-mgmt.

Option 1 – No Export Policy

The key properties for the first option are:

- Subscriber interface subnets and prefixes are distributed into the network by adding the subscriber interfaces as passive interfaces to the routing protocol.
- It is used in combination with IGP based route distribution.
- It works with the numbered model only.

In this option the BNG uses IS-IS as IGP and no export policy is needed.

```
configure
router
  isis
    area-id 48.0001
    multi-topology
      ipv6-unicast
    exit
    interface "system"
      no shutdown
    exit
    interface "sub-int-1"
      passive
      no shutdown
    exit
    interface "sub-int-2"
      passive
      no shutdown
    exit
    interface "toR1"
      interface-type point-to-point
      no shutdown
    exit
  no shutdown
exit
```

The corresponding IPv4 RIB on router R1 (from [Figure 155: Components of the Routed CO Model](#)) lists the subscriber-interface subnets, not the individual subscriber host addresses.

```
A:R1# show router "Base" route-table ipv4
=====
Route Table (Router: Base)
=====
Dest Prefix[Flags]      Type  Proto  Age      Pref
Next Hop[Interface Name] Metric
-----
10.1.1.0/24             Remote ISIS   00h14m11s 15
192.168.12.1            20
10.1.2.0/24             Remote ISIS   00h14m11s 15
192.168.12.1            20
10.2.1.0/24             Remote ISIS   00h14m05s 15
192.168.12.1            20
10.2.2.0/24             Remote ISIS   00h14m05s 15
192.168.12.1            20
192.0.2.75/32           Remote ISIS   00h14m17s 15
192.168.12.1            10
192.0.2.76/32           Local  Local    62d21h22m 0
system                  0
192.168.12.0/24         Local  Local    05d04h43m 0
toBNG                   0
-----
No. of Routes: 7
Flags: L = LFA nexthop available    B = BGP backup route available
      n = Number of times nexthop is repeated
=====
```

A:R1#

The corresponding IPv6 RIB on router R1 lists the subscriber-interface prefixes, not the individual subscriber host addresses/prefixes.

```
A:R1# show router "Base" route-table ipv6
=====
IPv6 Route Table (Router: Base)
=====
Dest Prefix[Flags]                                Type  Proto  Age          Pref
Next Hop[Interface Name]                        Metric
-----
2001:DB8::75/128                                Remote ISIS    00h04m25s  15
FE80::E84B:FFFF:FE00:0-"toBNG"                  10
2001:DB8::76/128                                Local  Local    02d05h54m   0
system                                             0
2001:DB8:101::/48                                Remote ISIS    00h04m25s  15
FE80::E84B:FFFF:FE00:0-"toBNG"                  20
2001:DB8:102::/48                                Remote ISIS    00h04m25s  15
FE80::E84B:FFFF:FE00:0-"toBNG"                  20
2001:DB8:201::/48                                Remote ISIS    00h04m25s  15
FE80::E84B:FFFF:FE00:0-"toBNG"                  20
2001:DB8:202::/48                                Remote ISIS    00h04m25s  15
FE80::E84B:FFFF:FE00:0-"toBNG"                  20
2001:DEAD::/64                                  Local  Local    01h42m18s   0
toBNG                                              0
-----
No. of Routes: 7
Flags: L = LFA nexthop available    B = BGP backup route available
      n = Number of times nexthop is repeated
=====
A:R1#
```

Alternatively the same result can be achieved with OSPF/OSPFv3.

Option 2 – Export Policy (from protocol direct)

The key properties for the second option are:

- Subscriber interface subnets and prefixes are distributed into the network by applying an export policy.
- It is most typically used in combination with BGP based route distribution.
- It works with the numbered model only.

The following export policy is used for this example.

```
configure
router
  policy-options
    policy-statement "local-subnets-out"
      entry 10
        from
          protocol direct
        exit
        action accept
      exit
    exit
  exit
```

In this option the BNG relies on BGP using the policy local-subnets-out as an export policy. The neighbor address is the IPv4 system address of router R1.

```
configure
router
  autonomous-system 65536
  bgp
    group "grp-1"
      family ipv4 ipv6
      export "local-subnets-out"
      peer-as 65536
      neighbor 192.0.2.76
        advertise-label ipv6
      exit
    exit
  no shutdown
exit
```

The following command shows the IPv4 routes advertised by applying the local-subnets-out policy. The subscriber interface subnets are advertised, as are some other local subnets.

```
*A:BNG# show router bgp neighbor 192.0.2.76 advertised-routes ipv4
=====
BGP Router ID:192.0.2.75      AS:65536      Local AS:65536
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP IPv4 Routes
=====
```

Flag	Network NextHop As-Path	LocalPref Path-Id	MED Label
i	10.1.1.0/24 192.0.2.75 No As-Path	100 None	None -
i	10.1.2.0/24 192.0.2.75 No As-Path	100 None	None -
i	10.2.1.0/24 192.0.2.75 No As-Path	100 None	None -
i	10.2.2.0/24 192.0.2.75 No As-Path	100 None	None -
i	10.11.11.1/32 192.0.2.75 No As-Path	100 None	None -
i	192.0.2.75/32 192.0.2.75 No As-Path	100 None	None -
i	192.168.12.0/24 192.0.2.75 No As-Path	100 None	None -
i	192.168.202.0/24 192.0.2.75 No As-Path	100 None	None -

```
-----
Routes : 8
=====
```

```
*A:BNG#
```

The same applies for IPv6.

```
*A:BNG# show router bgp neighbor 192.0.2.76 advertised-routes ipv6
=====
BGP Router ID:192.0.2.75      AS:65536      Local AS:65536
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP IPv6 Routes
=====
Flag  Network                               LocalPref  MED
      Nexthop                             Path-Id    Label
      As-Path
-----
i     2001:DB8::11/128                     100        None
      ::FFFF:C000:24B                     None        2
      No As-Path
i     2001:DB8::75/128                     100        None
      ::FFFF:C000:24B                     None        2
      No As-Path
i     2001:DB8:101::/48                   100        None
      ::FFFF:C000:24B                     None        2
      No As-Path
i     2001:DB8:102::/48                   100        None
      ::FFFF:C000:24B                     None        2
      No As-Path
i     2001:DB8:201::/48                   100        None
      ::FFFF:C000:24B                     None        2
      No As-Path
i     2001:DB8:202::/48                   100        None
      ::FFFF:C000:24B                     None        2
      No As-Path
i     2001:DEAD::/64                      100        None
      ::FFFF:C000:24B                     None        2
      No As-Path
-----
Routes : 7
=====
*A:BNG#
```

The corresponding IPv4 RIB on router R1 lists the subscriber-interface subnets, not the individual subscriber host addresses. Notice the list also includes other routes local to the BNG.

```
*A:R1# show router "Base" route-table ipv4
=====
Route Table (Router: Base)
=====
Dest Prefix[Flags]                Type  Proto  Age      Pref
Next Hop[Interface Name]          Metric
-----
10.1.1.0/24                        Remote BGP    00h13m34s 170
      192.168.12.1                  0
10.1.2.0/24                        Remote BGP    00h13m34s 170
      192.168.12.1                  0
10.2.1.0/24                        Remote BGP    00h13m34s 170
      192.168.12.1                  0
10.2.2.0/24                        Remote BGP    00h13m34s 170
      192.168.12.1                  0
```

```

10.11.11.1/32          Remote BGP      00h13m34s 170
    192.168.12.1
192.0.2.75/32         Remote ISIS    00h15m38s 15
    192.168.12.1
192.0.2.76/32         Local  Local    03h11m54s 0
    system
192.168.12.0/24       Local  Local    03h11m25s 0
    toBNG
192.168.202.0/24      Remote BGP      00h13m34s 170
    192.168.12.1
-----
No. of Routes: 9
Flags: L = LFA nexthop available    B = BGP backup route available
      n = Number of times nexthop is repeated
=====
*A:R1#

```

The corresponding IPv6 RIB on router R1 lists the subscriber-interface prefixes, not the individual subscriber host addresses/prefixes. They are tunneled through the IPv4 core.

```

*A:R1# show router "Base" route-table ipv6
=====
IPv6 Route Table (Router: Base)
=====
Dest Prefix[Flags]          Type  Proto  Age          Pref
  Next Hop[Interface Name]      Metric
-----
2001:DB8::11/128           Remote BGP      00h00m49s 170
    192.0.2.75 (tunneled)
2001:DB8::75/128           Remote ISIS    00h54m05s 15
    FE80::E84B:FFFF:FE00:0-"toBNG"
2001:DB8::76/128           Local  Local    05h18m12s 0
    system
2001:DB8:101::/48         Remote BGP      00h00m49s 170
    192.0.2.75 (tunneled)
2001:DB8:102::/48         Remote BGP      00h00m49s 170
    192.0.2.75 (tunneled)
2001:DB8:201::/48         Remote BGP      00h00m49s 170
    192.0.2.75 (tunneled)
2001:DB8:202::/48         Remote BGP      00h00m49s 170
    192.0.2.75 (tunneled)
2001:DEAD::/64             Local  Local    05h17m42s 0
    toBNG
-----
No. of Routes: 8
Flags: L = LFA nexthop available    B = BGP backup route available
      n = Number of times nexthop is repeated
=====
*A:R1# #

```

The same export policy can be used in combination with IGP based route distribution. However, when IGP based route distribution is needed option 1 is the preferred method.

Option 3 – Export Protocol (from protocol sub-mgmt)

The key properties for the third option are:

- Host addresses and prefixes are distributed into the network by applying an export policy.

- It is most typically used in combination with BGP based route distribution, as IGP based route distribution does not scale for a large number of subscribers.
- It is most typically used for the unnumbered model, and in some cases for the numbered model.

The following export policy is used for this option.

```
configure
router
  policy-options
    policy-statement "subsc-hosts-out"
      entry 10
        from
          protocol sub-mgmt
        exit
        action accept
      exit
    exit
  exit
```

In this option the BNG relies on BGP using the policy subsc-hosts-out as an export policy.

```
configure
router
  autonomous-system 65536
  bgp
    group "grp-1"
      family ipv4 ipv6
      export "subsc-hosts-out"
      peer-as 65536
      neighbor 192.0.2.76
        advertise-label ipv6
      exit
    exit
  no shutdown
  exit
```

The following command shows the IPv4 routes advertised by applying the subsc-hosts-out policy. Now the subscriber host addresses are advertised individually.

```
*A:BNG# show router bgp neighbor 192.0.2.76 advertised-routes ipv4
=====
BGP Router ID:192.0.2.75      AS:65536      Local AS:65536
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP IPv4 Routes
=====
```

Flag	Network Nexthop As-Path	LocalPref Path-Id	MED Label
?	10.1.1.101/32 192.0.2.75 No As-Path	100 None	0 -
?	10.1.1.102/32 192.0.2.75 No As-Path	100 None	0 -
?	10.2.2.1/32 192.0.2.75	100 None	0 -


```

?      No As-Path
      10.2.2.2/32
      192.0.2.75
      No As-Path
-----
Routes : 4
=====
*A:BN#

```

For IPv6, the host addresses and prefixes are advertised.

```

*A:BN# show router bgp neighbor 192.0.2.76 advertised-routes ipv6
=====
BGP Router ID:192.0.2.75      AS:65536      Local AS:65536
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP IPv6 Routes
=====
Flag  Network                               LocalPref  MED
      Nexthop                               Path-Id    Label
      As-Path
-----
?      2001:DB8:101::1/128
      ::FFFF:C000:24B
      No As-Path
?      2001:DB8:101:1::1/128
      ::FFFF:C000:24B
      No As-Path
?      2001:DB8:102::/56
      ::FFFF:C000:24B
      No As-Path
?      2001:DB8:102:100::/56
      ::FFFF:C000:24B
      No As-Path
-----
Routes : 4
=====
*A:BN#

```

The corresponding IPv4 RIB on router R1 looks as follows. Notice the individual host addresses do appear.

```

A:R1# show router route-table ipv4
=====
Route Table (Router: Base)
=====
Dest Prefix[Flags]                Type  Proto  Age          Pref
      Next Hop[Interface Name]      Metric
-----
10.1.1.101/32                     Remote BGP    00h40m49s  170
      192.168.12.1                    0
10.1.1.102/32                     Remote BGP    00h19m53s  170
      192.168.12.1                    0
10.2.2.1/32                       Remote BGP    00h44m49s  170
      192.168.12.1                    0
10.2.2.2/32                       Remote BGP    00h44m17s  170
      192.168.12.1                    0
192.0.2.75/32                     Remote ISIS   00h59m41s  15
      192.168.12.1                    10
192.0.2.76/32                     Local  Local    01h22m11s  0
      system                          0

```

```
192.168.12.0/24          Local   Local   01h21m42s  0
toBNG
-----
No. of Routes: 7
Flags: L = LFA nexthop available    B = BGP backup route available
      n = Number of times nexthop is repeated
=====
A:R1#
```

The corresponding IPv6 RIB on router R1 looks as follows. Notice the individual host addresses and prefixes are distributed in this case.

```
A:R1# show router route-table ipv6
=====
IPv6 Route Table (Router: Base)
=====
Dest Prefix[Flags]      Type   Proto   Age          Pref
Next Hop[Interface Name] Metric
-----
2001:DB8::76/128        Local   Local   01h22m16s    0
system
2001:DB8:101::1/128      Remote  BGP      00h36m40s   170
192.0.2.75 (tunneled)
2001:DB8:101:1::1/128    Remote  BGP      00h36m40s   170
192.0.2.75 (tunneled)
2001:DB8:102::/56        Remote  BGP      00h36m40s   170
192.0.2.75 (tunneled)
2001:DB8:102:100::/56    Remote  BGP      00h36m40s   170
192.0.2.75 (tunneled)
2001:DEAD::/64          Local   Local   01h21m47s    0
toBNG
-----
No. of Routes: 6
Flags: L = LFA nexthop available    B = BGP backup route available
      n = Number of times nexthop is repeated
=====
A:R1#
```

Conclusion

This example explains how to configure and use the Routed CO model. The subscriber and the group interfaces were configured for the numbered, unnumbered and hybrid scenario, showing the flexibility of the Routed CO model in terms of subnet/prefix assignment as well as the impact on the forwarding and the reachability to and from the subscriber hosts.

Subscriber Redundancy for Routed CO

This chapter provides information about Subscriber Redundancy for Routed CO (SRRP).

Topics in this chapter include:

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

Applicability

This chapter is applicable to SR OS routers and was initially based on Release 7.0.R5. The CLI is updated to Release 16.0.R6.

Overview

This chapter focuses on the delivery of redundant services in an enhanced subscriber management (ESM) routed CO environment using Internet enhanced service (IES) or virtual private routed network (VPRN).

It is applicable to delivering high speed Internet (HSI), voice-over-IP (VoIP) and video-on-demand (VoD) to subscribers.

Redundancy is provided at two levels:

- system redundancy
- network redundancy

The system redundancy is based on the high availability features of the SR OS routers, such as component redundancy (power, fans, control processor modules, and so on) and software redundancy (service and protocol redundancy and non-stop-routing), which are not discussed here.

The network redundancy for subscriber access in an ESM routed CO environment requires that each broadband service access node (BSAN) is dual-homed to two SR OS routers, either in a point-to-point fashion with the BSANs having direct physical connectivity to the SR OS routers, or by having Layer 2 aggregation between the BSANs and the SR OS routers.

This connection will operate in a primary-standby relationship providing both link and system redundancy for the subscribers on the BSAN when accessing the configured services.

Subscriber redundancy for routed-CO aims to minimize the outage due to a failure.

This chapter provides configuration and troubleshooting commands for SRRP with **static-host ip-mac**.


Knowledge of the triple play service delivery architecture (TPSDA) concepts is assumed throughout this document.

There are three components on SR OS routers to implement network redundancy:

1. Redundant access from the BSANs to two SR OS routers provided by the SRRP.
2. Mirroring of subscriber state between two SR OS routers is achieved through the multi-chassis synchronization (MCS) protocol.
3. Backup spoke SDP traffic path between two SR OS peers.

These components are shown in [Figure 159: Network Redundancy Components for ESM Routed CO](#).

Figure 159: Network Redundancy Components for ESM Routed CO



OSSG420

The following configuration tasks should be done first and are not explained in this chapter:

- Basic service router configuration (system interface, IGP, MPLS, BGP)
- Routed CO service topology: VPRN or IES service with subscriber and group interface on broadband network gateways (BNGs)
- ESM configuration
- Static host configuration

This chapter will focus on SRRP in a VPRN service subscriber-interface on BNG (routed-CO). In case of routed CO, it is also possible to configure SRRP in the base routing instance using an IES service.

SRRP Protocol

The SRRP protocol operates on a specific SAP within the group interface under the subscriber-interface of an IES/VPDN service. Through a method similar to the virtual router redundancy protocol (VRRP), it provides a set of default gateways to the subscribers on the BSAN. These are active on the SR OS router in the master state (primary) and inactive on the router in the backup state. Upstream and downstream traffic is forwarded through the master. If the backup loses connectivity with the primary (for example, fails to receive SRRP messages from the primary), it transitions to the master state and takes over the ownership of the default gateways and the responsibility for forwarding traffic to and from the subscriber. This provides redundancy from the BSAN toward the provider network.

If an SRRP fail over were to occur, it is important that the subscriber state (IP/MAC addresses, QOS profiles, etc.) be immediately available on the new SRRP primary; otherwise, subscriber traffic will be dropped due to the anti-spoofing security. The subscriber state is synchronized through the MCS protocol, which mirrors the subscriber state between the two peers. This allows both peers to know the details of the active subscribers and therefore forward traffic on their behalf with the correct QOS actions both to and from the BSAN if that peer is the SRRP primary.

The last topic relates to the forwarding from the provider network to the subscribers. If the IP routing causes this traffic to forward through the SRRP primary, then the traffic will automatically be forwarded to the subscriber. However, if the provider routing scheme causes traffic destined to a subscriber to arrive at the router in the SRRP backup state for that subscriber, it will be dropped as the backup does not forward traffic out of the subscriber SAPs.

To avoid this, a redundant interface is configured between the two SR OS routers under the subscriber/group-interfaces. Any traffic arriving on the router for an active subscriber, where its associated SRRP instance is in the backup state, will be forwarded through the redundant interface to the SRRP primary, which in turn forwards the traffic to the subscriber.

In addition to successful forwarding the traffic, this operation also maintains the subscriber QOS compliance as all traffic for a given subscriber enters and exits the routed-CO interface through a single SAP, allowing the associated IOM hardware to perform the correct QOS actions.

Configuration

Subscriber Interface Configuration

Redundant Default Gateway

The redundant default gateway IP addresses must be configured under the subscriber-interface (within the IES/VP RN service) for each subnet defined.

Three subnets are configured under the subscriber-interface sub-int-1, each with a **gw-ip-address** which is used as the default gateway by the subscribers in that subnet.

```
# on BNG-2
configure
service
    vprn 1 customer 1 create
    ---snip---
    subscriber-interface "sub-int-1" create
        address 10.2.0.2/16 gw-ip-address 10.2.0.254
        address 10.3.0.2/16 gw-ip-address 10.3.0.254
        address 10.4.0.2/16 gw-ip-address 10.4.0.254
    ---snip---
    exit
exit
exit
exit
```

The **gw-ip-address** could be a virtual (unused) address in this subnet or the address of one of the actual subscriber-interfaces on the two routers, but it must not be used as a subscriber address.

If DHCP were to be used, the associated subscriber-interface address should be used as the gi-address configured for DHCP under the group-interface (will not be covered here as static host is used). This ensures that the offer returned from the DHCP server and arriving at the SRRP backup (rather than primary) will be forwarded by the backup SRRP router to the primary SRRP router through the redundant interface.

In environments where there are many subscribers, it will take time to synchronize the subscriber state between the peers when the subscriber-interface is enabled (perhaps, after a reboot). In order to ensure that the state has time to get synchronized, a hold timer can be applied to the subscriber interface. The optional **init-only** parameter can be added to use this timer only after a reboot.

```
*A:BNG-2>config>service>vprn>red-if>hold-time#
[no] down          - Configure the hold time when the interface is coming up
[no] up            - Configure the hold time when the interface is going down

*A:BNG-2>config>service>vprn>red-if>hold-time#
```

```
# on BNG-2
configure
  service
    vprn 1 customer 1 create
      subscriber-interface "sub-int-1" create
        hold-time
          down ip 1200 init-only
        exit
      exit
    exit
  exit
exit
```

Group Interface Configuration

The group interface *group-int-1* providing connectivity to the BSAN is configured under the subscriber interface:

```
# on BNG-2
configure
  service
    vprn 1 customer 1 create
      subscriber-interface "sub-int-1" create
        group-interface "group-int-1" create
          ---snip---
        exit
      exit
    exit
  exit
exit
```

Static Host Configuration

First enable the sub-sla-mgmt and sub-ident-policy *sub-id-default* under **sap 1/1/3:1**, then define a static host (ip-mac) with **sla-profile sla-profile-1**, **sub-profile sub-profile-1** and **subscriber static-host-routed-10.2.0.3**.

```
# on BGG-2
configure
```

```

service
  vprn 1 customer 1 create
    subscriber-interface "sub-int-1" create
      group-interface "group-int-1" create
        sap 1/1/3:1 create
          sub-sla-mgmt
            sub-ident-policy "sub-id-default"
            multi-sub-sap
            no shutdown
          exit
          static-host ip 10.2.0.3 mac 00:00:00:00:00:01 create
            sla-profile "sla-profile-1"
            sub-profile "sub-profile-1"
            subscriber "static-host-routed-10.2.0.3"
            no shutdown
          exit
        exit
      exit
    exit
  exit
exit

```

SRRP Configuration

In order for the redundant gateway information to be used by subscribers through SAPs belonging to a particular group-interface, an SRRP instance must be added in the group interface context.

```

# on BNG-2
configure
  service
    vprn 1 customer 1 create
      subscriber-interface "sub-int-1" create
        group-interface "group-int-1" create
          srrp 1 create
            ---snip---
          exit
        exit
      exit
    no shutdown
  exit
exit

```

At this point, any subscriber ARPing for the gw-ip-address will receive a response from the SRRP primary with a source MAC of 00-00-5E-00-01-<xx>, where <xx> is the first byte of the SRRP identifier in hexadecimal, so in this case for SRRP=1 the source MAC will be 00-00-5E-00-01-01.

The redundant default gateway MAC address could be explicitly configured, if desired, by use of the **gw-mac** parameter.

```

*A:BNG-2>config>service>vprn>sub-if>grp-if>srrp# gw-mac
- gw-mac <mac-address>
- no gw-mac

<mac-address>          : xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx

*A:BNG-2>config>service>vprn>sub-if>grp-if>srrp#

```

There can be only one SRRP instance per group-interface and all SRRP identifiers must be unique per system.

When an SRRP instance is enabled, or when there is an SRRP failover from one device to another, gratuitous ARPs are sent on all VLANs associated with this instance for all gw-ip-addresses (for example, on all subscriber SAPs and on the SRRP message path SAP in the associated group interface). This allows all downstream devices to relearn the path to the new primary.

The SRRP messages are normally forwarded through the BSAN, thereby verifying the connectivity from one router through the BSAN to the other router. In order to achieve this, a non-subscriber SAP must be configured under the **group-interface** which is referenced in the SRRP configuration by the **message-path** parameter. This not only selects the SAP to be used for the SRRP messages but also avoids the subscriber anti-spoofing from automatically dropping the received messages (as there would be no subscriber IP-to-MAC entry corresponding to the received information) and causing both peers to become primary.

The message-path SAP configuration effectively disables IP-MAC anti-spoofing on that SAP.

```
# on BNG-2
configure
  service
    vprn 1 customer 1 create
      subscriber-interface "sub-int-1" create
        group-interface "group-int-1" create
          ----snip---
          sap 1/1/3:2 create
          exit
          srrp 1 create
            message-path 1/1/3:2
            no shutdown
          exit
        exit
      exit
    exit
  no shutdown
exit
exit
exit
exit
```

The SRRP messages are then not sent on the same SAP as the subscriber data traffic, under the assumption that the path traversed by the SRRP messages is the same as the path for the subscriber data; if this is not the case, then the SRRP state would not necessarily reflect a failure in the data path.

The primary of the SRRP instance generates advertisement messages at the keep-alive-interval (which is encoded in the message) ranging from one (1) to 100 in multiples of 100ms with a default of 10 (for example, 1 second). The SRRP backup will monitor the reception of these messages and assume the role of the primary if three consecutive messages are not received.

At all times the keep-alive-interval of the primary is used.

```
*A:BNG-2>config>service>vprn>sub-if>grp-if>srrp# keep-alive-interval
- keep-alive-interval <interval>
- no keep-alive-interval

<interval>          : [1..100] tenths of a second

*A:BNG-2>config>service>vprn>sub-if>grp-if>srrp#
```


Only two devices can participate in an SRRP protocol exchange for a given SRRP instance, this being another difference from VRRP which allows more potential backup devices. This is a consequence of the direct relationship between the SRRP instance and the associated redundant interface and MCS peering.

This protocol exchange is also used for the primary/backup election, based on the priority (1 to 254) configured in the SRRP instance. The device with the highest priority will become primary.

```
*A:BNG-2>config>service>vprn>sub-if>grp-if>srrp# priority
- no priority
- priority <priority>

<priority>          : [1..254]

*A:BNG-2>config>service>vprn>sub-if>grp-if>srrp#
```

The message source IP address (system IP address) is used as a tie break when the priorities are the same (the lower IP address becomes the primary). The primary/backup status is per SRRP instance (not per IP address). For example, the primary is the active gateway for all gw-ip-addresses under the subscriber interface for the associated group-interface (this is true even if the backup is the IP address owner for one of the gw-ip-addresses). Priority 0 is sent by the primary when it is transitioning to the backup role due to the appearance of a high priority peer. Higher priority backups always preempt a lower priority primary.

A basic form of load distribution can be achieved by having the primary SRRP for some group-interfaces on one peer and the primary for other group interfaces on the other peer. Clearly, a failure may cause all primary SRRPs to be active on a single peer, which must be taken into account when designing the network.

The minimum keep-alive-timer of 1 is configured, together with the message-path defining the SAP to be used for the SRRP messages. The priority is set to 250 (default is 100) in order to force this peer to be the SRRP primary when both peers are active.

```
# on BNG-2
configure
  service
    vprn 1 customer 1 create
      subscriber-interface "sub-int-1" create
      group-interface "group-int-1" create
      ---snip---
      srrp 1 create
        keep-alive-interval 1
        message-path 1/1/3:2
        priority 250
        no shutdown
      exit
    exit
  exit
exit
```

SRRP Configuration Notes

A VRRP policy statement can be added to the SRRP instance definition in order to dynamically adjust the SRRP priority based on certain non-SRRP related events occurring (for example, port down, LAG degrade, host unreachable or route unknown).

The gw-ip-addresses are accessible by active subscribers, for example, regardless which peer is the primary, an active subscriber can ping or telnet to its associated gw-ip-address (clearly, filters can be configured to control this).

Bi-directional Forwarding Detection

Bi-directional forwarding detection (BFD) can be configured with SRRP to speed up the convergence.

```
# on BNG-2
configure
service
  vprn 1 customer 1 create
  subscriber-interface "sub-int-1" create
  group-interface "group-int-1" create
  ---snip---
  srrp 1 create
  ---snip---
  bfd-enable 2 interface "bfd-1" dst-ip 10.1.1.1
  no shutdown
  exit
exit
exit
exit
exit
exit
```

An IES service needs to be created for the BFD session.

```
# on BNG-1
configure
service
  ies 2 customer 1 create
  interface "bfd-1" create
  address 10.1.1.1/31
  bfd 100 receive 100 multiplier 3
  sap 1/1/3:3 create
  exit
  exit
  no shutdown
exit
exit
exit
exit
```

Monitoring In-Band Communications Path

In order to monitor the in-band communications path between the subscribers and two routers, SRRP uses a slightly modified VRRP advertisement message.

The SRRP message destination IP address (224.0.0.18) and IP protocol number (112) are the same as for VRRP but there are changes in the following areas:

- The source IP address of the message is the system IP address, as opposed to the interface IP address.
- The protocol version has been set to eight (8) (the current version of VRRP is two (2)).
- The source MAC address is included instead of the virtual router IP addresses (this being 00-00-5E-00-01-<xx>, where <xx> is the first byte of the SRRP identifier in hexadecimal).

Synchronizing the SRRP Peer State

In order to troubleshoot an SRRP environment, the state of each peer is synchronized with the other peer through multi-chassis synchronization (MCS). MCS is a proprietary protocol used for synchronizing application state between two peers. SRRP will function without synchronizing its state but this synchronization allows for the current state of both the local and remote peers to be displayed and appropriate error messages to be reported when the peer state is not correct. It also allows the primary SRRP subscriber-interface to be pinged through the backup peer (through the redundant interface).

To link information being mirrored between two routers, a **sync-tag** value is configured to correspond to either an entire port/LAG or under a port/LAG for a VLAN range. This allows each router to know exactly which information should be in sync on each device. The sync-tag must be unique on the two peers involved.

This example configuration shows only the SRRP aspects. The SRRP instance has been configured for MCS peer 192.0.2.1 using VLANs 1-2 on port 1/1/3. Here, a sync-tag is associated with the SRRP instance.

```
# on BNG-2
configure
  redundancy
    multi-chassis
      peer 192.0.2.1 create
        authentication-key "sync-testing"
        sync
          srrp
            sub-mgmt ipoe
            port 1/1/3 "st1" create
              range 1-2 sync-tag "st1"
            exit
          no shutdown
        exit
      no shutdown
    exit
  no shutdown
exit
```

Alternatively, if the information needs to be synchronized for all VLANs on port 1/1/3, then the following port command could be used instead of the preceding port-plus-ranges shown.

```
# on BNG-2
configure
  redundancy
    multi-chassis
      peer 192.0.2.1 create
        authentication-key "sync-testing"
        sync
          srrp
            sub-mgmt ipoe
            port 1/1/3 sync-tag "st1" create
          exit
        no shutdown
      exit
    no shutdown
  exit
exit
```

The VLANs used within the group interfaces must match between the two peers, clearly the physical ports identifiers may differ.

Multi-Chassis Synchronization

Multi-chassis synchronization (MCS) is a general propriety protocol used to synchronize information between two peers. It can be used for the several applications, such as:

- IGMP
- IGMP snooping
- Subscriber management
- Subscriber router redundancy protocol

This chapter only covers the subscriber management and SRRP applications.

Subscriber Management Synchronization

In order to ensure that the QOS defined for a subscriber is adhered to, all traffic for a given subscriber needs to be forwarded by a single port. When an MSAN is dual-homed to two routers, this is achieved using the SRRP protocol (described above) and the redundant interface (described below); specifically, the traffic is forwarded through the SRRP primary of the related group-interface.

When a subscriber is created on the primary SRRP, a host route (/32) for its IP address is inserted in the FIB pointing towards the appropriate group-interface.

The same subscriber on the backup peer will have a host route in the IP FIB pointing to the redundant interface. On the backup peer, this requires the subscriber subnet to also be present in the FIB, which in turn requires one of the following:

- The local subscriber-interface is up.
- The subscriber subnets are learned from the active broadband network gateway (BNG) through the routing protocol.
- Forcing the subscriber-interface to stay up by creating an additional dummy group interface with the **oper-up-while-empty** command.

```
# on BNG-2
configure
service
  vprn 1 customer 1 create
    subscriber-interface "sub-int-1" create
    group-interface "dummy-group-int-1" create
      oper-up-while-empty
    exit
  exit
exit
exit
exit
exit
```

Redundant Interface

The requirement for the redundant interface stems from the situation where traffic destined to a subscriber arrives on the router where the associated SRRP state is not primary.

When the SRRP state is backup for a particular group-interface, subscriber traffic is normally not forwarded in/out of the associated subscriber SAPs. Also traffic could arrive where the specific group-interface for that subscriber is down. These situations could occur due to the regular routing within the provider network or temporarily during an SRRP failover. Note that as the subscriber subnets are configured under the subscriber-interface, it is not possible to stop advertising these subnets into the provider core in the case where only a subset of the group interfaces are down or the associated SRRP instances are in backup. The advertisement of the subscriber subnet could therefore attract traffic to the router, while not being the SRRP primary.

In these cases, traffic must be sent to the active SRRP router in order to be forwarded to the subscriber. This is achieved through the configuration of a redundant interface between two SRRP peers, protecting against failures of related group interfaces.

The redundant interface is configured under the IES/VRPN service. It must use a single pseudowire, configured as a spoke SDP, to provide connectivity to the peer router. This is essential as it avoids any possibility of the traffic being misrouted by any other system between the two peers. It can either use a /31 IP subnet mask or a longer mask with the remote IP being explicitly specified.

```
# on BRG-2
configure
service
  vprn 1 customer 1 create
  ---snip---
  redundant-interface "bng-2-bng-1-vprn-1" create
  address 192.168.4.1/31
  spoke-sdp 21:1 create
  exit
exit
---snip---
exit
exit
exit
```

If a non /31 address is used, the remote IP address will be required.

```
*A:BNG-2>config>service>vprn# redundant-interface "bng-2-bng-1-vprn-1"
*A:BNG-2>config>service>vprn>red-if# address 192.168.4.1/30
INFO: PIP #1399 Invalid or unspecified remote IP address - Non /31 address requires remote IP
address
```

The remote-ip address can be defined by the following command:

```
*A:BNG-2>config>service>vprn>red-if# address 192.168.4.1/30 remote-ip 192.4.1.2
```

Then each group-interface must be assigned a redundant interface, as follows:


```
# on BNG-2
configure
service
  vprn 1 customer 1 create
  ---snip---
  subscriber-interface "sub-int-1" create
  ---snip---
  group-interface "group-int-1" create
  redundant-interface "bng-2-bng-1-vprn-1"
  oper-up-while-empty
  ---snip---
exit
exit
```

```
no shutdown
exit
exit
exit
```

Only one redundant interface is required for a given IES/VP RN service on each peer, though it is possible to create multiple redundant interfaces and assign group interfaces to each individually. Clearly, each redundant interface needs to terminate on a matching redundant interface in the corresponding peer service.

When traffic arrives from the core network for an active subscriber on a SAP in group-interface *group-int-1*, and if the associated SRRP instance is in the backup state, then this traffic will be forwarded over the **redundant-interface bng-2-bng-1-vprn-1** to the peer router. It will then be forwarded to the subscriber as the associated SRRP instance will be in the primary state.

The information about the redundant interfaces is mirrored through MCS as part of the SRRP synchronization.



Note:

When changing the local and remote redundant-interface addresses afterward, **shutdown** and **no shutdown** of the redundant interface on one side is required in order to guarantee subscriber-interface IP reachability.

Show and Debug Commands

Routing Table Related Information

A host route (/32) for the static host is inserted in the FIB pointing toward the appropriate group-interface.

```
*A:BNG-2# show router 1 route-table protocol sub-mgmt

=====
Route Table (Service: 1)
=====
Dest Prefix[Flags]                                Type   Proto   Age           Pref
  Next Hop[Interface Name]                        Metric
-----
10.2.0.1/32                                         Remote Sub Mgmt 01h39m55s 0
  [bng-2-bng-1-vprn-1]                             0
10.2.0.3/32                                         Remote Sub Mgmt 02h05m21s 0
  [group-int-1]                                       0
10.3.0.1/32                                         Remote Sub Mgmt 01h39m55s 0
  [bng-2-bng-1-vprn-1]                             0
10.4.0.1/32                                         Remote Sub Mgmt 01h39m55s 0
  [bng-2-bng-1-vprn-1]                             0
-----
No. of Routes: 4
Flags: n = Number of times nexthop is repeated
      B = BGP backup route available
      L = LFA nexthop available
      S = Sticky ECMP requested
=====
*A:BNG-2#
```

The same subscriber on the backup peer will have a host route in the FIB pointing to the redundant interface.

```
*A:BNG-1# show router 1 route-table protocol sub-mgmt

=====
Route Table (Service: 1)
=====
Dest Prefix[Flags]                Type   Proto   Age           Pref
  Next Hop[Interface Name]              Metric
-----
10.2.0.2/32                        Remote Sub Mgmt 01h42m50s  0
      [bng-1-bng-2-vprn-1]              0
10.2.0.3/32                        Remote Sub Mgmt 02h08m31s  0
      [bng-1-bng-2-vprn-1]              0
10.3.0.2/32                        Remote Sub Mgmt 01h42m50s  0
      [bng-1-bng-2-vprn-1]              0
10.4.0.2/32                        Remote Sub Mgmt 01h42m50s  0
      [bng-1-bng-2-vprn-1]              0
-----
No. of Routes: 4
Flags: n = Number of times nexthop is repeated
      B = BGP backup route available
      L = LFA nexthop available
      S = Sticky ECMP requested
=====
*A:BNG-1#
```

Subscriber Related Information

To check the subscriber information:

```
*A:BNG-2# show service id 1 subscriber-hosts

=====
Subscriber Host table
=====
Sap      Subscriber
IP Address
MAC Address  PPPoE-SID Origin  Fwding State
-----
1/1/3:1      static-host-routed-10.2.0.3
10.2.0.3
00:00:00:00:00:01  N/A      Static      Fwding
-----
Number of subscriber hosts : 1
=====
*A:BNG-2#
```

To check the subscriber details:

```
*A:BNG-2# show service id 1 subscriber-hosts mac 00:00:00:00:00:01 detail

=====
Subscriber Host table
=====
Sap      Subscriber
IP Address
MAC Address  PPPoE-SID Origin  Fwding State
-----
```

```
1/1/3:1          static-host-routed-10.2.0.3
10.2.0.3
00:00:00:00:00:01  N/A      Static      Fwding
-----
Subscriber-interface : sub-int-1
Group-interface     : group-int-1
Sub Profile         : sub-profile-1
SLA Profile         : sla-profile-1
App Profile         : N/A
Egress Q-Group      : N/A
Egress Vport        : N/A
Acct-Session-Id     : 021BFF000000005CA701E9
Acct-Q-Inst-Session-Id: 021BFF0000000015CA701E9
Address Origin      : Static
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status  : N/A
OT HTTP Rdr Fltr Src : N/A
HTTP Rdr URL Override : N/A
GTP local break-out : No
DIAMETER session ID Gx: N/A
-----
Number of subscriber hosts : 1
=====
*A:BNG-2#
```

The same command on the peer would show the same information except for the port identifier part of the SAP, which is specific to the peer and may differ.

MCS Redundancy Related Information

The high-level state of MCS can be seen in the following output:

```
*A:BNG-2# show redundancy multi-chassis all

=====
Multi-Chassis Peers
=====
Peer Info          Client      Admin      Oper      State
-----
Peer Address : 192.0.2.1
Source Addre*: 192.0.2.2
Peer Admin   : inService
Authenticati*: None
* indicates that the corresponding row element may have been truncated.
MC-Sync:      inService    inService    inSync
MC-Ring:      --          --          --
MC-Endpt:     --          --          --
MC-Lag:       outOfService outOfService --
MC-IPsec:     --          --          Disabled
=====
*A:BNG-2#
```

Information about the peering, the use of authentication, the state of MCS (Enabled) and the fact that MCS is inSync is shown.

```
*A:BNG-2# show redundancy multi-chassis sync

=====
Multi-chassis Peer Table
=====
```



```

Peer
-----
Peer IP Address      : 192.0.2.1
Description          : (Not Specified)
Authentication       : Enabled
Source IP Address    : 192.0.2.2
Admin State          : Enabled
Warm standby         : No
Remote warm standby  : No
-----
Sync-status
-----
Client Applications  : SUBMGMT-IPOE SUBMGMT-PPPOE SRRP
Sync Admin State     : Up
Sync Oper State      : Up
Sync Oper Flags      :
DB Sync State        : inSync
Num Entries          : 26
Lcl Deleted Entries  : 0
Alarm Entries        : 0
OMCR Standby Entries : 0
OMCR Alarm Entries   : 0
Rem Num Entries      : 26
Rem Lcl Deleted Entries : 0
Rem Alarm Entries    : 0
Rem OMCR Standby Entries : 0
Rem OMCR Alarm Entries : 0
=====
*A:BNG-2#

```

In the output, it can be seen that SUBMGMT-IPOE, SUBMGMT-PPPOE and SRRP are client applications and they are inSync with 26 database entries both on this peer and the remote peer.

If the preceding command requested detailed output for peer 192.0.2.1, additional information would be shown.

```

*A:BNG-2# show redundancy multi-chassis sync peer 192.0.2.1 detail

=====
Multi-chassis Peer Table
=====
Peer
-----
Peer IP Address      : 192.0.2.1
Description          : (Not Specified)
Authentication       : Enabled
Source IP Address    : 192.0.2.2
Admin State          : Enabled
Warm standby         : No
Remote warm standby  : No
-----
Sync-status
-----
Client Applications  : SUBMGMT-IPOE SUBMGMT-PPPOE SRRP
Sync Admin State     : Up
Sync Oper State      : Up
Sync Oper Flags      :
DB Sync State        : inSync
Num Entries          : 26
Lcl Deleted Entries  : 0
Alarm Entries        : 0
OMCR Standby Entries : 0

```

```

OMCR Alarm Entries      : 0
Rem Num Entries        : 26
Rem Lcl Deleted Entries : 0
Rem Alarm Entries      : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries : 0

=====
MCS Application Stats
=====
Application            : igmp
Num Entries            : 0
Lcl Deleted Entries    : 0
Alarm Entries          : 0
OMCR Standby Entries   : 0
OMCR Alarm Entries     : 0

---snip---

-----
Application            : srrp
Num Entries            : 26
Lcl Deleted Entries    : 0
Alarm Entries          : 0
OMCR Standby Entries   : 0
OMCR Alarm Entries     : 0

-----
Rem Num Entries        : 26
Rem Lcl Deleted Entries : 0
Rem Alarm Entries      : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries : 0

-----

---snip---

=====
Ports synced on peer 192.0.2.1
=====
Port/Encap              Tag
-----
1/1/3                   st1
=====

---snip---

=====
Diameter proxy instances synced on peer 192.0.2.1
=====
Diameter-Peer-Policy      Tag
-----
No instances found
=====
=====
*A:BNG-2#

```

This shows that there are 26 entries on both peers for SRRP.

If the hold-time parameter is configured under the subscriber-interface, in order to allow time for the MCS to fully synchronize, its setting and current expiry time can be seen as follows:

```
A:BNG-2# show service id 1 interface "sub-int-1" detail
```

```

=====
Interface Table
=====

-----
Interface
-----
If Name       : sub-int-1
Admin State   : Up                               Oper (v4/v6)   : Down/--
Down Reason Code : delayedStartEnabled
Down Reason V4  : delayedStartEnabled
Down Reason V6  : ifProtoOperDown
Protocols      : None
IP Addr/mask    : 10.2.0.2/16
HoldUp-Time     : 0                               Track Srrp Inst : 0
IP Addr/mask    : 10.3.0.2/16
HoldUp-Time     : 0                               Track Srrp Inst : 0
IP Addr/mask    : 10.4.0.2/16
HoldUp-Time     : 0                               Track Srrp Inst : 0
Ignore Port State: None
-----

Details
-----
Description    : (Not Specified)
If Index       : 5                               Virt. If Index  : 5
Last Oper Chg  : 08/18/2016 10:22:20           Global If Index : 258
Mon Oper Grp   : None
Srrp En Rtng   : Disabled                       Hold time       : N/A
V4 Delay IfUp : 300 init-only                 V4 Time to IfUp : 278
                                                    Unmatching Subnet : No
                                                    Unmatching Pfxs   : No

If Type        : VPRN Sub

DHCP Details
Gi-Addr        : Not configured                   Gi-Addr as Src Ip : Disabled
Virt. subnet    : disabled

=====
Interface sub-int-1 group-interfaces
=====
Interface-Name      Adm      Opr(v4/v6)  Mode      Port/SapId
IP-Address          PfxState
-----
group-int-1         Up       Up/--      VPRN G*   1/1/3
-----
Group-Interfaces : 1
=====
* indicates that the corresponding row element may have been truncated.
-----
Interfaces : 1
=====
A:BNG-2#

```

Tool Dump Commands Related Information

The database entries can be view in more detail with the **tools dump redundancy multi-chassis** command.

```
*A:BNG-2# tools dump redundancy multi-chassis sync-database
```

The following totals are for:

```
peer ip ALL, port/lag/sdp ALL, sync-tag ALL, application ALL
Valid Entries:                26
Locally Deleted Entries:      0
Locally Deleted Alarmed Entries: 0
Pending Global Delete Entries: 0
Omcrr Alarmed Entries:        0
Omcrr Standby Entries:        0
Associated Shared Records (ALL): 0
Associated Shared Records (LD): 0
*A:BNG-2#
```

The output of the “tool dump redundancy multi-chassis” command is as follows:

```
*A:BNG-2# tools dump redundancy multi-chassis srrp-sync-database
Tag Key:  sap = 1/1/3:2
Key Info: (Type/Owner)
          SMK_BASE_CONFIG / 192.0.2.1
Data Info:
  srrpId: 1    svcId: 1    svcType: VPRN
  system IP: 0xc0000201    Group interface MAC: 02:19:01:01:00:03
  Gateway MAC: 00:00:5e:00:01:01
  Subscriber interface name: sub-int-1

Tag Key:  sap = 1/1/3:2
Key Info: (Type/Owner)
          SMK_BASE_CONFIG / 192.0.2.2
Data Info:
  srrpId: 1    svcId: 1    svcType: VPRN
  system IP: 0xc0000202    Group interface MAC: 02:1b:01:01:00:03
  Gateway MAC: 00:00:5e:00:01:01
  Subscriber interface name: sub-int-1

Tag Key:  sap = 1/1/3:2
Key Info: (Type/Owner)
          SMK_GRP_IF / 192.0.2.1
Data Info:
  Group interface name: group-int-1
  Redundant interface name: bng-1-bng-2-vprn-1
  Redundant Interface IP/Mask:
    IP: 0xc0a80400 Mask: 0xffffffffe
  AdminUp: Up, OperState: SRRP_STATE_BACKUP_SHUNT, InUsePriority: 100, Red-If OK:
  Yes, MessageSap OK: Yes

Tag Key:  sap = 1/1/3:2
Key Info: (Type/Owner)
          SMK_GRP_IF / 192.0.2.2
Data Info:
  Group interface name: group-int-1
  Redundant interface name: bng-2-bng-1-vprn-1
  Redundant Interface IP/Mask:
    IP: 0xc0a80401 Mask: 0xffffffffe
  AdminUp: Up, OperState: SRRP_STATE_MASTER, InUsePriority: 250, Red-If OK: Yes,
  MessageSap OK: Yes

Tag Key:  sap = 1/1/3:2
Key Info: (Type/Owner)
          SMK_GRP_IF_SAP_BUCKET0 / 192.0.2.1
Data Info:
---snip---
```

```
Tag Key:  sap = 1/1/3:2
Key Info: (Type/Owner)
```

```
SMK_GRP_IF_SAP_BUCKET9 / 192.0.2.2
Data Info:

Tag Key:  sap = 1/1/3:2
Key Info: (Type/Owner)
SMK_SUBNET_INFO / 192.0.2.1 vRtrId 2, ifIdx 5
Data Info:
Subscriber IP Addr: 10.2.0.1      Mask: 0xffff0000      Gateway: 10.2.0.254
Subscriber IP Addr: 10.3.0.1      Mask: 0xffff0000      Gateway: 10.3.0.254
Subscriber IP Addr: 10.4.0.1      Mask: 0xffff0000      Gateway: 10.4.0.254

Tag Key:  sap = 1/1/3:2
Key Info: (Type/Owner)
SMK_SUBNET_INFO / 192.0.2.2 vRtrId 2, ifIdx 5
Data Info:
Subscriber IP Addr: 10.2.0.2      Mask: 0xffff0000      Gateway: 10.2.0.254
Subscriber IP Addr: 10.3.0.2      Mask: 0xffff0000      Gateway: 10.3.0.254
Subscriber IP Addr: 10.4.0.2      Mask: 0xffff0000      Gateway: 10.4.0.254

*A:BNG-2#
```

Also shown are the port/VLANs synchronized with their respective sync-tags.

To further troubleshoot and debug this configuration, there are commands to both dump the sync and SRRP MCS information and to dump the SRRP database:

```
tools dump redundancy multi-chassis sync-database [application {sub-mgmt|srrp}]
```

The command provides the same information as the equivalent show commands. However, the detailed version gives more information about the contents of the sync-database.

For SRRP, there are entries for the base configuration, group interface and subnet information for each of the SRRP instances. This should show corresponding entries for the local and remote peer. Specifying the **sync-tag st1** shows only the information for SRRP instance 1.

```
*A:BNG-2# tools dump redundancy multi-chassis sync-database application srrp
                                     sync-tag st1 detail

If no entries are present for an application, no detail will be displayed.

FLAGS LEGEND: ld - local delete; da - delete alarm; pd - pending global delete;
              oal - omcr alarmed; ost - omcr standby

Peer Ip 192.0.2.1

Application SRRP
Sap-id      Client Key
SyncTag
deleteReason code and description
DLen  Flags
timeStamp
#ShRec
-----
1/1/3:2      SMK_BASE_CONFIG / 192.0.2.1
st1          88  -- -- -- -- -- 04/05/2019 09:46:39
0x0          0
1/1/3:2      SMK_BASE_CONFIG / 192.0.2.2
st1          88  -- -- -- -- -- 04/05/2019 09:46:39
0x0          0
1/1/3:2      SMK_GRP_IF / 192.0.2.1
st1         212  -- -- -- -- -- 04/05/2019 09:46:39
0x0          0
1/1/3:2      SMK_GRP_IF / 192.0.2.2
st1         212  -- -- -- -- -- 04/05/2019 09:46:39
```

```

0x0                                0
1/1/3:2      SMK_GRP_IF_SAP_BUCKET0 / 192.0.2.1
st1          4  -- -- -- -- 04/05/2019 09:46:39
0x0                                0

---snip---

1/1/3:2      SMK_GRP_IF_SAP_BUCKET9 / 192.0.2.2
st1          4  -- -- -- -- 04/05/2019 09:46:39
0x0                                0
1/1/3:2      SMK_SUBNET_INFO / 192.0.2.1 vRtrId 2, ifIdx 5
st1         40  -- -- -- -- 04/05/2019 09:46:39
0x0                                0
1/1/3:2      SMK_SUBNET_INFO / 192.0.2.2 vRtrId 2, ifIdx 5
st1         40  -- -- -- -- 04/05/2019 09:46:39
0x0                                0

The following totals are for:
peer ip ALL, port/lag/sdp ALL, sync-tag st1, application SRRP
Valid Entries:                26
Locally Deleted Entries:      0
Locally Deleted Alarmed Entries: 0
Pending Global Delete Entries: 0
Omcrr Alarmed Entries:        0
Omcrr Standby Entries:        0
Associated Shared Records (ALL): 0
Associated Shared Records (LD): 0
*A:BNG-2#

```

The same information can be seen in more detail by dumping the SRRP database. This output is for SRRP instance 1, and shows the detailed information for each peer. This should clearly reflect the configuration and current state of the SRRP instances. Again there are two entries (one for the local peer and the other for the remote peer) for the BASE_CONFIG, GRP_IF and SUBNET_INFO.

```

*A:BNG-2# tools dump redundancy multi-chassis srrp-sync-database instance 1
Tag Key:  sap = 1/1/3:2
Key Info: (Type/Owner)
          SMK_BASE_CONFIG / 192.0.2.1
Data Info:
  srrpId: 1   svcId: 1   svcType: VPRN
  system IP: 0xc0000201   Group interface MAC: 02:19:01:01:00:03
  Gateway MAC: 00:00:5e:00:01:01
  Subscriber interface name: sub-int-1

Tag Key:  sap = 1/1/3:2
Key Info: (Type/Owner)
          SMK_BASE_CONFIG / 192.0.2.2
Data Info:
  srrpId: 1   svcId: 1   svcType: VPRN
  system IP: 0xc0000202   Group interface MAC: 02:1b:01:01:00:03
  Gateway MAC: 00:00:5e:00:01:01
  Subscriber interface name: sub-int-1

Tag Key:  sap = 1/1/3:2
Key Info: (Type/Owner)
          SMK_GRP_IF / 192.0.2.1
Data Info:
  Group interface name: group-int-1
  Redundant interface name: bng-1-bng-2-vprn-1
  Redundant Interface IP/Mask:
    IP: 0xc0a80400 Mask: 0xffffffff
  AdminUp: Up, OperState: SRRP_STATE_BACKUP_SHUNT, InUsePriority: 100, Red-If OK:
  Yes, MessageSap OK: Yes

```

```

Tag Key: sap = 1/1/3:2
Key Info: (Type/Owner)
          SMK_GRP_IF / 192.0.2.2
Data Info:
  Group interface name: group-int-1
  Redundant interface name: bng-2-bng-1-vprn-1
  Redundant Interface IP/Mask:
    IP: 0xc0a80401 Mask: 0xfffffffffe
  AdminUp: Up, OperState: SRRP_STATE_MASTER, InUsePriority: 250, Red-If OK: Yes,
  MessageSap OK: Yes

Tag Key: sap = 1/1/3:2
Key Info: (Type/Owner)
          SMK_GRP_IF_SAP_BUCKET0 / 192.0.2.1
Data Info:

Tag Key: sap = 1/1/3:2
Key Info: (Type/Owner)
          SMK_GRP_IF_SAP_BUCKET0 / 192.0.2.2
Data Info:

---snip---

Tag Key: sap = 1/1/3:2
Key Info: (Type/Owner)
          SMK_GRP_IF_SAP_BUCKET9 / 192.0.2.1
Data Info:

Tag Key: sap = 1/1/3:2
Key Info: (Type/Owner)
          SMK_GRP_IF_SAP_BUCKET9 / 192.0.2.2
Data Info:

Tag Key: sap = 1/1/3:2
Key Info: (Type/Owner)
          SMK_SUBNET_INFO / 192.0.2.1 vRtrId 2, ifIdx 5
Data Info:
  Subscriber IP Addr: 10.2.0.1      Mask: 0xffff0000      Gateway: 10.2.0.254
  Subscriber IP Addr: 10.3.0.1      Mask: 0xffff0000      Gateway: 10.3.0.254
  Subscriber IP Addr: 10.4.0.1      Mask: 0xffff0000      Gateway: 10.4.0.254

Tag Key: sap = 1/1/3:2
Key Info: (Type/Owner)
          SMK_SUBNET_INFO / 192.0.2.2 vRtrId 2, ifIdx 5
Data Info:
  Subscriber IP Addr: 10.2.0.2      Mask: 0xffff0000      Gateway: 10.2.0.254
  Subscriber IP Addr: 10.3.0.2      Mask: 0xffff0000      Gateway: 10.3.0.254
  Subscriber IP Addr: 10.4.0.2      Mask: 0xffff0000      Gateway: 10.4.0.254

*A:BNG-2#

```

The following is an example of messages that could be seen due to this synchronization which otherwise would not be available.

An event will be generated in log 99, if the IP address was removed from the redundant interface on the remote peer.

```

*A:BNG-2# show log log-id 99 ascending

---snip---

89 2019/04/05 12:14:37.258 CEST WARNING: MC_REDUNDANCY #2012 vprn1 SRRP/MCS: Peer

```

```

Red i/f no addr
"SRRP ID 1: Redundant interface bng-1-bng-2-vprn-1 on peer 192.0.2.1 / interface
group-int-1 does not match local 192.0.2.2 / interface group-int-1."

90 2019/04/05 12:14:37.258 CEST WARNING: MC_REDUNDANCY #2012 vprn1 SRRP/MCS: Peer
Red i/f down
"SRRP ID 1: Redundant interface bng-1-bng-2-vprn-1 on peer 192.0.2.1 / interface
group-int-1 does not match local 192.0.2.2 / interface group-int-1."

*A:BNG-2#

```

The SRRP Instance Related Information

The SRRP instance information can be displayed by the following commands:

The primary BNS shows the **master** in the operation status.

```

*A:BNG-2# show srrp

=====
SRRP Table
=====
ID      Service      Group Interface      Admin      Oper
-----
1        1          group-int-1          Up         master
-----
No. of SRRP Entries: 1
=====
*A:BNG-2#

```

The backup BNG shows a *backupShunt* in the operation status.

```

*A:BNG-1# show srrp

=====
SRRP Table
=====
ID      Service      Group Interface      Admin      Oper
-----
1        1          group-int-1          Up         backupShunt
-----
No. of SRRP Entries: 1
=====
*A:BNG-1#

```

To check detailed information:

```

*A:BNG-2# show srrp 1 detail

=====
SRRP Instance 1
=====
Description      : (Not Specified)
Admin State      : Up                Oper State      : master
Preempt         : yes                One GARP per SAP : no
Monitor Oper Group : None
System IP       : 192.0.2.2
Service ID      : VPRN 1
Group If        : group-int-1      MAC Address     : 02:1b:01:01:00:03
Grp If Description : N/A

```



```

Grp If Admin State : Up          Grp If Oper State: Up
Subscriber If      : sub-int-1
Sub If Admin State : Up          Sub If Oper State: Up
Address           : 10.2.0.2/16  Gateway IP       : 10.2.0.254
Address           : 10.3.0.2/16  Gateway IP       : 10.3.0.254
Address           : 10.4.0.2/16  Gateway IP       : 10.4.0.254
Redundant If      : bng-2-bng-1-vprn-1
Red If Admin State : Up          Red If Oper State: Up
Address           : 192.168.4.1/31
Red Spoke-sdp     : 21:1
Msg Path SAP      : 1/1/3:2
Admin Gateway MAC :               Oper Gateway MAC : 00:00:5e:00:01:01
Config Priority    : 250          In-use Priority   : 250
Master Priority    : 250
Keep-alive Interval : 1 deci-seconds Master Since    : 04/05/2019 09:21:14
Fib Population Mode : all
VRRP Policy 1     : None          VRRP Policy 2    : None

-----
BFD interface
-----
Service ID        : 2
Interface Name     : bfd-1
Src IP            : 10.1.1.0
Dst IP            : 10.1.1.1
Session Oper State : connected

-----
Statistics
-----
Become Master      : 1          Master Changes   : 1
Become Bkup Routing : 1          Become Bkup Shunt: 1
Become Non-Master  : 0
Adv Sent           : 151493      Adv Received      : 3
Pri 0 Pkts Sent    : 0          Pri 0 Pkts Rcvd   : 0
Preempt Events     : 1          Preempted Events  : 0
Mesg Intvl Discards : 0         Mesg Intvl Errors : 0

=====
*A:BNG-2#

```

If this command is executed on the backup (BNG-1), an extra line appears after the keep-alive-interval showing the interval during which the receipt of no SRRP messages would cause the primary to be considered down, together with the instantaneous time to this interval expiring.

```

*A:BNG-1# show srrp 1 detail

=====
SRRP Instance 1
=====
Description       : (Not Specified)
Admin State       : Up          Oper State        : backupShunt
Preempt           : yes         One GARP per SAP  : no
Monitor Oper Group : None
System IP         : 192.0.2.1
Service ID        : VPRN 1
Group If          : group-int-1 MAC Address       : 02:19:01:01:00:03
Grp If Description : N/A
Grp If Admin State : Up          Grp If Oper State: Up
Subscriber If      : sub-int-1
Sub If Admin State : Up          Sub If Oper State: Up
Address           : 10.2.0.1/16  Gateway IP       : 10.2.0.254
Address           : 10.3.0.1/16  Gateway IP       : 10.3.0.254

```

```

Address      : 10.4.0.1/16      Gateway IP    : 10.4.0.254
Redundant If : bng-1-bng-2-vprn-1
Red If Admin State : Up        Red If Oper State: Up
Address      : 192.168.4.0/31
Red Spoke-sdp : 12:1
Msg Path SAP : 1/1/3:2
Admin Gateway MAC :              Oper Gateway MAC : 00:00:5e:00:01:01
Config Priority : 100            In-use Priority : 100
Master Priority : 250
Keep-alive Interval : 1 deci-seconds Master Since : 04/05/2019 09:21:14
Master Down Interval: 0.300 sec (Expires in 0.250 sec)
Fib Population Mode : all
VRRP Policy 1 : None            VRRP Policy 2 : None

-----
BFD interface
-----
Service ID      : 2
Interface Name   : bfd-1
Src IP          : 10.1.1.1
Dst IP          : 10.1.1.0
Session Oper State : connected

-----
Statistics
-----
Become Master      : 1            Master Changes : 2
Become Bkup Routing : 2          Become Bkup Shunt: 2
Become Non-Master  : 1
Adv Sent           : 148          Adv Received    : 151469
Pri 0 Pkts Sent    : 1            Pri 0 Pkts Rcvd : 0
Preempt Events     : 0            Preempted Events : 1
Mesg Intvl Discards : 0          Mesg Intvl Errors: 0

=====
*A:BNG-1#

```

Monitoring the Traffic on Redundant Interface

The Oper State reflects both the state of the SRRP instance and its action with respect to the redundant interface. Specifically, when the peer is SRRP primary the operational state is always primary – traffic is sent directly to the subscriber over its associated SAP. If the peer is SRRP backup and the redundant interface is Up then the Oper State will be backupShunt, if the redundant interface is down then the Oper State is *backupRouting*. In the *backupShunt* state, traffic to the subscriber is shunted (for example, forwarded) across the redundant interface to the peer (to the primary) in order to be forwarded to the subscriber.

When in the *backupRouting* state, the SRRP instance is in backup but the redundant interface is down, so the traffic is forwarded directly to the subscriber through its associated SAP.

A useful command to see the traffic on the redundant interface is:

```

*A:BNG-2# monitor service id 1 sdp 21:1 rate interval 11 repeat 3

=====
Monitor statistics for Service 1 SDP binding 21:1
=====
-----
At time t = 0 sec (Base Statistics)
-----

```

```
I. Fwd. Pkts.   : 193           I. Dro. Pkts.   : 0
I. Fwd. Octs.   : 6076          I. Dro. Octs.   : 0
E. Fwd. Pkts.   : 49           E. Fwd. Octets   : 3022

-----
At time t = 11 sec (Mode: Rate)
-----
I. Fwd. Pkts.   : 1           I. Dro. Pkts.   : 0
I. Fwd. Octs.   : 84          I. Dro. Octs.   : 0
E. Fwd. Pkts.   : 1           E. Fwd. Octets   : 98

-----
At time t = 22 sec (Mode: Rate)
-----
I. Fwd. Pkts.   : 1           I. Dro. Pkts.   : 0
I. Fwd. Octs.   : 84          I. Dro. Octs.   : 0
E. Fwd. Pkts.   : 1           E. Fwd. Octets   : 98

-----
At time t = 33 sec (Mode: Rate)
-----
I. Fwd. Pkts.   : 1           I. Dro. Pkts.   : 0
I. Fwd. Octs.   : 87          I. Dro. Octs.   : 0
E. Fwd. Pkts.   : 1           E. Fwd. Octets   : 98

=====
*A:BNB-2#
```

BFD-Related Information

To check the BFD session state.

```
*A:BNB-2# show router bfd session

=====
Legend:
  Session Id = Interface Name | LSP Name | Prefix | RSVP Sess Name | Service Id
  wp = Working path   pp = Protecting path
=====
BFD Session
=====
Session Id      State      Tx Pkts    Rx Pkts
Rem Addr/Info/SdpId:VcId  Multipl  Tx Intvl   Rx Intvl
Protocols       Type     LAG Port    LAG ID
-----
bfd-1           Up        143677     143677
10.1.1.1        3         100        100
srrp            iom       N/A        N/A
-----
No. of BFD sessions: 1
=====
*A:BNB-2#
```

To check the MAC addresses of the SRRP, this is can be done by checking the MACs table in the BSN.

```
*A:DSLAM# show service fdb-mac

=====
Service Forwarding Database
=====
```

ServId	MAC Transport:Tnl-Id	Source-Identifier	Type Age	Last Change
1	00:00:00:00:00:01	sap:1/2/1:1	L/0	04/05/19 14:53:54
1	00:00:5e:00:01:01	sap:1/1/2:1	L/0	04/05/19 14:53:51
1	02:19:01:01:00:03	sap:1/1/1:1	L/0	04/05/19 14:53:55
1	02:1b:01:01:00:03	sap:1/1/2:1	L/0	04/05/19 14:57:53
2	00:00:5e:00:01:01	sap:1/1/2:2	L/0	04/05/19 14:53:51

No. of Entries: 5				

Legend: L=Learned O=0am P=Protected-MAC C=Conditional S=Static Lf=Leaf				
=====				
*A:DSLAM#				

SRRP Debug Commands

There are debug command to show the SRRP protocol events and packets.

```
*A:BNG-2# debug router 1 srrp events
*A:BNG-2# debug router 1 srrp packets
```

To display the debugging information, a dedicated log should be created:

```
configure
log
    log-id 1
        from debug-trace
        to memory 3000
        no shutdown
    exit
exit
exit
exit
```

The following output displays a sample SRRP debug log:

```
2558 2019/04/05 13:57:55.104 CEST MINOR: DEBUG #2001 vprn1 SRRP
"SRRP: Receiving Pkt

Version (SRRP)      : 8
Type                : Advertisement (1)
Vr Id               : 1
Priority             : 250
Count Ip Addresses  : 3
Advertise Interval  : 10 centi-second
Checksum            : 0x21ef

Raw Pkt:

81 00 fe 00 00 00 00 01 00 00 5e 00 01 01 00 0a
00 03 21 ef "
```

An example SRRP message captured with tshark looks as follows:

```
Frame 1 (60 bytes on wire, 60 bytes captured)
  Arrival Time: Apr  8, 2019 09:20:52.989717000
    [Time delta from previous captured frame: 0.000000000 seconds]
    [Time delta from previous displayed frame: 0.000000000 seconds]
    [Time since reference or first frame: 0.000000000 seconds]
```

```

Frame Number: 1
Frame Length: 60 bytes
Capture Length: 60 bytes
[Frame is marked: False]
[Protocols in frame: eth:vlan:ip:vrrp]
Ethernet II, Src: IETF-VRRP-virtual-router-VRID_01 (00:00:5e:00:01:01),
                Dst: IPv4mcast_00:00:12 (01:00:5e:00:00:12)
Destination: IPv4mcast_00:00:12 (01:00:5e:00:00:12)
Address: IPv4mcast_00:00:12 (01:00:5e:00:00:12)
.... 1 .... = IG bit: Group address
                (multicast/broadcast)
.... 0 .... = LG bit: Globally unique address
                (factory default)
Source: IETF-VRRP-virtual-router-VRID_01 (00:00:5e:00:01:01)
Address: IETF-VRRP-virtual-router-VRID_01 (00:00:5e:00:01:01)
.... 0 .... = IG bit: Individual address (unicast)
.... 0 .... = LG bit: Globally unique address
                (factory default)
Type: 802.1Q Virtual LAN (0x8100)
802.1Q Virtual LAN, PRI: 0, CFI: 0, ID: 2
000. .... = Priority: 0
...0 .... = CFI: 0
.... 0000 0000 0010 = ID: 2
Type: IP (0x0800)
Trailer: 0000
Internet Protocol, Src: 192.0.2.2 (192.0.2.2), Dst: 224.0.0.18 (224.0.0.18)
Version: 4
Header length: 20 bytes
Differentiated Services Field: 0xc0 (DSCP 0x30: Class Selector 6; ECN: 0x00)
1100 00.. = Differentiated Services Codepoint: Class Selector 6 (0x30)
.... 0.. = ECN-Capable Transport (ECT): 0
.... 0.. = ECN-CE: 0
Total Length: 40
---snip---
Protocol: VRRP (0x70)
Header checksum: 0x1f66 [correct]
[Good: True]
[Bad : False]
Source: 192.0.2.2 (192.0.2.2)
Destination: 224.0.0.18 (224.0.0.18)
Virtual Router Redundancy Protocol
Version 8, Packet type 1 (Advertisement)
1000 .... = VRRP protocol version: 8
.... 0001 = VRRP packet type: Advertisement (1)
Virtual Rtr ID: 0
Priority: 250 (Non-default backup priority)
Count IP Adrs: 0
Auth Type: No Authentication (0)
Adver Int: 0
Checksum: 0x0001 [correct]

```

As an example, in the following output BNG-2 is the SRRP primary for instance 1 and sends SRRP advertisement messages. Then BNG-2 receives an SRRP message with a higher priority (254) from its peer BNG-1. This causes an event *Become Pending-Backup Shunt* where BNG-2 prepares to transition to the backup state. To achieve this, BNG-2 sends an SRRP message with priority 0. If BNG-2 continues to receive priority 254 SRRP messages from its peer BNG-1, it passes into the backup state with the event *Become Backup Shunt*.

```

3267 2019/04/08 08:57:13.388 CEST MINOR: DEBUG #2001 vprn1 SRRP
"SRRP: Sending Pkt

Version (SRRP)      : 8

```

```
Type                : Advertisement (1)
Vr Id               : 1
Priority            : 250
Count Ip Addresses  : 3
Advertise Interval  : 10 centi-second
Checksum           : 0x25ef

Raw Pkt:

81 00 fa 00 00 00 00 01 00 00 5e 00 01 01 00 0a
00 03 25 ef "

3268 2019/04/08 08:57:13.390 CEST MINOR: DEBUG #2001 vprn1 SRRP
"SRRP: Receiving Pkt

Version (SRRP)      : 8
Type                : Advertisement (1)
Vr Id               : 1
Priority            : 254
Count Ip Addresses  : 3
Advertise Interval  : 10 centi-second
Checksum           : 0x21ee

Raw Pkt:

81 01 fe 00 00 00 00 01 00 00 5e 00 01 01 00 0a
00 03 21 ee "

3269 2019/04/08 08:57:13.390 CEST MINOR: DEBUG #2001 vprn1 SRRP
"SRRP: Event
Become Pending-Backup Shunt: vRtrId 2, ifIdx 6, IPv4 vr_id 1, Master IP 192.0.2.
1"

3270 2019/04/08 08:57:13.390 CEST MINOR: DEBUG #2001 vprn1 SRRP
"SRRP: Sending Pkt

Version (SRRP)      : 8
Type                : Advertisement (1)
Vr Id               : 1
Priority           : 0
Count Ip Addresses  : 3
Advertise Interval  : 10 centi-second
Checksum           : 0x1ff0

Raw Pkt:

81 00 00 00 00 00 00 01 00 00 5e 00 01 01 00 0a
00 03 1f f0 "

3271 2019/04/08 08:57:13.392 CEST MINOR: DEBUG #2001 vprn1 SRRP
"SRRP: Receiving Pkt

Version (SRRP)      : 8
Type                : Advertisement (1)
Vr Id               : 1
Priority            : 254
Count Ip Addresses  : 3
Advertise Interval  : 10 centi-second
Checksum           : 0x21ef

---snip---
```

SRRP Traffic Marking

The SRRP messages are sent by default with DSCP of *nc1* and with 802.1p bits of 0, as can be seen in the following tshark snippet.

```
802.1Q Virtual LAN, PRI: 0, CFI: 0, ID: 2
 000. .... = Priority: 0
...0 .... = CFI: 0
.... 0000 0000 0010 = ID: 2
Type: IP (0x0800)
Trailer: 0000
Internet Protocol, Src: 192.0.2.2 (192.0.2.2), Dst: 224.0.0.18 (224.0.0.18)
Version: 4
Header length: 20 bytes
Differentiated Services Field: 0xc0 (DSCP 0x30: Class Selector 6; ECN: 0x00)
 1100 00.. = Differentiated Services Codepoint: Class Selector 6 (0x30)
.... ..0. = ECN-Capable Transport (ECT): 0
.... ...0 = ECN-CE: 0
Total Length: 40
```

```
*A:BNG-2# show qos dscp-table
```

DSCP Mapping

DSCP Name	DSCP Value	TOS (bin)	TOS (hex)
be	0	0000 0000	00
nc1	48	1100 0000	C0
cp63	63	1111 1100	FC

```
*A:BNG-2#
```

Where DSCP 0x30=48 (DSCP value). This can be changed to EF, for example, using the following command:

```
*A:BNG-2# configure service vprn 1 sgt-qos application srrp dscp ef
```

Conclusion

This chapter provides configuration and troubleshooting commands for SRRP with static (IP-MAC) host in a Layer 3 Routed-CO (IES/VP RN subscriber interface) context.

Virtual Residential Gateway Authentication Scenarios

This chapter describes virtual residential gateway authentication scenarios.

Topics in this chapter include:

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

Applicability

The information and configuration in this chapter are based on SR OS Release 14.0.R3.

Overview

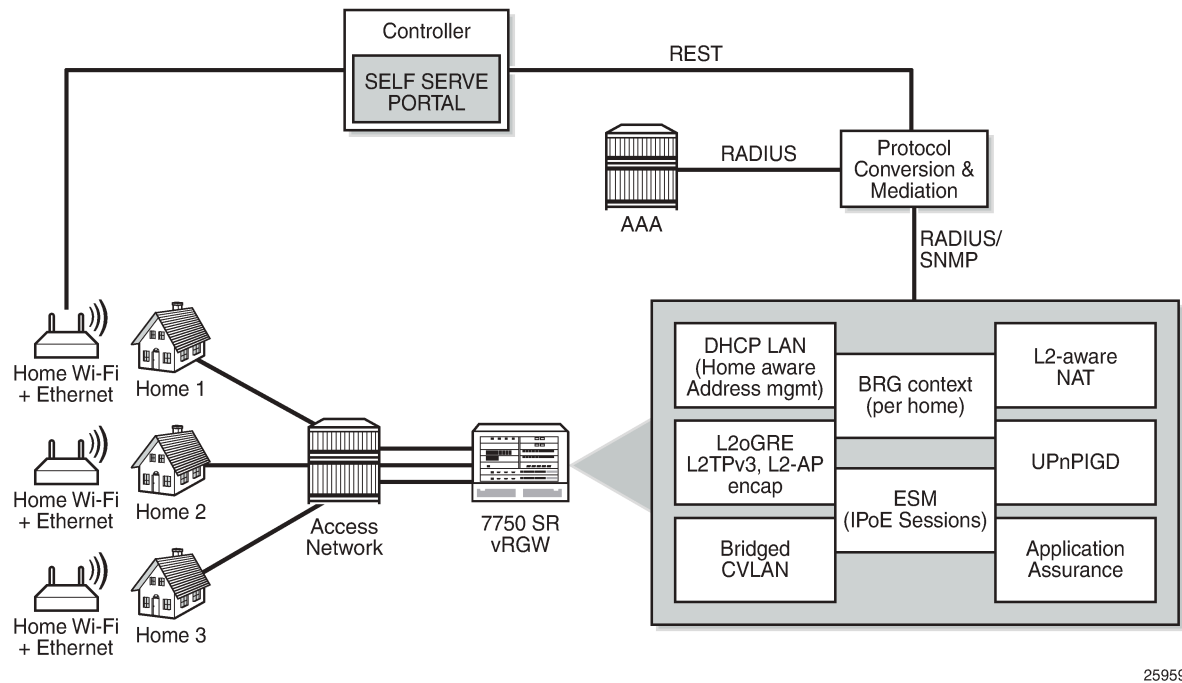
In the virtual residential gateway (vRGW) model, the Layer 3 (L3) functions are moved out of the traditional residential gateway (RGW) and into the network. The [Virtual Residential Gateway Home Pool Management](#) chapter provides the rationale for this scenario, and describes how services must be configured for the service router to support this model.

The home network can be self-managed through a service portal, where end users can connect and change home-specific settings; see [Figure 160: BRG and home device management](#). The portal logic is implemented in a controlling entity providing a RESTful interface. A protocol conversion and mediation platform (PCMP) is needed to translate the RESTful interface into RADIUS (and SNMP), and vice versa. The PCMP operates in conjunction with the 5620 SAM. In the remainder of this chapter, PCMP and the controller are represented as a single component.

Managing individual BRGs requires the vRGW to maintain a context for every BRG, so each BRG is identified through the BRG ID. This context is created when authenticating the BRG, and stores the home-level settings. The BRG context is deleted when the BRG is not deemed alive anymore.

The BRG ID can be derived from any of the parameters in the RADIUS Access-Request message, such as the SAP, tunnel-source, or called-station ID by a controlling entity; see the [Implicit authentication](#) and [Explicit authentication](#) sections that follow in this chapter. Typically, the MAC address of the BRG serves as the BRG ID, and that is what is used throughout this chapter.

Figure 160: BRG and home device management



25959

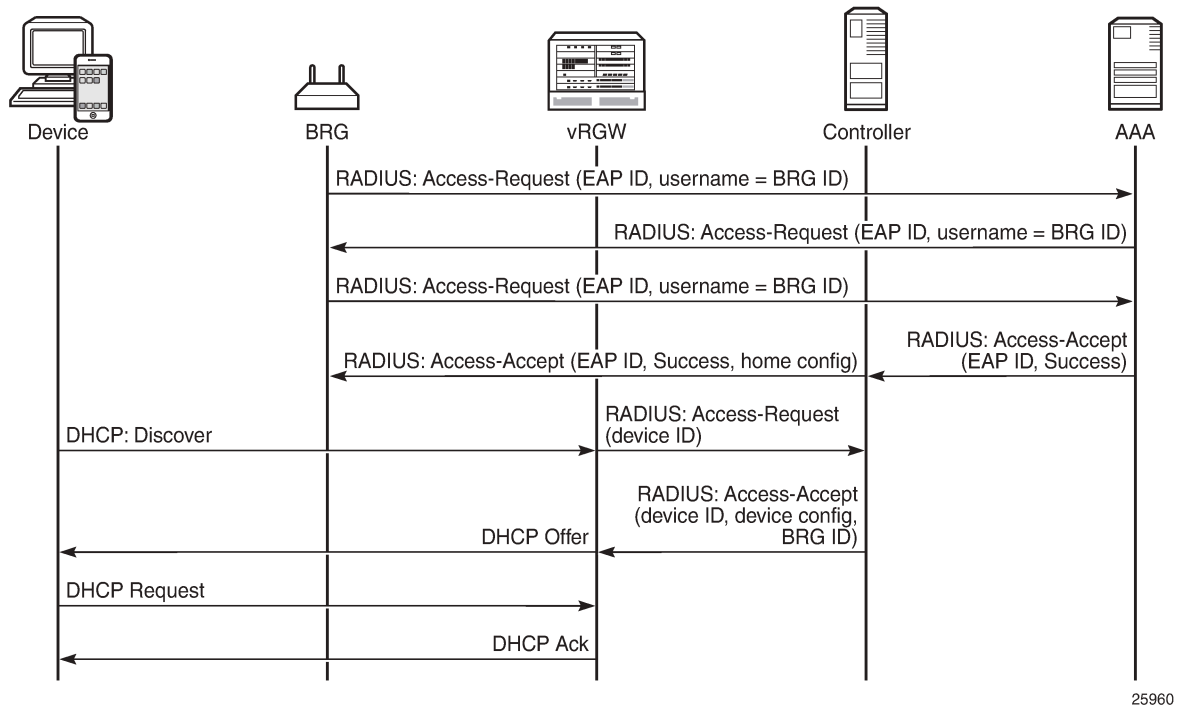
The vRGW supports two scenarios for authenticating bridged gateways and their hosts:

- Explicit authentication
- Implicit authentication

Explicit authentication

Two main phases are distinguished in the explicit authentication scenario; see [Figure 161: Explicit BRG authentication](#). The first phase is the BRG authentication phase, the second phase is the device authentication phase. The example in [Figure 161: Explicit BRG authentication](#) uses IPv4, but also works with IPv6.

Figure 161: Explicit BRG authentication



The first message in the first phase is an Access-Request message sent by the BRG toward the AAA/RADIUS server, and uses the extensible authentication protocol (EAP). This message is proxied by the vRGW as well as by the controller to the AAA server, and the BRG ID is used as the username. The last message of the first phase is the Access-Accept message. When the controller receives this message from the AAA server, it fetches and adds the per-home configuration parameters to the Access-Accept message before forwarding this message to the vRGW.

The second phase starts with the Discover message of a typical Discover-Offer-Request-Ack (DORA) message sequence. The vRGW then initiates device authentication toward the controller, which returns the BRG ID and, optionally, device-specific configuration data in an Access-Accept message. The controller usually will not proxy device authentication toward the AAA server. Typically, the RADIUS protocol is used between the vRGW and the controller. As such, the vRGW will send an Access-Request message to the controller for every new device, including static devices, to get the per-device configuration.

Usually, the vRGW combines the home-specific data with the device-specific data, where the more specific device data overrules the home-specific data. The combined data is then used to create the corresponding ESM hosts and IPE sessions.


If the home-specific data includes Allocated-Addresses attributes defining static devices, these devices are authenticated automatically as soon as at least one dynamic device is connected.

The explicit authentication scenario requires the BRG profile to include a RADIUS proxy server; see the Configuration section in this chapter for a practical example.

Implicit authentication

In the implicit authentication scenario, the BRG is authenticated when the first device connects to the vRGW; see [Figure 162: Implicit BRG authentication](#) for an example.

Figure 162: Implicit BRG authentication



In [Figure 162: Implicit BRG authentication](#), the Discover message triggers the vRGW to send an Access-Request message toward the controller. The controller returns device-specific data including the BRG ID. Because there is no context for this BRG yet, the vRGW starts BRG authentication toward the AAA/RADIUS server. The controller proxies this message, and on return adds the home-specific data to the Access-Accept message. The overall result is that the vRGW now knows the home- and device-specific data.

As with the explicit authentication scenario, the combined data is then used to create the corresponding ESM hosts and IPE sessions.

No separate BRG authentication is required when subsequent devices connect and device authentication returns a BRG ID that is already known to the vRGW.

As with the explicit authentication scenario, if the home-specific data includes Alc-Reserved-Addresses attributes defining static home devices, the static devices are authenticated automatically when at least one dynamic device is connected.

The implicit authentication scenario requires the BRG profile to include a RADIUS authentication context defining a RADIUS server policy and a password; see the [Configuration](#) section in this chapter for an example.

Connectivity verification and BRG deletion

For the purpose of clearing resources when these resources are not needed anymore, the vRGW performs the BRG connectivity verification and deletion process. When BRG connectivity is considered lost, the BRG and its hosts are deleted automatically.

Parameters controlling the BRG connectivity verification and deletion process are located in the BRG profile context:

```
*A:BNG>config>subscr-mgmt>brg-profile#
connectivity-verification count 3 timeout 30 retry-time 900
  count <nr-of-attempts>      : [1..5]      - default: 3
  timeout <timeout-seconds>   : [5..60]     - default: 30
  retry-time <retry-seconds>  : [300..3600] - default: 300
  hold-time <seconds> [30..86400] - default: no hold-time
  initial-hold-time <seconds> [0..900]    - default: 300
```

When the last dynamic host associated with a BRG is deleted, while at the same time connectivity-verification is enabled, the vRGW starts a liveness test toward the BRG through ICMP (v4 or v6) messages. These messages are sent to either the BRG tunnel source IP address or the BRG RADIUS source IP address. If the BRG has neither of these addresses (for example, because each BRG is managed through a unique VLAN in the implicit authentication scenario), connectivity verification is not executed and only the hold-timer applies. If no answer is returned by the BRG in time (timeout), the vRGW considers the BRG in a failed state, and tries again (retry-time) for some maximum number of times (count).

The vRGW starts the hold-timer when the maximum number of tries is reached, or when connectivity verification is disabled (no connectivity verification). When the hold-timer expires, the BRG context is deleted together with the associated hosts.

The initial-hold-timer is required in the scenario where operators want to use explicit authentication without connectivity-verification, and with no hold-timer defined. In this scenario, defining a non-zero initial-hold-timer value avoids BRG contexts from being deleted immediately after their creation.

If a new dynamic host connects while executing the liveness test or while the hold-timer or the initial-hold-timer is running, the connectivity verification and BRG deletion process is canceled.

The hold-timer is ignored when manually clearing BRGs and related hosts with the following command.

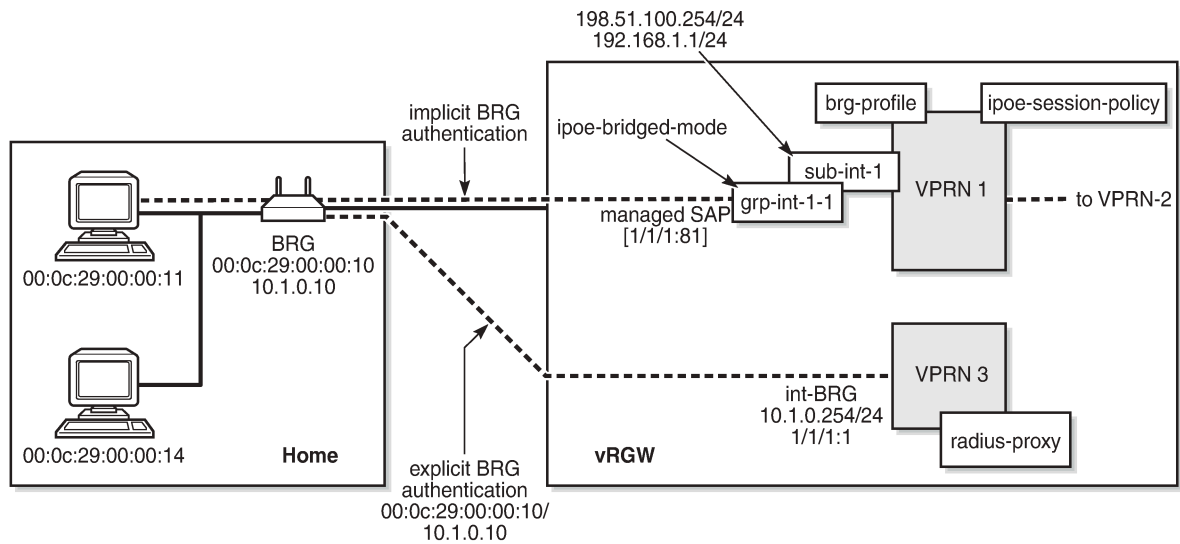
```
clear subscriber-mgmt brg gateway
  brg-id <brg-ident>
  host <ieee-address>
  all-hosts
  idle-bindings
  all-gateways
```

Configuration

The services configuration shown in [Figure 163: Example service configuration for explicit and implicit BRG authentication](#) applies to the examples throughout this chapter. Because the main focus of this chapter is on BRG and home device authentication, the detailed configuration of VPLS 10 containing the capture SAP and of VPRN-2 containing the outside L2-aware NAT range is not repeated here. See the [Virtual Residential Gateway Home Pool Management](#) chapter for those configurations.

In summary, VPRN-1 provides the connection toward the customer premises, and hosts the NAT inside addresses. VPRN-2 provides the connection toward the Internet, and hosts the NAT outside addresses. Also, VPRN-3 provides connection to the management interfaces of the BRGs, and is used for the explicit BRG authentication scenario.

Figure 163: Example service configuration for explicit and implicit BRG authentication



25962

Service configuration

An excerpt of the configuration of VPRN-1 follows. The group interface *grp-int-1-1* has RADIUS authentication enabled through authentication policy *radius-AUTH*, and BRG authentication through the default BRG profile *brg-prof-1*.

```
configure
service
  vprn 1 customer 1 create
  ---snip---
  subscriber-interface "sub-int-1" create
    address 198.51.100.254/24
    address 192.168.1.1/24
    ---snip---
  group-interface "grp-int-1-1" create
    ipv6
    ---snip---
    ipoe-bridged-mode
  exit
  ---snip---
  authentication-policy "radius-AUTH"
  ipoe-session
    ipoe-session-policy "sess-pol-SAP-MAC"
    sap-session-limit 128
    no shutdown
  exit
  brg
    default-brg-profile "brg-prof-1"
    no shutdown
```

```

                                exit
                                oper-up-while-empty
                            exit
                        exit
                    nat
                    inside
                    l2-aware
                    address 192.168.0.1/16
                    exit
                exit
            exit
        no shutdown
    exit
exit
exit

```

BRG profile

The BRG profile *brg-prof-1* is defined in the subscriber management context and provides an SLA profile, a subscriber profile, a DHCP pool, a RADIUS server policy plus the corresponding password, and a RADIUS proxy server.

For explicit BRG authentication, the RADIUS proxy server is used; for implicit BRG authentication, the RADIUS server policy and password defined in the RADIUS authentication context are used.

```

configure
  subscriber-mgmt
    ---snip---
    brg-profile "brg-prof-1" create
      description "default BRG-profile, demo purposes"
      sla-profile-string "sla-prof-1"
      sub-profile-string "sub-prof-1"
      dhcp-pool
        subnet 192.168.1.1/24 start 192.168.1.2 end 192.168.1.254
      exit
      radius-authentication
        password letmein
        radius-server-policy "rad-serv-pol-RSP"
      exit
      radius-proxy-server router 3 name "rad-prox-RPROX"
    exit
  exit
exit

```

RADIUS proxy configuration

VPRN-3 is defined for supporting explicit BRG authentication via a RADIUS proxy. The *int-BRG* interface is on SAP 1/1/1:1, and provides connectivity to the management interface of the physical BRGs. The RADIUS proxy listens on the *int-LB-PROXY* interface, and directs the incoming RADIUS messages to the server, as defined by the default authentication-server policy.

```

configure
  service
    vprn 3
      route-distinguisher 64496:3
      interface "int-LB-PROXY" create
        address 10.33.33.1/32

```

```
        loopback
    exit
    interface "int-BRG" create
        address 10.1.0.254/24
        sap 1/1/1:1 create
    exit
    exit
    radius-proxy
        server "rad-prox-RPROX" purpose authentication create
        default-authentication-server-policy "rad-serv-pol-RSP"
        interface "int-LB-PROXY"
        secret vsecret1
        no shutdown
    exit
    exit
    no shutdown
    exit
    exit
    exit
```

RADIUS policies

The RADIUS authentication and accounting policies are defined as follows, so authentication and accounting happens via the base router instance.

```
configure
    router
        radius-server
            server "radius-172.16.1.2" address 172.16.1.2 secret vsecret1 create
            accept-coa
        exit
    exit
    exit
    exit
    configure
        aaa
            radius-server-policy "rad-serv-pol-RSP" create
            servers
                router "Base"
                source-address 192.0.2.1
                server 1 name "radius-172.16.1.2"
            exit
        exit
    exit
    exit
    exit
    configure
        subscriber-mgmt
            authentication-policy "radius-AUTH" create
            description "RADIUS authentication policy"
            password letmein
            radius-server-policy "rad-serv-pol-RSP"
        exit
        radius-accounting-policy "radius-ACCT" create
        update-interval 5
        include-radius-attribute
            mac-address
            nat-port-range
            subscriber-id
        exit
```

```

        radius-accounting-server
        source-address 192.0.2.1
        router "Base"
        server 1 address 172.16.1.2 secret vsecret1
    exit
exit
exit
exit

```

RADIUS user configuration

Although a PCMP will be used in conjunction with an external controller, for demonstration purposes, this chapter relies on a RADIUS server only.

A sample RADIUS user configuration follows. MAC addresses are used for authentication. MAC address 00:0c:29:00:00:10 identifies the BRG. The addresses ranging from 00:0c:29:00:00:11 to 00:0c:29:00:00:1f identify the home devices connected to that BRG so they all return the same Alc-BRG-Id.

```

00:0c:29:00:00:10    Cleartext-Password := "letmein"
                    Alc-BRG-Id = "00:0c:29:00:00:10",
                    Framed-IPv6-Prefix = 2001:db8:101:1010::/64,
                    Alc-DMZ-address = 192.168.1.254,
                    Alc-Home-Aware-Pool =
                        "192.168.1.1/24 192.168.1.100-192.168.1.254",
                    Alc-Reserved-Addresses =
                        "sticky 00:0c:29:00:00:11 192.168.1.110",
                    Alc-Reserved-Addresses +=
                        "static 00:0c:29:00:00:1f 192.168.1.254",
                    Alc-Reserved-Addresses +=
                        "static 00:0c:29:00:00:1e 198.51.100.110",
                    Alc-Portal-Url = "http://11.11.11.11",
                    Alc-Primary-Dns = 1.1.1.1,
                    Alc-Secondary-Dns = 1.1.2.2,
                    Alc-Primary-Nbns = 2.2.1.1,
                    Alc-Secondary-Nbns = 2.2.2.2,
                    Alc-IPv6-Primary-DNS = 2001:db8:dddd:1::1,
                    Alc-IPv6-Secondary-DNS = 2001:db8:dddd:2::1,

00:0c:29:00:00:11    Cleartext-Password := "letmein"
                    Alc-BRG-Id = "00:0c:29:00:00:10",

00:0c:29:00:00:12    Cleartext-Password := "letmein"
                    Alc-BRG-Id = "00:0c:29:00:00:10",

00:0c:29:00:00:13    Cleartext-Password := "letmein"
                    Alc-BRG-Id = "00:0c:29:00:00:10",

00:0c:29:00:00:14    Cleartext-Password := "letmein"
                    Alc-BRG-Id = "00:0c:29:00:00:10",
                    Alc-Primary-Dns = 1.1.3.3,

```

Debug configuration

The following debug configuration can be used for troubleshooting purposes.

```

debug
router "Base"

```



```

        radius
            packet-type authentication accounting coa
            detail-level high
        exit
    exit
    router "1"
        ip
            dhcp
                detail-level medium
                mode egr-ingr-and-dropped
            exit
            icmp6
        exit
    exit
    router "3"
        radius-proxy
            server "rad-prox-RPROX"
            detail-level high
            direction both
            packet-type access-request access-accept access-reject
                    access-challenge accounting-request
                    accounting-response other
        exit
    exit
exit
exit
exit

```

Explicit authentication

In the explicit authentication scenario, the BRG is authenticated before any home device attempts to connect. The following trace shows that the RADIUS proxy server in router 3 receives an Access-Request from BRG with BRG ID 00:0c:29:00:00:10, and passes this to the AAA/RADIUS server, which returns the BRG specific data. Two static addresses and one sticky address are associated with this BRG.

```

1 2016/09/20 16:27:22.11 CEST MINOR: DEBUG #2001 vprn3 RADIUS
"RADIUS: Receive
  Proxy-server rad-prox-RPROX"

2 2016/09/20 16:27:22.11 CEST MINOR: DEBUG #2001 vprn3 RADIUS
"RADIUS: Receive
  Access-Request(1) id 18 len 81 from 10.1.0.10:49169 vrid 3 pol rad-serv-pol-RSP
    USER NAME [1] 17 00:0c:29:00:00:10
    PASSWORD [2] 16 bD5mfBsZr5M/aFg0a7iAtE
    NAS IP ADDRESS [4] 4 10.1.0.10
    MESSAGE AUTHENTICATOR [80] 16 0x6c9a951328e303920ba50fa8f7eea0c3
"

3 2016/09/20 16:27:22.11 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Access-Request(1) 172.16.1.2:1812 id 1 len 81 vrid 1 pol rad-serv-pol-RSP
    USER NAME [1] 17 00:0c:29:00:00:10
    PASSWORD [2] 16 m9Lo8f7y.V35pw3KzGBs.U
    NAS IP ADDRESS [4] 4 10.1.0.10
    MESSAGE AUTHENTICATOR [80] 16 0xbe54c468e3e6952f0cedfadd3477683b

  Hex Packet Dump:
  ---snip---
"

4 2016/09/20 16:27:22.11 CEST MINOR: DEBUG #2001 Base RADIUS

```

```
"RADIUS: Receive
Access-Accept(2) id 1 len 388 from 172.16.1.2:1812 vrid 1 pol rad-serv-pol-RSP
VSA [26] 19 Alcatel(6527)
  BRG ID [225] 17 00:0c:29:00:00:10
FRAMED IPV6 PREFIX [97] 18 2001:db8:101:1010::/64
VSA [26] 6 Alcatel(6527)
  BRG DMZ ADDRESS [221] 4 192.168.1.254
VSA [26] 44 Alcatel(6527)
  BRG HOME AWARE POOL [220] 42 192.168.1.1/24 192.168.1.100-192.168.1.254
VSA [26] 40 Alcatel(6527)
  BRG RESERVED ADDRESS [223] 38 sticky 00:0c:29:00:00:11 192.168.1.110
VSA [26] 40 Alcatel(6527)
  BRG RESERVED ADDRESS [223] 38 static 00:0c:29:00:00:1f 192.168.1.254
VSA [26] 41 Alcatel(6527)
  BRG RESERVED ADDRESS [223] 39 static 00:0c:29:00:00:1e 198.51.100.110
VSA [26] 20 Alcatel(6527)
  PORTAL URL [177] 18 http://11.11.11.11
VSA [26] 6 Alcatel(6527)
  PRIMARY DNS [9] 4 1.1.1.1
VSA [26] 6 Alcatel(6527)
  SECONDARY DNS [10] 4 1.1.2.2
VSA [26] 6 Alcatel(6527)
  PRIMARY NBNS [29] 4 2.2.1.1
VSA [26] 6 Alcatel(6527)
  SECONDARY NBNS [30] 4 2.2.2.2
VSA [26] 18 Alcatel(6527)
  IPV6 PRIMARY DNS [105] 16 2001:db8:dddd:1::1
VSA [26] 18 Alcatel(6527)
  IPV6 SECONDARY DNS [106] 16 2001:db8:dddd:2::1
```

Hex Packet Dump:

---snip---

5 2016/09/20 16:27:22.12 CEST MINOR: DEBUG #2001 vprn3 RADIUS

"RADIUS: Transmit

Proxy-server rad-prox-RPROX"

6 2016/09/20 16:27:22.12 CEST MINOR: DEBUG #2001 vprn3 RADIUS

"RADIUS: Transmit

```
Access-Accept(2) 10.1.0.10:49169 id 18 len 388 vrid 3
VSA [26] 19 Alcatel(6527)
  BRG ID [225] 17 00:0c:29:00:00:10
FRAMED IPV6 PREFIX [97] 18 2001:db8:101:1010::/64
VSA [26] 6 Alcatel(6527)
  BRG DMZ ADDRESS [221] 4 192.168.1.254
VSA [26] 44 Alcatel(6527)
  BRG HOME AWARE POOL [220] 42 192.168.1.1/24 192.168.1.100-192.168.1.254
VSA [26] 40 Alcatel(6527)
  BRG RESERVED ADDRESS [223] 38 sticky 00:0c:29:00:00:11 192.168.1.110
VSA [26] 40 Alcatel(6527)
  BRG RESERVED ADDRESS [223] 38 static 00:0c:29:00:00:1f 192.168.1.254
VSA [26] 41 Alcatel(6527)
  BRG RESERVED ADDRESS [223] 39 static 00:0c:29:00:00:1e 198.51.100.110
VSA [26] 20 Alcatel(6527)
  PORTAL URL [177] 18 http://11.11.11.11
VSA [26] 6 Alcatel(6527)
  PRIMARY DNS [9] 4 1.1.1.1
VSA [26] 6 Alcatel(6527)
  SECONDARY DNS [10] 4 1.1.2.2
VSA [26] 6 Alcatel(6527)
  PRIMARY NBNS [29] 4 2.2.1.1
VSA [26] 6 Alcatel(6527)
  SECONDARY NBNS [30] 4 2.2.2.2
```

```
VSA [26] 18 Alcatel(6527)
  IPV6 PRIMARY DNS [105] 16 2001:db8:dddd:1::1
VSA [26] 18 Alcatel(6527)
  IPV6 SECONDARY DNS [106] 16 2001:db8:dddd:2::1
"
```

As a result, the vRGW creates and stores context for this BRG, which can be displayed using the following command. The Proxy authenticated flag is set to "yes".

```
*A:BNG# show subscriber-mgmt brg gateways

=====
Bridged Residential Gateways
=====
Identifier                : 00:0c:29:00:00:10
SLAAC prefix              : 2001:db8:101:1010::/64
Subnet                    : 192.168.1.1/24
Subnet start address      : 192.168.1.100
Subnet end address        : 192.168.1.254
DMZ address               : 192.168.1.254
DNS 1 v4                  : 1.1.1.1
DNS 1 v6                  : 2001:db8:dddd:1::1
DNS 2 v4                  : 1.1.2.2
DNS 2 v6                  : 2001:db8:dddd:2::1
NBNS 1                    : 2.2.1.1
NBNS 2                    : 2.2.2.2
DHCP lease time           : 21600
DHCP stream destination  : (Not Specified)
IPv4 portal URL           : http://11.11.11.11
IPv6 portal URL           : (Not Specified)
BRG profile               : brg-prof-1
Subscriber profile        : sub-prof-1
SLA profile               : sla-prof-1
UPnP policy override      : (Not Specified)
DMZ address in use        : no
Proxy authenticated       : yes
Ingress IPv4 filter override : N/A
Egress IPv4 filter override : N/A
Ingress IPv6 filter override : N/A
Egress IPv6 filter override : N/A
No QoS overrides found.
No Filter rules received.

-----
No. of gateways: 1
=====
*A:BNG#
```

Initially, no hosts are created and associated with this BRG, as the following command shows.

```
*A:BNG# show subscriber-mgmt brg gateway brg-id "00:0c:29:00:00:10" hosts
No entries found.
*A:BNG#
```

To show the static and sticky addresses associated with the BRG, use the following command. Even without any device connected to this BRG, some bindings are created.

```
*A:BNG# show subscriber-mgmt brg gateway brg-id "00:0c:29:00:00:10" bindings

=====
Bridged Residential Gateway home-aware pool address bindings
```

```
=====
Home-aware pool           : 00:0c:29:00:00:10
-----
MAC address               : 00:0c:29:00:00:11
IP address                : 192.168.1.110
Allocation type           : sticky-ip-address
DHCP lease                : false
Remaining lease time      : (Unknown)
Lease start time          : N/A

MAC address               : 00:0c:29:00:00:1e
IP address                : 198.51.100.110
Allocation type           : static
DHCP lease                : (Unknown)
Remaining lease time      : (Unknown)
Lease start time          : N/A

MAC address               : 00:0c:29:00:00:1f
IP address                : 192.168.1.254
Allocation type           : static
DHCP lease                : (Unknown)
Remaining lease time      : (Unknown)
Lease start time          : N/A

-----
No. of bindings: 3
=====
*A:BNG#
```

When the first device connects, in this example using DHCPv4 (DORA), this device is authenticated using RADIUS, and the controller returns the corresponding BRG ID and device-specific primary DNS server (messages 7 and 8). Because the BRG has two static addresses associated with it, at the same time these are also authenticated (messages 12 and 15 for the first static host, 13 and 16 for the second static host). The device with MAC address 00:0c:29:00:00:14 has a dedicated primary DNS server, which overrules the primary DNS server defined at BRG level (messages 8, 10, and 14).

```
7 2016/09/20 16:29:08.53 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Access-Request(1) 172.16.1.2:1812 id 2 len 79 vrid 1 pol rad-serv-pol-RSP
    USER NAME [1] 17 00:0c:29:00:00:14
    PASSWORD [2] 16 Lh/0pV5SVw5Cp7gBJe75s.
    NAS IP ADDRESS [4] 4 192.0.2.1
    NAS PORT TYPE [61] 4 Ethernet(15)
    NAS PORT ID [87] 8 1/1/1:81

  Hex Packet Dump:
  ---snip---
"
```

```
8 2016/09/20 16:29:08.53 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Access-Accept(2) id 2 len 57 from 172.16.1.2:1812 vrid 1 pol rad-serv-pol-RSP
    VSA [26] 19 Alcatel(6527)
    BRG ID [225] 17 00:0c:29:00:00:10
    VSA [26] 6 Alcatel(6527)
    PRIMARY DNS [9] 4 1.1.3.3

  Hex Packet Dump:
  ---snip---
"
```

```
9 2016/09/20 16:29:08.53 CEST MINOR: DEBUG #2001 vprn1 PIP
```

```
"PIP: DHCP
instance 2 (1), interface index 6 (grp-int-1-1),
  received DHCP Boot Request on Interface grp-int-1-1 (1/1/1:81) Port 67

H/W Type: Ethernet(10Mb)  H/W Address Length: 6
ciaddr: 0.0.0.0            yiaddr: 0.0.0.0
siaddr: 0.0.0.0            giaddr: 0.0.0.0
chaddr: 00:0c:29:00:00:14  xid: 0x2

DHCP options:
[53] Message type: Discover
[255] End
"

10 2016/09/20 16:29:08.53 CEST MINOR: DEBUG #2001 vprn1 PIP
"PIP: DHCP
instance 2 (1), interface index 6 (grp-int-1-1),
  transmitted DHCP Boot Reply to Interface grp-int-1-1 (1/1/1:81) Port 68

H/W Type: Ethernet(10Mb)  H/W Address Length: 6
ciaddr: 0.0.0.0            yiaddr: 192.168.1.100
siaddr: 192.168.1.1        giaddr: 192.168.1.1
chaddr: 00:0c:29:00:00:14  xid: 0x2

DHCP options:
[53] Message type: Offer
[54] DHCP server addr: 192.168.1.1
[51] Lease time: 21600
[1] Subnet mask: 255.255.255.0
[3] Router: 192.168.1.1
[6] Domain name server: length = 8
    1.1.3.3
    1.1.2.2
[44] NETBIOS name server: length = 8
    2.2.1.1
    2.2.2.2
[255] End
"

11 2016/09/20 16:29:08.64 CEST MINOR: DEBUG #2001 vprn1 PIP
"PIP: DHCP
instance 2 (1), interface index 6 (grp-int-1-1),
  received DHCP Boot Request on Interface grp-int-1-1 (1/1/1:81) Port 67

H/W Type: Ethernet(10Mb)  H/W Address Length: 6
ciaddr: 0.0.0.0            yiaddr: 0.0.0.0
siaddr: 0.0.0.0            giaddr: 0.0.0.0
chaddr: 00:0c:29:00:00:14  xid: 0x2

DHCP options:
[53] Message type: Request
[50] Requested IP addr: 192.168.1.100
[54] DHCP server addr: 192.168.1.1
[255] End
"

12 2016/09/20 16:29:08.64 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
Access-Request(1) 172.16.1.2:1812 id 3 len 79 vrid 1 pol rad-serv-pol-RSP
  USER NAME [1] 17 00:0c:29:00:00:1e
  PASSWORD [2] 16 hW.TR6SdCXxM0/3iZ.3WNk
  NAS IP ADDRESS [4] 4 192.0.2.1
  NAS PORT TYPE [61] 4 Ethernet(15)
  NAS PORT ID [87] 8 1/1/1:81
```

```
Hex Packet Dump:
---snip---
"

13 2016/09/20 16:29:08.64 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
Access-Request(1) 172.16.1.2:1812 id 4 len 79 vrid 1 pol rad-serv-pol-RSP
  USER NAME [1] 17 00:0c:29:00:00:1f
  PASSWORD [2] 16 vXTCIxeAZzeGRFAQy8eS/k
  NAS IP ADDRESS [4] 4 192.0.2.1
  NAS PORT TYPE [61] 4 Ethernet(15)
  NAS PORT ID [87] 8 1/1/1:81

Hex Packet Dump:
---snip---
"

14 2016/09/20 16:29:08.64 CEST MINOR: DEBUG #2001 vprn1 PIP
"PIP: DHCP
instance 2 (1), interface index 6 (grp-int-1-1),
  transmitted DHCP Boot Reply to Interface grp-int-1-1 (1/1/1:81) Port 68

H/W Type: Ethernet(10Mb) H/W Address Length: 6
ciaddr: 0.0.0.0 yiaddr: 192.168.1.100
siaddr: 192.168.1.1 giaddr: 192.168.1.1
chaddr: 00:0c:29:00:00:14 xid: 0x2

DHCP options:
[53] Message type: Ack
[54] DHCP server addr: 192.168.1.1
[51] Lease time: 21600
[1] Subnet mask: 255.255.255.0
[3] Router: 192.168.1.1
[6] Domain name server: length = 8
    1.1.3.3
    1.1.2.2
[44] NETBIOS name server: length = 8
    2.2.1.1
    2.2.2.2
[255] End
"

15 2016/09/20 16:29:08.63 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
Access-Accept(2) id 3 len 45 from 172.16.1.2:1812 vrid 1 pol rad-serv-pol-RSP
  VSA [26] 19 Alcatel(6527)
  BRG ID [225] 17 00:0c:29:00:00:10

Hex Packet Dump:
---snip---
"

16 2016/09/20 16:29:08.63 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
Access-Accept(2) id 4 len 45 from 172.16.1.2:1812 vrid 1 pol rad-serv-pol-RSP
  VSA [26] 19 Alcatel(6527)
  BRG ID [225] 17 00:0c:29:00:00:10

Hex Packet Dump:
---snip---
"
```

Displaying the bindings again shows that there is now an additional dynamic host, for which the allocation type is dynamic.

```
*A:BNG# show subscriber-mgmt brg gateway brg-id "00:0c:29:00:00:10" bindings

=====
Bridged Residential Gateway home-aware pool address bindings
=====
Home-aware pool      : 00:0c:29:00:00:10
-----
MAC address          : 00:0c:29:00:00:11
IP address            : 192.168.1.110
Allocation type       : sticky-ip-address
DHCP lease            : false
Remaining lease time  : (Unknown)
Lease start time      : N/A

MAC address          : 00:0c:29:00:00:14
IP address            : 192.168.1.100
Allocation type       : dynamic
DHCP lease            : true
Remaining lease time  : 21580
Lease start time      : 2016/09/20 16:29:08

MAC address          : 00:0c:29:00:00:1e
IP address            : 198.51.100.110
Allocation type       : static
DHCP lease            : (Unknown)
Remaining lease time  : (Unknown)
Lease start time      : N/A

MAC address          : 00:0c:29:00:00:1f
IP address            : 192.168.1.254
Allocation type       : static
DHCP lease            : (Unknown)
Remaining lease time  : (Unknown)
Lease start time      : N/A

-----
No. of bindings: 4
=====
*A:BNG#
```

The following command shows the hosts associated with this BRG. The sticky address is not in the list because the host is not online. For the sticky address to appear, that device must send a DHCPv4 Discover message, initiating its own authentication.

```
*A:BNG# show subscriber-mgmt brg gateway brg-id "00:0c:29:00:00:10" hosts

=====
Bridged Residential Gateway hosts
=====
Identifier           : 00:0c:29:00:00:10
MAC address          : 00:0c:29:00:00:14
IP address            : 192.168.1.100
Service              : 1 (VPRN)
Allocation type       : dynamic
Home-aware pool      : 00:0c:29:00:00:10
DHCP lease            : true
Remaining lease time  : 21567
Lease start time      : 2016/09/20 16:29:08
```

```
Identifier           : 00:0c:29:00:00:10
MAC address         : 00:0c:29:00:00:1e
IP address          : 198.51.100.110
Service             : 1 (VPRN)
Allocation type      : static
Home-aware pool     : 00:0c:29:00:00:10
DHCP lease          : (Unknown)
Remaining lease time : (Unknown)
Lease start time     : N/A
```

```
Identifier           : 00:0c:29:00:00:10
MAC address         : 00:0c:29:00:00:1f
IP address          : 192.168.1.254
Service             : 1 (VPRN)
Allocation type      : static
Home-aware pool     : 00:0c:29:00:00:10
DHCP lease          : (Unknown)
Remaining lease time : (Unknown)
Lease start time     : N/A
```

```
-----
No. of BRG hosts: 3
=====
```

```
*A:BNG#
```

Even when the last dynamic host disconnects, the static hosts remain, as shown by the following command.

```
*A:BNG# show subscriber-mgmt brg gateway brg-id "00:0c:29:00:00:10" hosts
```

```
=====
Bridged Residential Gateway hosts
=====
```

```
Identifier           : 00:0c:29:00:00:10
MAC address         : 00:0c:29:00:00:1e
IP address          : 198.51.100.110
Service             : 1 (VPRN)
Allocation type      : static
Home-aware pool     : 00:0c:29:00:00:10
DHCP lease          : (Unknown)
Remaining lease time : (Unknown)
Lease start time     : N/A
```

```
Identifier           : 00:0c:29:00:00:10
MAC address         : 00:0c:29:00:00:1f
IP address          : 192.168.1.254
Service             : 1 (VPRN)
Allocation type      : static
Home-aware pool     : 00:0c:29:00:00:10
DHCP lease          : (Unknown)
Remaining lease time : (Unknown)
Lease start time     : N/A
```

```
-----
No. of BRG hosts: 2
=====
```

```
*A:BNG#
```


As long as the BRG is still alive, these hosts will remain, and so will the BRG context. For that reason, the vRGW might start connectivity verification and eventually delete the BRG and its hosts, depending on the configuration.

The BRG context can also be cleared manually using the following command, after which the BRG and the associated hosts are deleted.

```
*A:BNG# clear subscriber-mgmt brg gateway brg-id "00:0c:29:00:00:10"

*A:BNG# show subscriber-mgmt brg gateway brg-id "00:0c:29:00:00:10" hosts
No entries found.

*A:BNG# show subscriber-mgmt brg gateways
No entries found.
*A:BNG#
```

Implicit authentication

In the implicit authentication scenario, the BRG is authenticated when the first host is authenticated, which requires two phases. In the first phase, the RADIUS server is accessed for authenticating the device (00:0c:29:00:00:11), returning the device-specific data including the BRG ID. In the second phase, the RADIUS server is accessed for authenticating the BRG (00:0c:29:00:00:10), returning the home-specific data. Because the RADIUS server returns two reserved static addresses, SR OS additionally authenticates the static devices.

```
18 2016/09/20 16:30:59.88 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Access-Request(1) 172.16.1.2:1812 id 5 len 79 vrid 1 pol rad-serv-pol-RSP
    USER NAME [1] 17 00:0c:29:00:00:11
    PASSWORD [2] 16 0hUXqLDfL7XPKlx8EmoMak
    NAS IP ADDRESS [4] 4 192.0.2.1
    NAS PORT TYPE [61] 4 Ethernet(15)
    NAS PORT ID [87] 8 1/1/1:81

  Hex Packet Dump:
  ---snip---
"
```

```
19 2016/09/20 16:30:59.89 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Access-Accept(2) id 5 len 45 from 172.16.1.2:1812 vrid 1 pol rad-serv-pol-RSP
    VSA [26] 19 Alcatel(6527)
    BRG ID [225] 17 00:0c:29:00:00:10

  Hex Packet Dump:
  ---snip---
"
```

```
20 2016/09/20 16:30:59.89 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Access-Request(1) 172.16.1.2:1812 id 6 len 88 vrid 1 pol rad-serv-pol-RSP
    USER NAME [1] 17 00:0c:29:00:00:10
    PASSWORD [2] 16 u87pDAyiQx.TVIGEHFV22U
    NAS IP ADDRESS [4] 4 192.0.2.1
    VSA [26] 19 Alcatel(6527)
    BRG ID [225] 17 00:0c:29:00:00:10

  Hex Packet Dump:
  ---snip---
```

```
"
21 2016/09/20 16:30:59.89 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
Access-Accept(2) id 6 len 388 from 172.16.1.2:1812 vrid 1 pol rad-serv-pol-RSP
VSA [26] 19 Alcatel(6527)
  BRG ID [225] 17 00:0c:29:00:00:10
FRAMED IPV6 PREFIX [97] 18 2001:db8:101:1010::/64
VSA [26] 6 Alcatel(6527)
  BRG DMZ ADDRESS [221] 4 192.168.1.254
VSA [26] 44 Alcatel(6527)
  BRG HOME AWARE POOL [220] 42 192.168.1.1/24 192.168.1.100-192.168.1.254
VSA [26] 40 Alcatel(6527)
  BRG RESERVED ADDRESS [223] 38 sticky 00:0c:29:00:00:11 192.168.1.110
VSA [26] 40 Alcatel(6527)
  BRG RESERVED ADDRESS [223] 38 static 00:0c:29:00:00:1f 192.168.1.254
VSA [26] 41 Alcatel(6527)
  BRG RESERVED ADDRESS [223] 39 static 00:0c:29:00:00:1e 198.51.100.110
VSA [26] 20 Alcatel(6527)
  PORTAL URL [177] 18 http://11.11.11.11
VSA [26] 6 Alcatel(6527)
  PRIMARY DNS [9] 4 1.1.1.1
VSA [26] 6 Alcatel(6527)
  SECONDARY DNS [10] 4 1.1.2.2
VSA [26] 6 Alcatel(6527)
  PRIMARY NBNS [29] 4 2.2.1.1
VSA [26] 6 Alcatel(6527)
  SECONDARY NBNS [30] 4 2.2.2.2
VSA [26] 18 Alcatel(6527)
  IPV6 PRIMARY DNS [105] 16 2001:db8:dddd:1::1
VSA [26] 18 Alcatel(6527)
  IPV6 SECONDARY DNS [106] 16 2001:db8:dddd:2::1

Hex Packet Dump:
---snip---
"

22 2016/09/20 16:30:59.89 CEST MINOR: DEBUG #2001 vprn1 PIP
"PIP: DHCP
instance 2 (1), interface index 6 (grp-int-1-1),
  received DHCP Boot Request on Interface grp-int-1-1 (1/1/1:81) Port 67

H/W Type: Ethernet(10Mb)  H/W Address Length: 6
ciaddr: 0.0.0.0            yiaddr: 0.0.0.0
siaddr: 0.0.0.0            giaddr: 0.0.0.0
chaddr: 00:0c:29:00:00:11  xid: 0x1

DHCP options:
[53] Message type: Discover
[255] End
"

23 2016/09/20 16:30:59.89 CEST MINOR: DEBUG #2001 vprn1 PIP
"PIP: DHCP
instance 2 (1), interface index 6 (grp-int-1-1),
  transmitted DHCP Boot Reply to Interface grp-int-1-1 (1/1/1:81) Port 68

H/W Type: Ethernet(10Mb)  H/W Address Length: 6
ciaddr: 0.0.0.0            yiaddr: 192.168.1.110
siaddr: 192.168.1.1        giaddr: 192.168.1.1
chaddr: 00:0c:29:00:00:11  xid: 0x1

DHCP options:
[53] Message type: Offer
```

```
[54] DHCP server addr: 192.168.1.1
[51] Lease time: 21600
[1] Subnet mask: 255.255.255.0
[3] Router: 192.168.1.1
[6] Domain name server: length = 8
    1.1.1.1
    1.1.2.2
[44] NETBIOS name server: length = 8
    2.2.1.1
    2.2.2.2
[255] End
"

24 2016/09/20 16:31:00.00 CEST MINOR: DEBUG #2001 vprn1 PIP
"PIP: DHCP
instance 2 (1), interface index 6 (grp-int-1-1),
    received DHCP Boot Request on Interface grp-int-1-1 (1/1/1:81) Port 67

H/W Type: Ethernet(10Mb)  H/W Address Length: 6
ciaddr: 0.0.0.0            yiaddr: 0.0.0.0
siaddr: 0.0.0.0            giaddr: 0.0.0.0
chaddr: 00:0c:29:00:00:11  xid: 0x1

DHCP options:
[53] Message type: Request
[50] Requested IP addr: 192.168.1.110
[54] DHCP server addr: 192.168.1.1
[255] End
"

25 2016/09/20 16:31:00.00 CEST MINOR: DEBUG #2001 vprn1 PIP
"PIP: DHCP
instance 2 (1), interface index 6 (grp-int-1-1),
    transmitted DHCP Boot Reply to Interface grp-int-1-1 (1/1/1:81) Port 68

H/W Type: Ethernet(10Mb)  H/W Address Length: 6
ciaddr: 0.0.0.0            yiaddr: 192.168.1.110
siaddr: 192.168.1.1        giaddr: 192.168.1.1
chaddr: 00:0c:29:00:00:11  xid: 0x1

DHCP options:
[53] Message type: Ack
[54] DHCP server addr: 192.168.1.1
[51] Lease time: 21600
[1] Subnet mask: 255.255.255.0
[3] Router: 192.168.1.1
[6] Domain name server: length = 8
    1.1.1.1
    1.1.2.2
[44] NETBIOS name server: length = 8
    2.2.1.1
    2.2.2.2
[255] End
"

26 2016/09/20 16:31:00.00 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
Access-Request(1) 172.16.1.2:1812 id 7 len 79 vrid 1 pol rad-serv-pol-RSP
  USER NAME [1] 17 00:0c:29:00:00:1e
  PASSWORD [2] 16 hz9D.7GnC2H.LSaqMZ4TiE
  NAS IP ADDRESS [4] 4 192.0.2.1
  NAS PORT TYPE [61] 4 Ethernet(15)
  NAS PORT ID [87] 8 1/1/1:81
```

```
Hex Packet Dump:
---snip---
"

27 2016/09/20 16:31:00.00 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Access-Request(1) 172.16.1.2:1812 id 8 len 79 vrid 1 pol rad-serv-pol-RSP
    USER NAME [1] 17 00:0c:29:00:00:1f
    PASSWORD [2] 16 zIWsyfz7KatfZ42IB4/uY.
    NAS IP ADDRESS [4] 4 192.0.2.1
    NAS PORT TYPE [61] 4 Ethernet(15)
    NAS PORT ID [87] 8 1/1/1:81

Hex Packet Dump:
---snip---
"

28 2016/09/20 16:30:59.99 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Access-Accept(2) id 7 len 45 from 172.16.1.2:1812 vrid 1 pol rad-serv-pol-RSP
    VSA [26] 19 Alcatel(6527)
    BRG ID [225] 17 00:0c:29:00:00:10

Hex Packet Dump:
---snip---
"

29 2016/09/20 16:30:59.99 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Access-Accept(2) id 8 len 45 from 172.16.1.2:1812 vrid 1 pol rad-serv-pol-RSP
    VSA [26] 19 Alcatel(6527)
    BRG ID [225] 17 00:0c:29:00:00:10

Hex Packet Dump:
---snip---
"
```

Subsequent connections of additional devices connected to the BRG result in one single RADIUS access per device. There is no need to reauthenticate the BRG.

```
30 2016/09/20 16:32:10.74 CEST MINOR: DEBUG #2001 vprn1 PIP
"PIP: DHCP
instance 2 (1), interface index 6 (grp-int-1-1),
  received DHCP Boot Request on Interface grp-int-1-1 (1/1/1:81) Port 67

H/W Type: Ethernet(10Mb)  H/W Address Length: 6
ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
siaddr: 0.0.0.0           giaddr: 0.0.0.0
chaddr: 00:0c:29:00:00:14  xid: 0x2

DHCP options:
[53] Message type: Discover
[255] End
"

31 2016/09/20 16:32:10.74 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Access-Request(1) 172.16.1.2:1812 id 9 len 79 vrid 1 pol rad-serv-pol-RSP
    USER NAME [1] 17 00:0c:29:00:00:14
    PASSWORD [2] 16 TcHWkAooYb5Tcnpj9KPR3M.
    NAS IP ADDRESS [4] 4 192.0.2.1
    NAS PORT TYPE [61] 4 Ethernet(15)
    NAS PORT ID [87] 8 1/1/1:81
```

```
Hex Packet Dump:
---snip---
"

32 2016/09/20 16:32:10.74 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
Access-Accept(2) id 9 len 57 from 172.16.1.2:1812 vrid 1 pol rad-serv-pol-RSP
VSA [26] 19 Alcatel(6527)
BRG ID [225] 17 00:0c:29:00:00:10
VSA [26] 6 Alcatel(6527)
PRIMARY DNS [9] 4 1.1.3.3

Hex Packet Dump:
---snip---
"

33 2016/09/20 16:32:10.74 CEST MINOR: DEBUG #2001 vprn1 PIP
"PIP: DHCP
instance 2 (1), interface index 6 (grp-int-1-1),
transmitted DHCP Boot Reply to Interface grp-int-1-1 (1/1/1:81) Port 68

H/W Type: Ethernet(10Mb) H/W Address Length: 6
ciaddr: 0.0.0.0 yiaddr: 192.168.1.100
siaddr: 192.168.1.1 giaddr: 192.168.1.1
chaddr: 00:0c:29:00:00:14 xid: 0x2

DHCP options:
[53] Message type: Offer
[54] DHCP server addr: 192.168.1.1
[51] Lease time: 21600
[1] Subnet mask: 255.255.255.0
[3] Router: 192.168.1.1
[6] Domain name server: length = 8
1.1.3.3
1.1.2.2
[44] NETBIOS name server: length = 8
2.2.1.1
2.2.2.2
[255] End
"

34 2016/09/20 16:32:10.84 CEST MINOR: DEBUG #2001 vprn1 PIP
"PIP: DHCP
instance 2 (1), interface index 6 (grp-int-1-1),
received DHCP Boot Request on Interface grp-int-1-1 (1/1/1:81) Port 67

H/W Type: Ethernet(10Mb) H/W Address Length: 6
ciaddr: 0.0.0.0 yiaddr: 0.0.0.0
siaddr: 0.0.0.0 giaddr: 0.0.0.0
chaddr: 00:0c:29:00:00:14 xid: 0x2

DHCP options:
[53] Message type: Request
[50] Requested IP addr: 192.168.1.100
[54] DHCP server addr: 192.168.1.1
[255] End
"

35 2016/09/20 16:32:10.84 CEST MINOR: DEBUG #2001 vprn1 PIP
"PIP: DHCP
instance 2 (1), interface index 6 (grp-int-1-1),
transmitted DHCP Boot Reply to Interface grp-int-1-1 (1/1/1:81) Port 68
```

```
H/W Type: Ethernet(10Mb) H/W Address Length: 6
ciaddr: 0.0.0.0          yiaddr: 192.168.1.100
siaddr: 192.168.1.1      giaddr: 192.168.1.1
chaddr: 00:0c:29:00:00:14  xid: 0x2
```

```
DHCP options:
[53] Message type: Ack
[54] DHCP server addr: 192.168.1.1
[51] Lease time: 21600
[1] Subnet mask: 255.255.255.0
[3] Router: 192.168.1.1
[6] Domain name server: length = 8
    1.1.3.3
    1.1.2.2
[44] NETBIOS name server: length = 8
    2.2.1.1
    2.2.2.2
[255] End
```

"

The following command shows the corresponding BRG hosts.

```
*A:BNG# show subscriber-mgmt brg gateway brg-id "00:0c:29:00:00:10" hosts
```

===== Bridged Residential Gateway hosts =====

```
Identifier          : 00:0c:29:00:00:10
MAC address         : 00:0c:29:00:00:11
IP address          : 192.168.1.110
Service             : 1 (VPRN)
Allocation type     : sticky-ip-address
Home-aware pool     : 00:0c:29:00:00:10
DHCP lease          : true
Remaining lease time : 21514
Lease start time    : 2016/09/20 16:31:00
```

```
Identifier          : 00:0c:29:00:00:10
MAC address         : 00:0c:29:00:00:14
IP address          : 192.168.1.100
Service             : 1 (VPRN)
Allocation type     : dynamic
Home-aware pool     : 00:0c:29:00:00:10
DHCP lease          : true
Remaining lease time : 21584
Lease start time    : 2016/09/20 16:32:10
```

```
Identifier          : 00:0c:29:00:00:10
MAC address         : 00:0c:29:00:00:1e
IP address          : 198.51.100.110
Service             : 1 (VPRN)
Allocation type     : static
Home-aware pool     : 00:0c:29:00:00:10
DHCP lease          : (Unknown)
Remaining lease time : (Unknown)
Lease start time    : N/A
```

```
Identifier          : 00:0c:29:00:00:10
MAC address         : 00:0c:29:00:00:1f
IP address          : 192.168.1.254
```

```
Service                : 1 (VPRN)
Allocation type         : static
Home-aware pool         : 00:0c:29:00:00:10
DHCP lease              : (Unknown)
Remaining lease time    : (Unknown)
Lease start time        : N/A
```

```
-----
No. of BRG hosts: 4
=====
```

```
*A:BNG#
```

Releasing both dynamic devices results in all BRG hosts being deleted, including the static hosts.

```
*A:BNG# show subscriber-mgmt brg gateway brg-id "00:0c:29:00:00:10" hosts
No entries found.
*A:BNG#
```

Also, the BRG is deleted automatically.

```
*A:BNG# show subscriber-mgmt brg gateways
No entries found.
*A:BNG#
```

Conclusion

This chapter describes the explicit and the implicit authentication models for the BRG. The explicit authentication model requires the BRG to contain an embedded RADIUS client, and offers better security in comparison with the implicit model, which does not require an embedded RADIUS client.

Virtual Residential Gateway Home LAN Extension

This chapter describes the home LAN extension of the Virtual Residential Gateway.

Topics in this chapter include:

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

Applicability

The information and configuration in this chapter are based on SR OS Release 15.0.R6. Hardware that supports the Wireless LAN Gateway (WLAN-GW) must be used (ISA1/ISA2 cards).

Overview

With Virtual Residential Gateway (vRGW), all L3 routing and services (DHCP server, NAT, firewall, Application Assurance (AA), and so on) are moved out of the home CPE to the service provider network.

The home CPE (known as Bridged Residential Gateway (BRG)) runs in a bridged mode and acts as an L2 switch between all connected home devices and the vRGW. Therefore, in-home traffic will be switched locally while other traffic will be sent to the vRGW.

However, there are some services/applications implemented on the cloud (for example, Data Center (DC)) that need to appear as if they are on the same home LAN and need L2 bridged access, so they need a Home LAN Extension (HLE).

Examples:

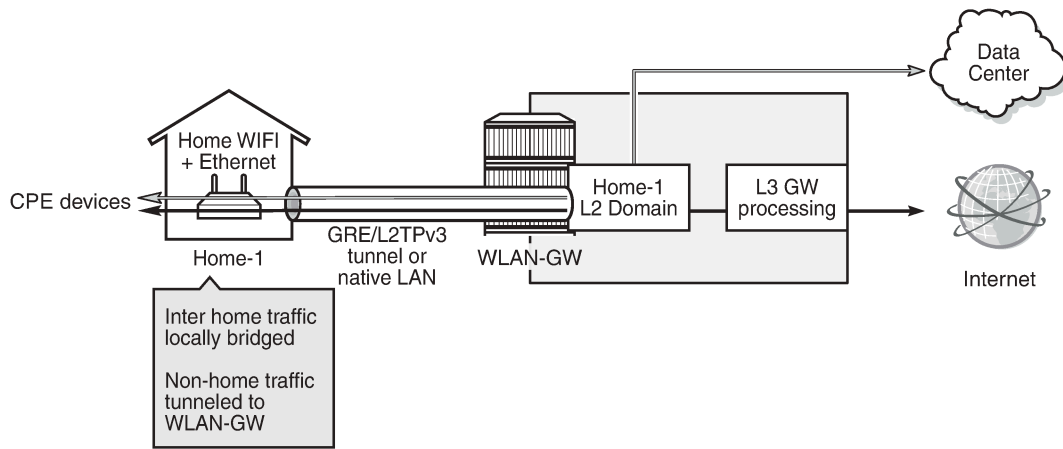
- Digital Living Network Alliance (DLNA) server in the cloud for media streaming
- Storage server in the cloud, discoverable via Server Message Block (SMB) protocol

HLE provides the capability to deploy new services in a DC that require L2 reachability to the home and being reachable on L2 to each individual host at home.

A unique per-home Bridge Domain (BD) is created on the WLAN-GW ISA1/ISA2. This acts as an L2 switch before any IP routing functionality is performed.

By using BGP-EVPN VPLS services, a BD is extended via a Virtual Extensible Local Area Network (VXLAN) tunnel to a virtual machine (VM) in a DC. Access from the home CPE can be a soft-GRE or an L2 Tunnel Protocol version 3 (L2TPv3) tunnel or Native VLAN terminated on the WLAN-GW group ISA via L2 Access Point (L2-AP). The HLE network side relies on MP-BGP for the control plane, whereas the data plane is based on VXLAN; see [Figure 164: vRGW-HLE](#).

Figure 164: vRGW-HLE



27632

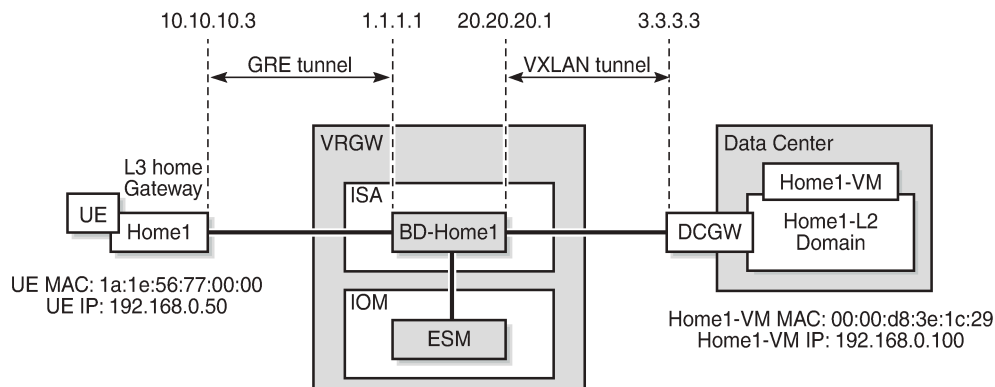
Bridge Domain Characteristics

The BD bridges traffic between the following connections:

- access-facing connection (for example, home) - GRE/L2TPv3/L2-AP
- network-facing connection (for example, DC) - VXLAN tunnel
- ESM SAP-facing connection - each home has its own ESM SAP

Figure 165: BD Connections in the Data Plane shows the BD connections in the data plane.

Figure 165: BD Connections in the Data Plane



27633

With HLE services, each home host constitutes a WLAN-GW User Equipment (UE) object and an ESM host object. Each network host (such as a VM in a DC) constitutes a WLAN-GW UE object but not an ESM host.

Each BD maintains the following tables:

- MAC table: this table contains the learned MAC address by access and network connections

- flood table: this table contains flood destinations for broadcast, unknown unicast, multicast (BUM) traffic; these are typically the access and network connections
- IPv4 ARP table: this table contains learned ARP entries by access and network connections. The ARP and neighbor tables are populated when Assistive Address Resolution (AAR) is enabled (see the Assistive Address Resolution section for more information).
- IPv6 neighbor table: this table contains learned neighbor entries by access and network connections.

Configuration

vRGW Configuration

The following prerequisite configurations are based on vRGW, in the order to be created.

WLAN-GW Group

```
configure wlan-gw-group <wlan-gw-group-id> [create] [redundancy <unit>]

<wlan-gw-group-id>  : [1..4]
<create>            : keyword
<unit>              : iom|mda

configure
  isa
    wlan-gw-group 1 create redundancy mda
    active-mda-limit 1
    mda 3/1
    no shutdown
  exit
exit
exit
```

RADIUS Policies

ISA RADIUS Policy

```
configure
  aaa
    isa-radius-policy "isa-rad-pol-1" create
    password nokia
    servers
      router "Base"
      source-address-range 44.44.44.44
      server 1 create
        authentication
        coa
        ip-address 192.77.77.2
        secret Nokia123
```

```
                no shutdown
            exit
        exit
    exit
exit
```

RADIUS Server Policy

```
configure
router
    radius-server
        server "FreeRadius" address 192.77.77.2 secret Nokia123 create
        accept-coa
    exit
exit
exit
exit
```

```
configure
aaa
    radius-server-policy "rad-serv-pol-1" create
    servers
        router "Base"
        source-address 10.10.10.4
        server 1 name "FreeRadius"
    exit
exit
exit
exit
```

RADIUS Accounting Policy

```
configure
subscriber-mgmt
    radius-accounting-policy "rad-acc-pol-1" create
    ---snip---
    radius-server-policy "rad-serv-pol-1"
exit
exit
```

L2-Aware NAT

```
configure
service
    vprn 333 customer 1 create
    nat
        inside
            l2-aware
            address 192.168.0.101/24
        exit
    exit
    outside
        pool "nat-pool-1" nat-group 1 type l2-aware create
```

```
        port-reservation blocks 1
        address-range 120.1.0.1 120.1.0.100 create
        exit
        no shutdown
    exit
exit
exit
exit
nat
    nat-policy "nat-pol-1" create
    pool "nat-pool-1" router 333
    exit
exit
exit
exit
exit
```

Subscriber Management Parameters

```
configure
    subscriber-mgmt
        authentication-policy "auth-pol-1" create
        exit
        sla-profile "sla-prof-def" create
        exit
        sub-profile "sub-prof-def" create
            nat-policy "nat-pol-1"
            radius-accounting
                policy "rad-acc-pol-1"
            exit
        exit
        sub-ident-policy "sub-ident-pol-def" create
            sub-profile-map
                use-direct-map-as-default
            exit
            sla-profile-map
                use-direct-map-as-default
            exit
        exit
    exit
exit
exit
```

BRG Profile

```
configure
    subscriber-mgmt
        vrgw
            brg
                brg-profile "brg-prof-1" create
                dhcp-pool
                    subnet 192.168.0.1/24 start 192.168.0.50 end 192.168.0.99
                exit
                radius-authentication
                    password nokia
                    radius-server-policy "rad-serv-pol-1"
                exit
            exit
        exit
    exit
```

```
exit
exit
```

vRGW BRG Service

A service can be either an IES or a VPRN; in this example, a VPRN is shown.

```
*A:SR7-CMPT-vRGW# configure service vprn 333
*A:SR7-CMPT-vRGW>config>service>vprn# info
-----
route-distinguisher 333:333
subscriber-interface "sub-int-1" create
  allow-unmatching-subnets
  address 192.168.1.1/24
  group-interface "grp-int-1" wlangw create
  sap-parameters
    sub-sla-mgmt
      def-sla-profile "sla-prof-def"
      def-sub-id use-auto-id
      def-sub-profile "sub-prof-def"
      sub-ident-policy "sub-ident-pol-def"
    exit
  exit
dhcp
  proxy-server
    emulated-server 192.168.1.1
    no shutdown
  exit
  trusted
  lease-populate 1000
  gi-address 192.168.1.1
  no shutdown
exit
authentication-policy "auth-pol-1"
wlan-gw
  gw-address 1.1.1.1
  learn-ap-mac
  router "Base"
  wlan-gw-group 1
  l2-access-points
    l2-ap 4/2/1 create
      encap-type dot1q
      no shutdown
    exit
  exit
  vlan-tag-ranges
    range default
      authentication
        authentication-policy "isa-rad-pol-1"
      exit
      authenticate-on-dhcp
    vrgw
      brg
        default-brg-profile "brg-prof-1"
        no shutdown
      exit
    exit
  exit
  no shutdown
exit
```

```

        exit
    exit
    nat
    ---snip---
    exit
    no shutdown
    -----

```

BGP Configuration

Because HLE is based on BGP-EVPN, the operator needs to configure BGP with address family EVPN and define the corresponding neighbors, where a neighbor can, for example, be a Nuage Virtualized Service Controller (VSC) in a DC, a DC Gateway (DCGW), or a Route Reflector (RR) to avoid full mesh iBGP between iBGP speakers.

```

configure
  router
    bgp
      min-route-advertisement 1
      rapid-withdrawal
      rapid-update evpn
      group "EVPN"
        family vpn-ipv4 evpn
        type internal
        neighbor 10.10.10.3
      exit
    exit
    no shutdown
  exit
exit
exit
exit

```

HLE RADIUS Attributes and User Configuration

Each BD has a unique ID, which is an integer returned by the RADIUS server as the mandatory attribute [241.26.6527.9] **Alc-Bridge-Id** during BRG and device authentication. The returned value must be equal on both authentication levels.

The Alc-Bridge-Id is the ID of the per-subscriber BD on the ISA. It is different from the Alc-BRG-ID, which is the ID of the BRG.

Optional attributes can be returned by the RADIUS server during BRG authentication, which are used for BGP-EVPN VPLS to extend the BD to the remote network in the DC:

- HLE BGP-EVPN route target (RT): [241.26.6527.14] **Alc-RT**, which defines which BD the route belongs to. If not returned, the system defines the RT as "target:<configured_lanext_as>:<Alc-Bridge-Id>". Configuring LANEXT AS will be covered in the [VXLAN/EVPN Parameters](#) section later in this chapter.
- HLE BGP-EVPN route distinguisher (RD): [241.26.6527.15] **Alc-RD**; if not returned, the system defines the RD as "<configured_lanext_as>:<Alc-Bridge-Id>"
- HLE BGP-EVPN VXLAN VNI: [241.26.6527.10] **Alc-Vxlan-VNI** (VXLAN Virtual Network Identifier), which is encoded in the MPLS label field in the EVPN routes, and is used to demux the VXLAN packet into the correct BD. If not returned, the system automatically assigns a VNI.

Examples for user entries on RADIUS server (FreeRadius) follow:

Device (Host) authentication

```
1a:1e:56:55:00:00 Cleartext-Password := "nokia"  
Alc-Subsc-ID-str := "BRG-ID-5",  
Alc-Bridge-Id := 555
```

Home (BRG) authentication

```
BRG-ID-5 Cleartext-Password := "nokia"  
Alc-Bridge-Id := 555,  
Alc-RT := target:100:555,  
Alc-RD := 100:555,  
Alc-Vxlan-VNI := 3
```

For accounting, in addition to the previous attributes, the operator can add the following attribute into accounting messages sent to the accounting server:

[241.26.6527.28] **Alc-HLE-Device-Type** to indicate the type of HLE host; the value is fixed to "home".

An example RADIUS accounting policy including LANEXT attributes follows:

```
configure  
  subscriber-mgmt  
    radius-accounting-policy "rad-acc-pol-1" create  
      ---snip---  
      include-radius-attribute  
      ---snip---  
        lanext-bridge-id  
        lanext-device-type  
        lanext-route-distinguisher  
        lanext-route-target  
        lanext-vni  
      ---snip---  
    exit  
  radius-server-policy "rad-serv-pol-1"  
exit  
exit  
exit
```

VXLAN/EVPN Parameters

```
A:WLANGW# configure router vrgw lanext  
- lanext  
[no] vxlan-port          - Configure the remote VXLAN UDP port  
[no] vxlan-vtep-ran*     - Configure a range of VXLAN VTEP addresses  
[no] wlan-gw-group       - Configure the ISA WLAN Gateway group  
[no] shutdown            - Enable/disable the Home LAN Extension functionality
```

Where:

```
- vxlan-port <4789|8472>          # default is 4789  
- vxlan-vtep-range start <ip-address|ipv6-address> end <ip-address|ipv6-address>  
- wlan-gw-group <wlan-gw-group-id> # value configured in previous step
```

Example:

```
configure
```

```

router
  vrgw
    lanext
      vxlan-vtep-range start 20.20.20.1 end 20.20.20.20
      wlan-gw-group 1
      no shutdown
    exit
  exit
exit

```

HLE is enabled on the Base routing instance.

As mentioned, RT is derived from the target:LANEXT_AS:bridge-id.

The LANEXT Autonomous System (AS) number can be configured in the following context:

```

A:WLANGW# configure subscriber-mgmt vrgw lanext router-target-as-number
- router-target-as-number <as-number>

<as-number>          : [1..65535]

```

Example:

```

configure
  subscriber-mgmt
    vrgw
      lanext
        router-target-as-number 100
      exit
    exit
  exit
exit

```

If the router AS number is not configured on the WLAN-GW, it must be returned from RADIUS in both Alc-RT & Alc-RD attributes during BRG authentication; otherwise, the host DHCP fails and the system shows the following error message in log 99.

```

53 2018/01/03 15:48:53.236 PST WARNING: DHCP #2005 vprn333 Lease State Population
Error
"Lease state table population error on SAP 3/1/nat-out-ip:2145.22 in service 333 - Could not
link IPoE session (3/1/nat-out-ip:2145.22 - 1a:1e:56:77:00:00) with BRG BRG-ID-7: Router AS
number is not configured"

```

BD-related Settings

On the group-interface level, the overall number of allowed BDs created can be limited. By default, the limit is zero, but it can be set between 1 and 131 071.

```

A:WLANGW>config>service>vprn>sub-if>grp-if>wlan-gw# max-lanext-bd
- max-lanext-bd <[1..131071]>

```

When HLE is enabled (under the VLAN vRGW LANEXT hierarchy) and **max-lanext-bd** on the group-interface level is zero, the following error message is displayed when the WLAN-GW is enabled in the **group-interface** context:

```

*A:WLANGW>config>service>vprn>sub-if>grp-if>wlan-gw# no shutdown

```



```
MINOR: SVCNMR #8456 The group-interface's max-lanext-bd must be greater than 0 to enable lanext
in the VLAN range - VLAN range 4096-4096
```

If the number of BDs reaches that limit, the following message is displayed in log 99:

```
114 2018/01/03 16:04:14.215 PST WARNING: DHCP #2029 Base Miscellaneous DHCP Problem
"Failed to create bridge domain: max-lanext-bd limit reached - Bridge ID: 555, MAC:
1a:1e:56:55:00:00"
```

BD Access-Side Settings

The **max-mac** parameter in the **access** context is used to limit the number of MAC addresses learned from home devices:

```
*A:SR7-CMPT-vRGW>grp-if>wlan-gw>ranges>range>vrgw>lanext# access
- access

[no] max-mac          - Configure maximum number of MAC entries in bridged domain
[no] multi-access     - Allow multiple access
```

The **max-mac** ranges between 1 and 128, with a default value of 20.

The **multi-access** parameter is used for Multi-Dwelling Unit (MDU) scenarios.

To override the default parameters, the vRGW LANEXT hierarchy should be in the shutdown state; an error is raised when trying to configure these parameters when LANEXT is enabled.

```
*A:WLANGW>grp-if>wlan-gw>ranges>range>vrgw>lanext>access# max-mac 128
MINOR: SVCNMR #8452 Not allowed when lanext is no shut

*A:WLANGW>grp-if>wlan-gw>ranges>range>vrgw>lanext>access# multi-access
MINOR: SVCNMR #8452 Not allowed when lanext is no shut
```

BD Network-Side Settings

The **max-mac** parameter in the **network** context is used to limit the number of MAC addresses learned from EVPN/DC.

```
*A:SR7-CMPT-vRGW>grp-if>wlan-gw>ranges>range>vrgw>lanext# network
- network

[no] max-mac          - Configure maximum number of VM MAC entries in bridged domain
[no] shutdown         - Enable/disable data center connections
```

The **max-mac** ranges between 1 and 64, with a default value of 20. LANEXT is by default enabled (no shutdown).

```
*A:SR7-CMPT-vRGW>grp-if>wlan-gw>ranges>range>vrgw>lanext>network# max-mac
<[1..64]>
```

The operator can shut down the network side to disable the connection to the DC.

To override the default configuration, LANEXT should be shut down, as follows:

```
*A:SR7-CMPT-vRGW>grp-if>wlan-gw>ranges>range>vrgw>lanext>network# max-mac 64
```

```
MINOR: SVCMGR #8452 Not allowed when lanext is no shut
```

```
*A:SR7-CMPT-vRGW>grp-if>wlan-gw>ranges>range>vrgw>lanext>network# shutdown
MINOR: SVCMGR #8452 Not allowed when lanext is no shut
```

Finally, the LANEXT functionality is enabled under the following context:

```
*A:WLANGW>grp-if>wlan-gw>ranges>range>vrgw>lanext#
[no] shutdown          - Enable/disable the Home LAN Extension functionality
```

ARP/ND Handling and MAC Advertisement Optimizations

Assistive Address Resolution

AAR is an optional HLE feature used to avoid sending ARP/ND requests from the home across a WAN to a remote network, or sending ARP/ND requests from the remote network across an access network to a home. The system responds to the ARP/ND request from the network or home with the learned ARP/ND entries, instead of flooding them.

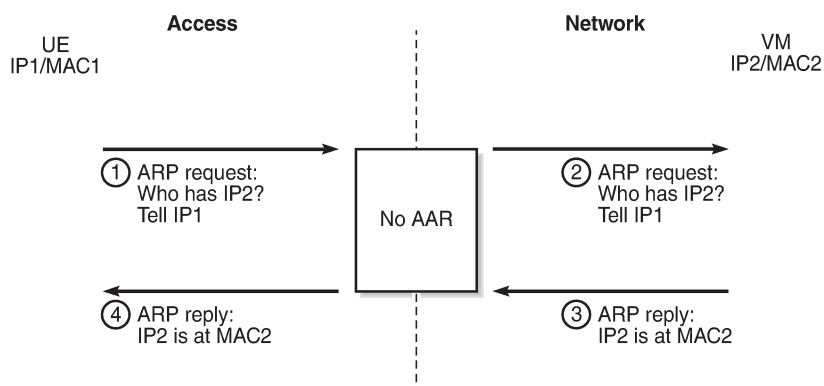
With AAR enabled, the system populates the ARP and neighbor tables with the learned ARP and neighbor entries via either:

- BGP EVPN MAC routes that contain an IP address, or
- (G)ARP/ND/NS packets

AAR Disabled

AAR is disabled by default. With AAR disabled, the ARP request is flooded from one side to the other; see [Figure 166: ARP Requests Flooded with AAR Disabled](#). The BD ARP table will only be populated for the GW address and any IP/MAC learned via BGP EVPN MAC routes that contain an IP address.

Figure 166: ARP Requests Flooded with AAR Disabled



27634

AAR Enabled

To enable AAR, **assistive-address-resolution** is configured in the **lanext** context, as follows:

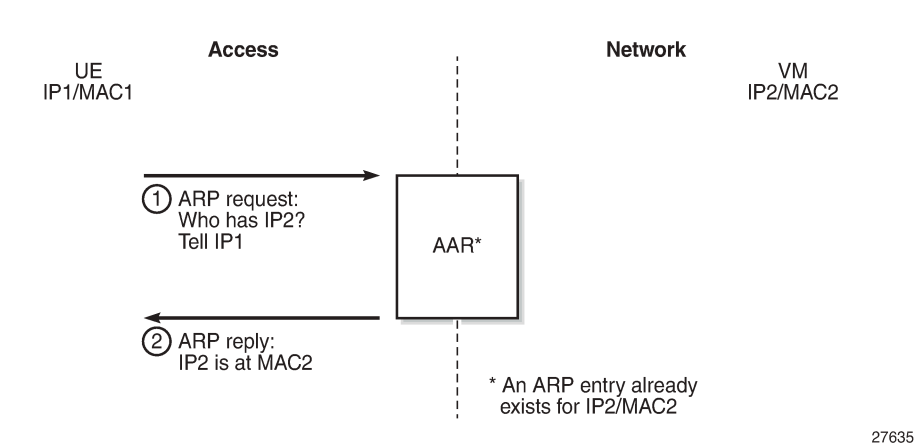
```
*A:SR7-CMPT-vRGW>grp-if>wlan-gw>ranges>range>vrgw>lanext#  
[no] assistive-addr* - Enable/disable assistive address resolution in bridged domains
```

AAR can only be enabled when **lanext** is shut down, as follows:

```
A:WLANGW>grp-if>wlan-gw>ranges>range>vrgw>lanext# assistive-address-resolution  
MINOR: SVCNMR #8452 Not allowed when lanext is no shut
```

In [Figure 167: No ARP Request Flooding with AAR Enabled](#), when an ARP request is received, the ISA performs a lookup in the ARP table with the IP address as a key, and if found, it answers without flooding to the other side.

Figure 167: No ARP Request Flooding with AAR Enabled



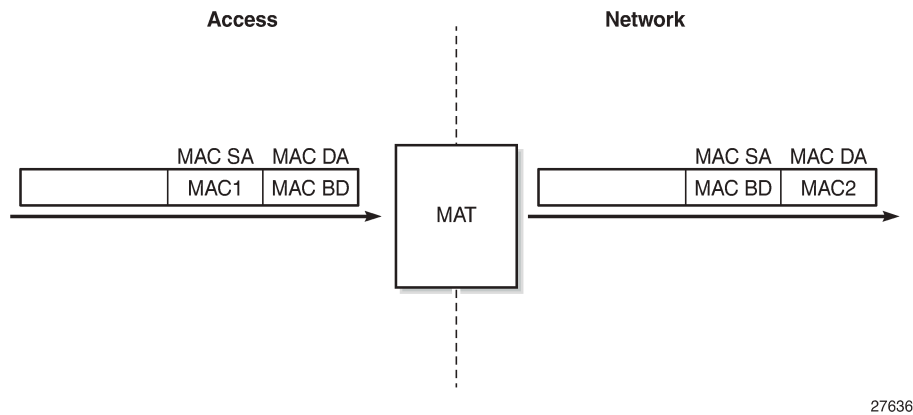
MAC Address Translation

MAC Address Translation (MAT) is an optional HLE feature that translates the host MAC address of a subscriber into a single BD MAC address. This feature decreases the number of BGP EVPN MAC routes to advertise per subscriber to one, and prevents BGP sending update messages when hosts are created and removed, which increases BGP stability.

MAT is performed on traffic in both directions:

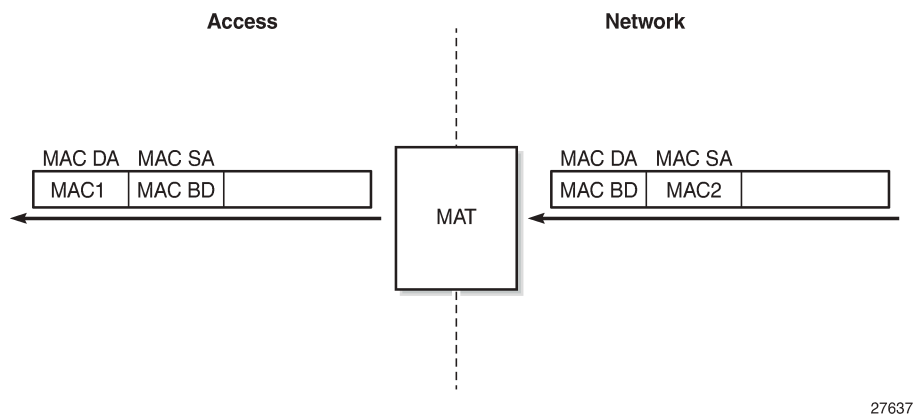
- traffic from access to network - the system changes the source MAC address to BD MAC address, as shown in [Figure 168: MAT - Access to Network Direction](#).

Figure 168: MAT - Access to Network Direction



- traffic from network to access - the system changes the destination MAC address (BD MAC) to the real host MAC address based on the ARP table or neighbor table lookup using the IP address as key, as shown in [Figure 169: MAT - Network to Access Direction](#).

Figure 169: MAT - Network to Access Direction



MAT is disabled by default. MAC is enabled by configuring **mac-translation** and **bd-mac-prefix** under the **lanext** context, as follows.

```
*A:SR7-CMPT-vRGW>grp-if>wlan-gw>ranges>range>vrgw>lanext#
[no] assistive-addr* - Enable/disable assistive address resolution in bridged domains
[no] bd-mac-prefix   - Configure MAC translation prefix in bridged domains
[no] mac-translation - Enable/disable MAC translation in bridged domains
```

To configure MAT, the following conditions apply:

- Assistive address resolution must be enabled

```
*A:WLANGW>grp-if>wlan-gw>ranges>range>vrgw>lanext# no shutdown
MINOR: SVCMMGR #8455 mac-translation requires assistive-address-resolution
```

- **lanext** context must be shut down

```
A:WLANGW>grp-if>wlan-gw>ranges>range>vrgw>lanext# mac-translation
MINOR: SVCMMGR #8452 Not allowed when lanext is no shut
```

- BD-MAC-Prefix must be configured

```
*A:WLANGW>grp-if>wlan-gw>ranges>range>vrgw>lanext# shutdown
*A:WLANGW>grp-if>wlan-gw>ranges>range>vrgw>lanext# mac-translation
*A:WLANGW>grp-if>wlan-gw>ranges>range>vrgw>lanext# no shutdown
MINOR: SVCMMGR #8454 mac-translation requires a bd-mac-prefix to be set
*A:WLANGW>grp-if>wlan-gw>ranges>range>vrgw>lanext#
```

When MAT is configured, the BD ARP table will use a system-generated MAC address with MAC prefix as the configured BD-MAC-prefix, instead of the group-interface MAC address for the default GW entry.

This MAC address will be advertised toward the network as a BGP EVPN MAC route (route type 2), and used for MAC translation, as shown in [Figure 168: MAT - Access to Network Direction](#) and [Figure 169: MAT - Network to Access Direction](#).



Note:

- When MAT is disabled (default), ARP replies contain the real MAC address, in case there is an existing entry in the ARP table; otherwise, the ARP requests are flooded to the other side. Data traffic will contain the real MAC destination address (DA).
- When MAT is enabled, ARP replies use the BD MAC address in both directions. No MAC routes from the host side are learned at the network side. Data traffic will contain the BD MAC DA.
- For two hosts within the same BD, and for traffic in the network to access direction, the ISA does a lookup for ARP entries with the IP DA as a key. If the IP DA is found, the MAC DA is replaced by the MAC DA of the corresponding host.

[Table 38: ARP/MAT/BD-MAC-Prefix Possible Combinations](#) shows possible combinations for ARP/MAT/BD-MAC-Prefix, where Combo1 is the default combination.

Table 38: ARP/MAT/BD-MAC-Prefix Possible Combinations

vRGW lanext	Combo1	Combo2	Combo3	Combo4	Combo5
assistive-address-resolution	Disabled	Disabled	Enabled	Enabled	Enabled
bd-mac-prefix	Disabled	Enabled	Disabled	Enabled	Enabled
mac-translation	Disabled	Disabled	Disabled	Disabled	Enabled

Show Commands

The following command displays BD info, or a specific BD:

```
show subscriber-mgmt vrgw lanext bd [<bridge-id>]
```

```
A:WLANGW# show subscriber-mgmt vrgw lanext bd 555
```

```
=====
```

```
Bridge Domain(s)
=====
Bridge-id           : 555
VNI                 : 3
Route Target        : target:100:555
Route Distinguisher : 100:555
WlanGw GrpId        : 1
ISA MemberId        : 1
WlanGw Bd VlanTag    : N/A
WlanGw Bd Service    : 333
WlanGw Bd Interface  : grp-int-1
WlanGw Bd Mac Translation : enabled
WlanGw Bd Mac        : aa:bb:cc:00:00:03
WlanGw Bd Assist.Addr.Res : enabled
WlanGw Bd Netw MaxMac : 20
WlanGw Bd Netw State : enabled
WlanGw Bd Accs MaxMac : 20

-----
No. of Bridge Domains: 1
=====
```

The following command displays UE information for a specific BD:

```
show subscriber-mgmt wlan-gw ue bd <bridge-id>
```

Each host at home or VM in a DC is considered a WLAN-GW UE object, as follows:

```
A:WLANGW# show subscriber-mgmt wlan-gw ue bd 777

=====
User Equipments
=====
Bridge ID           : 777
MAC address         : 00:00:d8:3e:1c:29
-----
VLAN Q-tag          : (Not Specified)
MPLS label          : (Not Specified)
Tunnel router        : "Base"
Tunnel remote IP address : 3.3.3.3
Tunnel local IP address  : 20.20.20.1
Tunnel encapsulation    : vxlan
Retail service       : N/A
SSID                 : (Not Specified)
Previous Access Point IP : (Not Specified)
IMSI                 : (Not Specified)
Subscriber host service : N/A
Subscriber host SAP    : N/A
Last move time       : 2018/01/04 19:33:35

Bridge ID           : 777
MAC address         : 00:21:05:a1:ba:6a
-----
VLAN Q-tag          : (Not Specified)
MPLS label          : (Not Specified)
Tunnel router        : "Base"
Tunnel remote IP address : 3.3.3.3
Tunnel local IP address  : 20.20.20.1
Tunnel encapsulation    : vxlan
Retail service       : N/A
SSID                 : (Not Specified)
Previous Access Point IP : (Not Specified)
IMSI                 : (Not Specified)
```

```
Subscriber host service : N/A
Subscriber host SAP     : N/A
Last move time         : 2018/01/04 19:31:57

Bridge ID              : 777
MAC address            : 1a:1e:56:77:00:00
-----
VLAN Q-tag             : 777
MPLS label             : 1020
Tunnel router          : "Base"
Tunnel remote IP address : 10.10.10.3
Tunnel local IP address  : 1.1.1.1
Tunnel encapsulation    : gre
Retail service         : N/A
SSID                   : (Not Specified)
Previous Access Point IP : (Not Specified)
IMSI                   : (Not Specified)
Subscriber host service : 333
Subscriber host SAP     : 3/1/nat-out-ip:2145.24
Last move time         : 2018/01/04 19:31:55
-----
```

```
No. of UE: 3
=====
```

The previous example was for GRE access (tunnel encapsulation GRE for UE on the access side); the following example is for the VLAN native L2-AP case.

```
*A:WLANGW# show subscriber-mgmt wlan-gw ue bd 555
```

```
=====
User Equipments
=====
```

```
Bridge ID              : 555
MAC address            : 00:00:d8:3e:1c:27
-----
VLAN Q-tag             : (Not Specified)
MPLS label             : (Not Specified)
Tunnel router          : "Base"
Tunnel remote IP address : 10.10.10.3
Tunnel local IP address  : 20.20.20.1
Tunnel encapsulation    : vxlan
Retail service         : N/A
SSID                   : (Not Specified)
Previous Access Point IP : (Not Specified)
IMSI                   : (Not Specified)
Subscriber host service : N/A
Subscriber host SAP     : N/A
Last move time         : 2018/01/05 10:44:41

Bridge ID              : 555
MAC address            : 1a:1e:56:55:00:00
-----
VLAN Q-tag             : (Not Specified)
MPLS label             : (Not Specified)
Tunnel router          : "Base"
Tunnel remote IP address : fe80::22bf:ffff:fe00:1901
Tunnel local IP address  : fe80::ff:fe02:202
Tunnel encapsulation    : vlan
Retail service         : N/A
SSID                   : (Not Specified)
Previous Access Point IP : (Not Specified)
IMSI                   : (Not Specified)
```

```
Subscriber host service      : 333
Subscriber host SAP         : 3/1/nat-out-ip:2145.29
Last move time             : 2018/01/05 10:43:54
```

```
-----
No. of UE: 2
=====
```

Tools Commands

The following lists **tools** commands dump info regarding BD data plane tables, as well as BD statistics:

```
*A:WLANGW# tools dump wlan-gw lanext bd
- bd <bridge-id>

arp-table      - Dump bridge domain ARP table
flood-table    - Dump bridge domain flood table
mac-table      - Dump bridge domain MAC table
neighbor-table - Dump bridge domain neighbor table
statistics     - Dump bridge domain statistics
```

The following **tools** command is used to clear ARP/MAC/Neighbor table entries, as well as BD statistics:

```
*A:WLANGW# tools perform wlan-gw lanext bd
- bd <bridge-id>

clear-arp      - Clear bridge domain ARP table
clear-mac      - Clear bridge domain MAC table
clear-neighbor - Clear bridge domain neighbor table
clear-statisti* - Clear bridge domain statistics
```

ARP Table

The following **tools** command shows the ARP table when AAR is disabled and ARP requests are flooded when the ARP table does not contain an entry for the IP DA:

```
A:WLANGW# tools dump wlan-gw lanext bd 777 arp-table

=====
Matched 2 ARP entries for Bridge Domain 777 on Slot #3 MDA #1
=====
```

IP Address	MAC Address	Class	Timestamp
192.168.0.101	00:21:05:a1:ba:6a	Network	01/03/2018 11:59:39
192.168.0.1	00:00:00:02:02:02	Access	01/03/2018 11:59:37

```
=====
```

Where:

- 192.168.0.1 00:00:00:02:02:02 is the entry learned from the BRG subnet default GW and group-interface MAC
- 192.168.0.101 00:21:05:a1:ba:6a is an entry learned from the network side via the BGP EVPN MAC route (type 2), which contains IP address 192.168.0.101, as shown in the following output:

```
A:WLANGW# show router bgp routes evpn mac mac-address 00:21:05:a1:ba:6a
```



```

=====
BGP Router ID:10.10.10.4      AS:65100      Local AS:65100
=====
Legend -
Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                  l - leaked, x - stale, > - best, b - backup, p - purge
Origin codes  : i - IGP, e - EGP, ? - incomplete
=====
BGP EVPN MAC Routes
=====
Flag  Route Dist.      MacAddr      ESI
      Tag              Mac Mobility  Label1
              Ip Address
              NextHop
-----
u*>i  200:777          00:21:05:a1:ba:6a ESI-0
      0                Static        VNI 777
                  192.168.0.101
                  3.3.3.3
-----
Routes : 1
=====

```

The following **tools** command shows the ARP table when AAR is enabled and bd-mac-prefix configured:

```

*A:WLANGW# tools dump wlan-gw lanext bd 666 arp-table
=====
Matched 4 ARP entries for Bridge Domain 666 on Slot #3 MDA #1
=====
IP Address      MAC Address      Class      Timestamp
-----
192.168.0.100    00:00:d8:3e:1c:28 Network 01/05/2018 11:27:59
192.168.0.51     1a:1e:56:66:00:01 Access 01/05/2018 11:27:59
192.168.0.1      aa:bb:cc:00:00:02 Access 01/05/2018 10:43:54
192.168.0.50     1a:1e:56:66:00:00 Access 01/05/2018 11:27:59
=====

```

Where:

- 192.168.0.1 aa:bb:cc:00:00:02 is the entry for the BRG subnet default GW with bd-mac-prefix AA:BB:CC configured
- 192.168.0.100 00:00:d8:3e:1c:28 is an entry learned from the network side via ARP

MAC Table

The following **tools** command shows the MAC table for GRE access:

```

*A:WLANGW# tools dump wlan-gw lanext bd 777 mac-table
=====
Matched 3 MAC entries for Bridge Domain 777 on Slot #3 MDA #1
=====
UE-Mac      Vlan  Bridge-Id  Description      L2-Svc  L2-Vlan
Tunnel      Class  Type      Src-IP           Dst-IP
Anchor
-----
1a:1e:56:77:00:00 777    777        ESM-user         N/A     N/A

```

3/1	Access	GRE	Base	10.10.10.3		
3/1/nat-out-ip:2145.28				1.1.1.1		

00:00:d8:3e:1c:29	N/A	777	ESM-user		N/A	N/A
3/1	Network	VXLAN	Base	3.3.3.3		
3/1/nat-out-ip:2145.28				20.20.20.1		

00:21:05:a1:ba:6a	N/A	777	ESM-user		N/A	N/A
3/1	Network	VXLAN	Base	3.3.3.3		
3/1/nat-out-ip:2145.28				20.20.20.1		

=====						

The following **tools** command shows the MAC table for VLAN native L2-AP access:

```
A:WLANGW# tools dump wlan-gw lanext bd 555 mac-table
```

=====						
Matched 2 MAC entries for Bridge Domain 555 on Slot #3 MDA #1						
=====						
UE-Mac		Vlan	Bridge-Id	Description	L2-Svc	L2-Vlan
Tunnel	Class	Type	Router	Src-IP		
Anchor				Dst-IP		

00:00:d8:3e:1c:27	N/A	555	ESM-user		N/A	N/A
3/1	Network	VXLAN	Base	10.10.10.3		
3/1/nat-out-ip:2145.29				20.20.20.1		

1a:1e:56:55:00:00	N/A	555	ESM-user		2147483692	555
3/1	Access	L2	Base	fe80::22bf:ffff:fe00:1901		
3/1/nat-out-ip:2145.29				fe80::ff:fe02:202		

=====						

The type is VXLAN for MAC addresses learned from the network side.

Flood Table

The following **tools** command shows the flood table for GRE:

```
*A:WLANGW# tools dump wlan-gw lanext bd 777 flood-table
```

=====				
Matched 2 flood entries for Bridge Domain 777 on Slot #3 MDA #1				
=====				
Tunnel	Class	Type	Vlan	Src-IP
Port		Encap		Dst-IP

3/1	Access	GRE	777	10.10.10.3
3/1/nat-in-ip		2049.1		1.1.1.1

3/1	Network	VXLAN	N/A	3.3.3.3
3/1/nat-out-ip		2081.3		20.20.20.1

=====				

The following **tools** command shows the flood table for VLAN native L2-AP:

```
*A:WLANGW# tools dump wlan-gw lanext bd 555 flood-table
```

```
=====
Matched 2 flood entries for Bridge Domain 555 on Slot #3 MDA #1
=====
```

Tunnel Port	Class	Type Encap	Vlan	Src-IP Dst-IP
3/1	Access	L2	N/A	fe80::22bf:ffff:fe00:1901
3/1/nat-in-ip		2305.1		fe80::ff:fe02:202
3/1	Network	VXLAN	N/A	10.10.10.3
3/1/nat-out-ip		2081.3		20.20.20.1

```
=====
```

Statistics

To view statistics for traffic crossing a specific BD, the following **tools** command is used:

```
*A:WLANGW# tools dump wlan-gw lanext bd 555 statistics
```

```
=====
Statistics for Bridge Domain 555 on Slot #3 MDA #1
=====
```

upstream rx frames	: 9703260
downstream rx frames	: 2
upstream rx octets	: 9664443740
downstream rx octets	: 692
arp entry dropped	: 0
arp entry no resources	: 0
arp entry expired	: 0
arp entry conflict	: 0
neighbor entry dropped	: 0
neighbor entry no resources	: 0
neighbor entry expired	: 0
neighbor entry conflict	: 0
flood entry dropped	: 0
flood no resources	: 0
flood throttled	: 0
flood upstream forward	: 1
flood upstream dropped	: 0
flood downstream forward	: 0
flood downstream dropped	: 0

```
=====
```

Where:

- Upstream direction means traffic coming from the DC or home going toward the BD.
- Downstream direction means traffic coming from the Internet toward the BD.

Conclusion

This chapter describes HLE based on vRGW and BGP EVPN to extend home LAN to DC VMs via VXLAN, for different access models (softGRE and VLAN L2-AP).

Virtual Residential Gateway Home Pool Management

This chapter describes virtual residential gateway home pool management.

Topics in this chapter include:

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

Applicability

The information and configuration in this chapter are applicable to SR OS Release 14.0.R3.

Overview

In the virtual residential gateway model, the Layer 3 (L3) functions are moved out of the traditional residential gateway (RGW) and into the network; see [Figure 170: Virtual Residential Gateway in the Network with Bridged Residential Gateway at Home](#). Examples of L3 functions moved to the network are:

- DHCPv4
- Network Address Translation (NAT)
- Firewalling
- Universal Plug and Play

The in-home equipment interconnecting all devices in the home is referred to as the bridged residential gateway (BRG). The BRG only handles Layer 2 (L2) connectivity (for example, Ethernet and WiFi) and is always operating in bridged mode. The BRG has an L2 uplink connecting it to the virtual residential gateway (vRGW), either through a direct link or through tunneling technology. The vRGW handles all L3 connectivity.

For the vRGW to offer IPv4 connectivity to the devices in the home network, the vRGW provides the following features:

- Private addresses from a single home address pool are offered. The address pools can overlap between homes.
- Sticky or static addresses provide a fixed device to IP mapping.
- Public addresses can be assigned, to enable home servers to be publicly accessible.
- A demilitarized (DMZ) host can be defined.


Because multiple homes are allowed to use the same private subnet, the vRGW requires L2-aware NAT. L2-aware NAT is handled in the "L2-aware NAT (with dNAT and MNPs)" chapter in *7450 ESS and 7750 SR Multiservice ISA and ESA Advanced Configuration Guide for Classic CLI*.

For the vRGW to offer IPv6 connectivity to the devices in the home network, the vRGW provides the following features:

- A /64 SLAAC prefix is assigned per home.
- IA_NA address allocation using DHCPv6 relay or proxy is supported, following standard ESM rules.
- Prefix Delegation (PD) is not supported.

Using the vRGW model, the ISP now has visibility on the MAC and IP addresses used by the in-home devices.

Figure 170: Virtual Residential Gateway in the Network with Bridged Residential Gateway at Home



25957

DHCP and IP Address Management

The vRGW has the following characteristics for DHCPv4:

- One pool per home
- IP overlap between homes
- Sticky IP addresses

Sticky IP addresses are DHCPv4 addresses assigned to devices that need to have the same address all the time and are provided through the DHCP protocol. Home servers, network-attached storage (NAS), and network printers are examples of devices that typically are configured with sticky IP addresses. The vRGW sets a flag indicating that the IP address is reserved, to avoid assigning the sticky address to devices that do not have this requirement.

Static IP addresses can be used and configured for devices that do not use or support the DHCP protocol, and are configured manually on the device. The static IP address used can be a public or a private address. Traffic to and from private addresses undergoes NAT, whereas traffic to and from public addresses does not undergo NAT. The vRGW drops DHCP messages originating from devices that are considered to be static.

The vRGW requires the use of IPoE sessions for supporting BRGs, and creates an IPoE session per device. Static devices require at least one dynamic host to be created first so that the associated SAP or tunnel can be defined to send traffic over. Each static device is authenticated individually using RADIUS to retrieve the per-device parameters.

The vRGW has the following characteristics for IPv6:

- Both SLAAC and IA_NA allocation are supported.

- Prefix delegation (PD) is not supported in the vRGW.

25958

Services Configuration

An excerpt of the configuration of VPRN-1 follows. This VPRN contains subscriber interface *sub-int-1* with group interface *grp-int-1-1*, thereby using the routed central office (CO) model. No static SAPs are configured; managed SAPs will be created when triggers are received on VPLS-10. The IPoE session policy required for the vRGW is *sess-pol-SAP-MAC*, and the default BRG profile used is *brg-prof-1*.

Relay and proxy scenarios are configured for DHCPv4 and DHCPv6. SLAAC prefixes are taken from the wan-host prefix range, and must be advertised to devices using router-advertisement (RA) messages in response to router-solicit (RS) messages. For that purpose, group interface *grp-int-1-1* is configured in its IPv6 context to support router advertisement and router solicitation. Also IPoE bridged mode is enabled in this context, so that the same SLAAC prefix can be allocated multiple times to the same SAP.

RADIUS authentication is enabled through authentication policy *radius-AUTH*. The inside L2-aware NAT range is 192.168.0.1/16, and the DHCP pool subnets defined in the BRG profile must belong to that range; see the [ESM Configuration](#) part in this chapter.

```
configure
service
  vprn 1 customer 1 create
  ---snip---
  route-distinguisher 64496:1
  interface "int-DHCP" create
    address 10.11.11.1/32
    ipv6
      address 2001:db8::11/128
      local-dhcp-server "dhcp6-SRVC1"
    exit
  loopback
exit
subscriber-interface "sub-int-1" create
  address 198.51.100.254/24
  address 192.168.1.1/24
  ipv6
    link-local-address fe80::1
    subscriber-prefixes
      prefix 2001:db8:101::/48 wan-host
    exit
  exit
group-interface "grp-int-1-1" create
  ipv6
    router-advertisements
      prefix-options
        autonomous
      exit
      no shutdown
    exit
    dhcp6
      proxy-server
        no shutdown
      exit
      relay
        server 2001:db8::11
        no shutdown
      exit
    exit
    router-solicit
      inactivity-timer hrs 2
      no shutdown
    exit
```

```
        ipoe-bridged-mode
    exit
    ---snip---
    arp-populate
    dhcp
        proxy-server
            emulated-server 198.51.100.254
            no shutdown
        exit
        trusted
        lease-populate 128
        gi-address 198.51.100.254
        no shutdown
    exit
    authentication-policy "radius-AUTH"
    ipoe-session
        ipoe-session-policy "sess-pol-SAP-MAC"
        sap-session-limit 128
        no shutdown
    exit
    brg
        default-brg-profile "brg-prof-1"
        no shutdown
    exit
    oper-up-while-empty
exit
exit
nat
    inside
        l2-aware
        address 192.168.0.1/16
    exit
    exit
exit
    no shutdown
exit
    exit
exit
exit
```

VPRN-2 defines an interface to the outside world, *int-VPRN2-INTERNET*, as well as the outside L2-aware NAT range, which is 203.0.113.1 up to 203.0.113.10. Port-reservation blocks is set to 1, to ensure a unique outside IP per subscriber (home) and correct operation of the DMZ feature.

```
configure
    service
        vprn 2 customer 1 create
            route-distinguisher 64496:2
            interface "int-VPRN2-INTERNET" create
            ---snip---
        exit
        nat
            outside
                pool "nat-outside-1" nat-group 1 type l2-aware create
                port-reservation blocks 1
                address-range 203.0.113.1 203.0.113.10 create
                exit
                no shutdown
            exit
        exit
        no shutdown
    exit
exit
```



```
exit
```

VPLS-10 defines the capture SAP on port 1/1/1. The triggers configured are dhcp, dhcp6, and rtr-solicit. The authentication policy and the IPoE session policy used are the same as the ones used on VPRN-1, and are *radius-AUTH* and *sess-pol-SAP-MAC*, respectively. The MSAP defaults indicate that the managed SAPs must be created on service 1, group interface *grp-int-1-1*, using MSAP policy *msap-pol-MSAP*.

```
configure
service
  vpls 10 customer 1 create
    description "VPLS for capture SAPs - BRG-demo"
    stp
      shutdown
    exit
    sap 1/1/1:* capture-sap create
      description "capture SAP for MSAP creation on 1/1/1"
      trigger-packet dhcp dhcp6 rtr-solicit
      ipoe-session
        ipoe-session-policy "sess-pol-SAP-MAC"
        no shutdown
      exit
      msap-defaults
        group-interface "grp-int-1-1"
        policy "msap-pol-MSAP"
        service 1
      exit
      authentication-policy "radius-AUTH"
      no shutdown
    exit
  no shutdown
exit
exit
exit
exit
```

RADIUS User Configuration

Although a protocol conversion and mediation platform (PCMP) will be used in conjunction with an external controller, for demonstration purposes this chapter relies on a RADIUS server only.

A sample RADIUS user configuration follows. The user's MAC address is used for authenticating purposes. MAC address 00:0c:29:00:00:10 identifies the BRG. The addresses ranging from 00:0c:29:00:00:11 to 00:0c:29:00:00:1f identify the home devices connected to the same BRG so they all return the same Alc-BRG-Id. For the BRG, RADIUS provides primary and secondary DNS and NBNS servers, a set of reserved addresses, a DMZ address, a framed IPv6 prefix (used for SLAAC), a home-aware pool, and the BRG profile. The home pool subnet must be a subnet of the L2-aware inside subnet.

```
00:0c:29:00:00:10    Cleartext-Password := "letmein"
                    Alc-BRG-Id = "00:0c:29:00:00:10",
                    Alc-BRG-Profile = "brg-prof-1",
                    Framed-IPv6-Prefix = 2001:db8:101:1010::/64,
                    Alc-DMZ-address = 192.168.1.254,
                    Alc-Home-Aware-Pool =
                      "192.168.1.1/24 192.168.1.100-192.168.1.254",
                    Alc-Reserved-Addresses =
                      "sticky 00:0c:29:00:00:11 192.168.1.110",
                    Alc-Reserved-Addresses +=
                      "static 00:0c:29:00:00:1f 192.168.1.254",
                    Alc-Reserved-Addresses +=
```

```

                                "static 00:0c:29:00:00:1e 198.51.100.110",
                                Alc-Portal-Url = "http://11.11.11.11",
                                Alc-Primary-Dns = 1.1.1.1,
                                Alc-Secondary-Dns = 1.1.2.2,
                                Alc-Primary-Nbns = 2.2.1.1,
                                Alc-Secondary-Nbns = 2.2.2.2,
                                Alc-Ipv6-Primary-DNS = 2001:db8:dddd:1::1,
                                Alc-Ipv6-Secondary-DNS = 2001:db8:dddd:2::1,

00:0c:29:00:00:11      Cleartext-Password := "letmein"
                                Alc-BRG-Id = "00:0c:29:00:00:10",

00:0c:29:00:00:12      Cleartext-Password := "letmein"
                                Alc-BRG-Id = "00:0c:29:00:00:10",

00:0c:29:00:00:13      Cleartext-Password := "letmein"
                                Alc-BRG-Id = "00:0c:29:00:00:10",

00:0c:29:00:00:14      Cleartext-Password := "letmein"
                                Alc-BRG-Id = "00:0c:29:00:00:10",
                                Alc-Primary-Dns = 1.1.3.3,

---snip---
```

The RADIUS server is allowed to return a SLAAC pool name using the **Alc-SLAAC-IPv6-Pool** attribute, but in that case local address assignment needs to be configured on the vRGW. See the [ESM SLAAC Prefix Assignment via Local Address Server](#) chapter for more information.

ESM Configuration

The subscriber management policies and profiles are defined as follows. The BRG profile *brg-prof-1* provides an SLA profile, a sub-profile, a RADIUS server policy, plus the corresponding password, a RADIUS proxy server, and a DHCP pool. The DHCP pool is within the inside L2-aware NAT range defined for VPRN-1.

```

configure
  subscriber-mgmt
    ipoe-session-policy "sess-pol-SAP-MAC" create
    exit
    sla-profile "sla-prof-1" create
    exit
    sub-profile "sub-prof-1" create
      nat-policy "nat-pol-1"
    exit
    sub-ident-policy "sub-ident-DIRECT" create
      sub-profile-map
        use-direct-map-as-default
      exit
      sla-profile-map
        use-direct-map-as-default
      exit
    exit
    msap-policy "msap-pol-MSAP" create
      sub-sla-mgmt
        def-sub-id use-sap-id
        def-sub-profile "sub-prof-1"
        def-sla-profile "sla-prof-1"
        sub-ident-policy "sub-ident-DIRECT"
      exit
    exit
    brg-profile "brg-prof-1" create
```

```
        description "default BRG-profile, demo purposes"
        sla-profile-string "sla-prof-1"
        sub-profile-string "sub-prof-1"
        dhcp-pool
            subnet 192.168.1.1/24 start 192.168.1.2 end 192.168.1.254
        exit
        radius-authentication
            password letmein
            radius-server-policy "rad-serv-pol-RSP"
        exit
    exit
exit
```

The RADIUS authentication and accounting policies are defined as follows.

```
configure
  router
    radius-server
      server "radius-172.16.1.2" address 172.16.1.2 secret vsecret1 create
      accept-coa
    exit
  exit
exit

configure
  aaa
    radius-server-policy "rad-serv-pol-RSP" create
    servers
      router "Base"
      source-address 192.0.2.1
      server 1 name "radius-172.16.1.2"
    exit
  exit
exit

configure
  subscriber-mgmt
    authentication-policy "radius-AUTH" create
    description "Radius authentication policy"
    password letmein
    radius-authentication-server
      source-address 192.0.2.1
    exit
    radius-server-policy "rad-serv-pol-RSP"
  exit
  radius-accounting-policy "radius-ACCT" create
  update-interval 5
  include-radius-attribute
    mac-address
    nat-port-range
    subscriber-id
  exit
  radius-accounting-server
    source-address 192.0.2.1
    router "Base"
    server 1 address 172.16.1.2 secret vsecret1
  exit
exit
exit
```

NAT Policy configuration

The NAT policy used in support for the BRGs is *nat-pol-1*, and refers to the outside address pool defined in VPRN-2.

```
configure
  service
    nat
      nat-policy "nat-pol-1" create
      pool "nat-outside-1" router 2
    exit
  exit
exit
exit
exit
```

Operation and Verification

The following command shows the current BRG hosts. Six hosts are connected. The first host has a sticky address, the third a plain dynamic address, the fifth a public static address, and the last a private static address. The second and the fourth hosts correspond to the SLAAC hosts, and their allocation type is "not-applicable". For the static hosts, no DHCP lease information is maintained.

```
*A:BNB# show subscriber-mgmt brg brg-hosts

=====
Bridged Residential Gateway hosts
=====
Identifier          : 00:0c:29:00:00:10
MAC address         : 00:0c:29:00:00:11
IP address          : 192.168.1.110
Service             : 1 (VPRN)
Allocation type      : sticky-ip-address
Home-aware pool     : 00:0c:29:00:00:10
DHCP lease          : true
Remaining lease time : 11479
Lease start time     : 2016/09/19 20:51:22

Identifier          : 00:0c:29:00:00:10
MAC address         : 00:0c:29:00:00:11
IP address          : 2001:db8:101:1010::
Service             : 1 (VPRN)
Allocation type      : not-applicable

Identifier          : 00:0c:29:00:00:10
MAC address         : 00:0c:29:00:00:14
IP address          : 192.168.1.100
Service             : 1 (VPRN)
Allocation type      : dynamic
Home-aware pool     : 00:0c:29:00:00:10
DHCP lease          : true
Remaining lease time : 11480
Lease start time     : 2016/09/19 20:51:24

Identifier          : 00:0c:29:00:00:10
MAC address         : 00:0c:29:00:00:14
IP address          : 2001:db8:101:1010::
Service             : 1 (VPRN)
Allocation type      : not-applicable
```

```
Identifier          : 00:0c:29:00:00:10
MAC address        : 00:0c:29:00:00:1e
IP address         : 198.51.100.110
Service           : 1 (VPRN)
Allocation type    : static
Home-aware pool    : 00:0c:29:00:00:10
DHCP lease        : (Unknown)
Remaining lease time : (Unknown)
Lease start time   : N/A
```

```
Identifier          : 00:0c:29:00:00:10
MAC address        : 00:0c:29:00:00:1f
IP address         : 192.168.1.254
Service           : 1 (VPRN)
Allocation type    : static
Home-aware pool    : 00:0c:29:00:00:10
DHCP lease        : (Unknown)
Remaining lease time : (Unknown)
Lease start time   : N/A
```

```
-----
No. of BRG hosts: 6
=====
```

```
*A:BNG#
```

The following command shows the active BRGs. Because all the hosts from the previous command belong to the same BRG, only a single BRG gateway exists. The SLAAC prefix, subnet, start address, end address, DMZ address, DNS addresses for IPv4 and IPv6, NBNS-1 and NBNS-2 addresses, and IPv4 and IPv6 portal addresses are obtained from the RADIUS server. "DMZ address in use" is set to "yes" because RADIUS returned the Alc-DMZ-address, and the outside L2 NAT pool has a single port-reservation block configured (ref. VPRN-2).

```
*A:BNG# show subscriber-mgmt brg gateways
```

```
=====
Bridged Residential Gateways
=====
```

```
Identifier          : 00:0c:29:00:00:10
SLAAC prefix        : 2001:db8:101:1010::/64
Subnet              : 192.168.1.1/24
Subnet start address : 192.168.1.100
Subnet end address   : 192.168.1.254
DMZ address         : 192.168.1.254
DNS 1 v4            : 1.1.1.1
DNS 1 v6            : 2001:db8:dddd:1::1
DNS 2 v4            : 1.1.2.2
DNS 2 v6            : 2001:db8:dddd:2::1
NBNS 1              : 2.2.1.1
NBNS 2              : 2.2.2.2
DHCP lease time     : 21600
DHCP stream destination : (Not Specified)
IPv4 portal URL     : http://11.11.11.11
IPv6 portal URL     : (Not Specified)
BRG profile         : brg-prof-1
Subscriber profile   : sub-prof-1
SLA profile         : sla-prof-1
UPnP policy override : (Not Specified)
DMZ address in use   : yes
Proxy authenticated : no
Ingress IPv4 filter override : N/A
Egress IPv4 filter override : N/A
```

```
Ingress IPv6 filter override : N/A
Egress IPv6 filter override : N/A
No QoS overrides found.
No Filter rules received.
```

```
-----
No. of gateways: 1
=====
```

```
*A:BNG#
```

The following command shows the active subscribers. Only a single subscriber exists, with two SLAAC hosts (origin is SLAAC), two dynamic hosts (origin is DHCP), and two "static" hosts (origin is AAA). Also, the NAT policy used and the outside IP address and the ports are shown.

```
*A:BNG# show service active-subscribers
```

```
=====
Active Subscribers
=====
```

```
Subscriber 00:0c:29:00:00:10 (sub-prof-1)
-----
```

```
NAT Policy: nat-pol-1
Outside IP: 203.0.113.1 (vprn2)
Ports      : 1024-5119
```

```
-----
(1) SLA Profile Instance sap:[1/1/1:81] - sla:sla-prof-1
-----
```

IP Address	MAC Address	Session	Origin	Svc	Fwd
192.168.1.110	00:0c:29:00:00:11	IPoE	DHCP	1	Y
2001:db8:101:1010::/64	00:0c:29:00:00:11	IPoE	SLAAC	1	Y
192.168.1.100	00:0c:29:00:00:14	IPoE	DHCP	1	Y
2001:db8:101:1010::/64	00:0c:29:00:00:14	IPoE	SLAAC	1	Y
198.51.100.110	00:0c:29:00:00:1e	IPoE	AAA	1	Y
192.168.1.254	00:0c:29:00:00:1f	IPoE	AAA	1	Y

```
-----
Number of active subscribers : 1
-----
```

```
*A:BNG#
```

The following command shows the subscriber hierarchy for a single subscriber. This way it is apparent which host belongs to which IPoE session and on which SAP. The subscriber ID used is the MAC address of the BRG and is accompanied with the sub-profile used. The bridge ID is accompanied with the BRG profile. The NAT outside IP address is accompanied with the service number and the NAT policy.

```
*A:BNG# show service active-subscribers hierarchy subscriber "00:0c:29:00:00:10"
```

```
=====
Active Subscribers Hierarchy
=====
```

```
Hierarchy
-----
-- 00:0c:29:00:00:10 (sub-prof-1)
|   brg-id: 00:0c:29:00:00:10 - brg-profile: brg-prof-1
|   NAT Outside IP: 203.0.113.1 (vprn2) policy nat-pol-1
|
+-- sap:[1/1/1:81] - sla:sla-prof-1
|   |
|   |-- IPOE-session - mac:00:0c:29:00:00:11 - svc:1
|   |   |
|   |   |-- 192.168.1.110 - DHCP - L2Aware
|   |   |
|   |   +-- 2001:db8:101:1010::/64 - SLAAC
|   |
|   |-- IPOE-session - mac:00:0c:29:00:00:14 - svc:1
|   |   |
|   |   |-- 192.168.1.100 - DHCP - L2Aware
|   |   |
|   |   +-- 2001:db8:101:1010::/64 - SLAAC
|   |
|   |-- IPOE-session - mac:00:0c:29:00:00:1e - svc:1
|   |   |
|   |   +-- 198.51.100.110 - AAA
|   |
|   +-- IPOE-session - mac:00:0c:29:00:00:1f - svc:1
|       |
|       +-- 192.168.1.254 - AAA - L2Aware
|
=====
*A:BNG#
```

The following command shows the subscriber hosts on VPRN-1. All these hosts belong to the same subscriber, with subscriber ID 00:0c:29:00:00:10. The subscriber ID used is the MAC address of the BRG.

```
*A:BNG# show service id 1 subscriber-hosts

=====
Subscriber Host table
=====
Sap      Subscriber
IP Address
MAC Address      PPPoE-SID Origin      Fwding State
-----
[1/1/1:81]      00:0c:29:00:00:10
192.168.1.100
00:0c:29:00:00:14      N/A      DHCP      Fwding
[1/1/1:81]      00:0c:29:00:00:10
192.168.1.110
00:0c:29:00:00:11      N/A      DHCP      Fwding
[1/1/1:81]      00:0c:29:00:00:10
192.168.1.254
00:0c:29:00:00:1f      N/A      AAA      Fwding
[1/1/1:81]      00:0c:29:00:00:10
198.51.100.110
00:0c:29:00:00:1e      N/A      AAA      Fwding
[1/1/1:81]      00:0c:29:00:00:10
2001:db8:101:1010::/64
00:0c:29:00:00:11      N/A      IPoE-SLAAC      Fwding
[1/1/1:81]      00:0c:29:00:00:10
2001:db8:101:1010::/64
00:0c:29:00:00:14      N/A      IPoE-SLAAC      Fwding
-----
Number of subscriber hosts : 6
```

```
=====
*A:BNG#
```

The DHCP lease state table on VRPN-1 can then be shown with the following command. These addresses are taken from the Alc-Home-Aware-Pool as defined by the RADIUS server. The 192.168.1.110 address is a sticky address returned in the Alc-Reserved-Addresses attribute. Obviously, no entries appear for the static hosts.

```
*A:BNG# show service id 1 dhcp lease-state
```

```
=====
DHCP lease state table, service 1
=====
```

IP Address	Mac Address	Sap/Sdp Id	Remaining LeaseTime	Lease Origin	MC Stdby
192.168.1.100	00:0c:29:00:00:14	[1/1/1:81]	03h11m20s	Radius	
192.168.1.110	00:0c:29:00:00:11	[1/1/1:81]	03h11m18s	Radius	

```
-----
Number of lease states : 2
=====
```

```
*A:BNG#
```

Because a public static IP address is returned by the RADIUS server when authenticating the BRG, a 'static' BRG host was created, and the routing table is adjusted accordingly.

```
*A:BNG# show router 1 route-table
```

```
=====
Route Table (Service: 1)
=====
```

Dest Prefix[Flags] Next Hop[Interface Name]	Type	Proto	Age	Metric	Pref
10.11.11.1/32 int-DHCP	Local	Local	11h48m56s	0	0
192.168.1.0/24 sub-int-1	Local	Local	11h48m56s	0	0
198.51.100.0/24 sub-int-1	Local	Local	11h48m56s	0	0
198.51.100.110/32 [grp-int-1-1]	Remote	Sub Mgmt	11h48m41s	0	0

```
-----
No. of Routes: 4
```

```
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
=====
```

```
*A:BNG#
```

The following command shows the corresponding L2-aware hosts. Only three L2-aware hosts are created, using private IP addresses on the 'inside', and public address on the 'outside'. Traffic for these hosts passes through the ISA. Traffic for the static public address (10.0.10.110) does not undergo NAT and does not pass through the ISA.

```
*A:BNG# show service nat l2-aware-hosts
```

```
=====
Layer-2-Aware NAT hosts
=====
```



```

=====
Subscriber          : 00:0c:29:00:00:10
Inside IP address   : 192.168.1.100
-----
Policy              : nat-pol-1
Outside router      : 2
Outside IP address   : 203.0.113.1
-----
Subscriber          : 00:0c:29:00:00:10
Inside IP address   : 192.168.1.110
-----
Policy              : nat-pol-1
Outside router      : 2
Outside IP address   : 203.0.113.1
-----
Subscriber          : 00:0c:29:00:00:10
Inside IP address   : 192.168.1.254
-----
Policy              : nat-pol-1
Outside router      : 2
Outside IP address   : 203.0.113.1
-----
No. of hosts: 3
=====
*A:BN#

```

The following command shows the corresponding IPoE sessions. There is one session per device/MAC-address, so there are four IPoE sessions for this subscriber.

```

*A:BN# show service id 1 ipoe session

=====
IPoE sessions for svc-id 1
=====
Sap Id      Subscriber-Id      Mac Address      Up Time      MC-Stdby
  [CircuitID] | [RemoteID]
-----
[1/1/1:81]  00:0c:29:00:00:11    0d 11:48:42
00:0c:29:00:00:10
[1/1/1:81]  00:0c:29:00:00:14    0d 11:48:40
00:0c:29:00:00:10
[1/1/1:81]  00:0c:29:00:00:1e    0d 11:48:42
00:0c:29:00:00:10
[1/1/1:81]  00:0c:29:00:00:1f    0d 11:48:42
00:0c:29:00:00:10
-----
CID | RID displayed when included in session-key
Number of sessions : 4
=====
*A:BN#

```

With this single subscriber connected, the following command shows the managed hosts that are created.

```

*A:BN# show service id 1 managed-hosts type aaa

=====
Managed aaa hosts
=====
IP address      MAC address
-----

```

```
192.168.1.254/32          00:0c:29:00:00:1f
198.51.100.110/32        00:0c:29:00:00:1e
-----
No. of Managed hosts: 2
=====
*A: BNG#
```

Conclusion

Using a BRG in the home network instead of a full-fledged L3 RGW offers network operators a view on the IP addresses and MAC addresses used by home devices, enabling them to provide per-device service offerings. Integrating BRGs in their networks can help the operators to optimize the revenue stream.

WiFi Aggregation and Offload — Basic Open SSID

This chapter provides information about WiFi Aggregation and Offload — Basic Open SSID.

Topics in this chapter include:

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

Applicability

The information and configuration in this chapter are based on SR OS Release 12.0.R5.

Overview

WiFi Aggregation and Offload functionality for the 7750 SR is supported on SR OS 10.0.R3 or later. The functionality includes a RADIUS proxy server with a RADIUS proxy cache and support for soft-GRE tunnels.

WLAN-GW subscribers are implemented using Enhanced Subscriber Management (ESM) on the Control Processing Module (CPM), to benefit from the extensive ESM features available on the 7750 SR platform. Many different WiFi Offload configurations are possible, with the two most versatile configurations being open and secure Service Set Identifier (SSID).

This configuration should be used as a starting point for operators who need to offer an open SSID, where any client can connect to an Access Point (AP) and obtain an IP address without authentication. In most cases, operators want users to go through an authentication process before allowing full Internet access using the open SSID; therefore, this configuration also includes a web portal.

IP address assignment and Internet connectivity can be achieved using various methods in SR OS. In this configuration, a local DHCP server provides IP addresses to the User Equipment (UE) and routing to the Internet is performed using Global Routing Table (GRT) leaking.

Several considerations typically affect the choice of a WiFi Offload solution:

- Access can be free or paid.
- Equipment can be preconfigured or users can bring their own WiFi device.

When there is no pre-existing subscription, an open SSID is the most obvious solution. To provide a paywall or to have the user acknowledge certain terms of use due to legal reasons, a web portal may also be required.

When a web portal is implemented, users who connect to the open SSID which are not yet authenticated have all their web traffic redirected to the web portal landing page. This is performed using an http-redirect filter applied to the initial (limited) Service Level Agreement (SLA) profile assigned to the UE. Typically, the operating system of the UE will detect the presence of the web portal and automatically open the login

page for the user. When the user logs in, the web portal sends a RADIUS Change of Authorization (CoA) request to the WLAN-GW, changing the SLA profile to one that does not contain an http-redirect filter.

Besides authentication, a major consideration is the method used to achieve Internet connectivity. Will the users require public addresses or are private addresses sufficient? In case few public IP addresses are available, private IP addresses can be assigned to UEs and the WLAN-GW can perform a Network Address Translation (NAT) function. If public routable IP addresses can be made available to all UEs, traffic from the UEs can be routed by the WLAN-GW to the Internet.

When a UE connects to an open SSID (as shown in [Figure 172: Call Flow for Open SSID](#)), typically the UE attempts to obtain an IP address using Dynamic Host Configuration Protocol (DHCP). The WLAN-GW can serve as a DHCP relay or proxy and may obtain the IP address from an external source, or use a local DHCP server function. A DHCP Discover or Request packet from a UE will trigger a form of authentication where the WLAN-GW requests information about the UE, such as SLA profile or DHCP local pool name. This authentication is separate from the web portal authentication and occurs immediately when a UE connects.

In summary:

- DHCP Discover triggering RADIUS authentication
- DHCP completes and UE has SLA profile with limited access
- UE logs into a web portal
- Successful login causes the portal to send a RADIUS CoA which assigns an SLA profile with full access

```
sequenceDiagram
    participant UE
    participant AP
    participant WLAN_GW as WLAN GW
    participant AAA
    participant Portal
    participant Internet

    Note over UE, AP: UE is Connected
    UE->>AP: DHCP Discover (chaddr=UE-MAC)
    AP->>WLAN_GW: GRE (DHCP Discover (chaddr=UE-MAC))
    WLAN_GW->>AAA: RADIUS Access-Request (NAS-ID)
    AAA->>WLAN_GW: RADIUS Access-Accept (Sub-ID, Sub-Profile, SLA-Profile, Framed-IP-Address, Framed-IP-Netmask)
    WLAN_GW->>AP: GRE (DHCP Offer (chaddr=UE-MAC, your-ip=UE-IP, Subnet-Mask, Router=WLAN-GW-IP, Lease-Time))
    AP->>UE: DHCP Offer (chaddr=UE-MAC, your-ip=UE-IP, Subnet-Mask, Router=WLAN-GW-IP, Lease-Time)
    UE->>AP: DHCP Request (chaddr=UE-MAC, Requested-IP-Address=UE-IP)
    AP->>WLAN_GW: GRE (DHCP Request (chaddr=UE-MAC, Requested-IP-Address=UE-IP))
    Note over WLAN_GW: ESM subscriber created on CPM
    WLAN_GW->>AP: GRE (DHCP Ack (chaddr=UE-MAC, your-ip=UE-IP, Subnet-Mask, Router=WLAN-GW-IP, Lease-Time))
    AP->>UE: DHCP Ack (chaddr=UE-MAC, your-ip=UE-IP, Subnet-Mask, Router=WLAN-GW-IP, Lease-Time)
    UE->>AP: HTTP GET (Internet URL)
    AP->>WLAN_GW: GRE (HTTP GET (Internet URL))
    WLAN_GW->>AP: GRE (HTTP Redirect/302 (Portal URL))
    AP->>UE: HTTP Redirect/302 (Portal URL)
    UE->>Portal: HTTP GET (Portal URL)
    Portal->>Portal: Portal authentication
    Portal->>WLAN_GW: RADIUS CoA (SLA-Profile)
    Note over WLAN_GW: SLA-Profile updated, redirect filter removed
    WLAN_GW->>UE: HTTP GET (Internet URL)
    Note over UE: Disconnect
```

The SR OS is flexible in allowing the operator to separate the various WiFi Offload functions between different routing instances. All functions can be configured in the same routing instance, or as shown in [Figure 173: WiFi Offload Scenario with Open SSID and Local DHCP Server](#), the connectivity to the APs (and soft-GRE tunnels) can be provided in one Virtual Private Routed Network (VPRN), the users can be instantiated in another VPRN, and Authentication, Authorization and Accounting (AAA) access can be


provided in yet another routing instance (in this case, the Base router). This clear separation of functions can enhance security; for example, by separating user traffic from authentication traffic.

Configuration

The WiFi offload scenario shown in [Figure 173: WiFi Offload Scenario with Open SSID and Local DHCP Server](#) has following characteristics:

- Open SSID with web portal authentication
- Local breakout to Internet using GRT leaking, routing through Base routing instance
- Same private IP address assigned to all UEs, with L2-aware NAT
- AP access in VPRN 1000
- UEs terminated in VPRN 2005
- Local DHCP server assigning public IP addresses

Figure 173: WiFi Offload Scenario with Open SSID and Local DHCP Server



WLAN-GW

Note that the uplink interface, Interior Gateway Protocol (IGP), and system configuration is outside the scope of this document.

The following card and Media Dependent Adapter (MDA) configuration shows only the WLAN-Input/Output Module (IOM). An IOM card containing two Multi-Service Integrated Service Adapter (MS-ISA) cards provides the WLAN-GW functionality. The MDA type for the ISA cards is isa-bb, the same type that is used for NAT.

```
*A:WLAN-GW# /configure card 2
*A:WLAN-GW>config>card# info
-----
card-type iom3-xp-b
mda 1
mda-type isa-bb
no shutdown
```

```
exit
mda 2
    mda-type isa-bb
    no shutdown
exit
no shutdown
-----
```

The following ISA configuration defines a wlan-gw-group referencing the IOM in slot two which hosts the two MS-ISA cards and providing the WLAN-GW functions.

```
A:WLAN-GW# /configure isa
A:WLAN-GW>config>isa# info
-----
    wlan-gw-group 1 create
        active-iom-limit 1
        iom 2
        no shutdown
    exit
-----
```

The following is a RADIUS server configuration, where the secret must match the secret configured on the external RADIUS server. The accept-coa option must be configured to allow the change of SLA profile by the web portal using a CoA request.

```
*A:WLAN-GW# /configure router radius-server
*A:WLAN-GW>config>router>radius-server# info
-----
    server "Server2" address 10.43.186.2 secret "zmLYVgt8UOLypJamceNSSHDWbZproq7Y"
    hash2 create
        accept-coa
    exit
-----
```

The following AAA configuration contains a RADIUS server policy used in the authentication policy. The source address must match the IP address configured for this client on the RADIUS server.

```
*A:WLAN-GW# /configure aaa
*A:WLAN-GW>config>aaa# info
-----
    radius-server-policy "RS_5" create
        servers
            router "Base"
            source-address 10.10.10.165
            server 1 name "Server2"
        exit
    exit
-----
```

The following policy configuration is used for exporting routes so that they are reachable by the public network. These policies are used for exporting UE routes in subsequent configuration sections.

```
*A:WLAN-GW# /configure router policy-options
*A:WLAN-GW>config>router>policy-options# info
-----
    prefix-list "WiFi-clients"
        prefix 10.165.0.0/16 longer
    exit
    policy-statement "WiFi-clients"
        entry 10
-----
```

```

        from
        prefix-list "WiFi-clients"
        exit
        action accept
        exit
    exit
exit
-----

```

The uplink network configuration is outside the scope of this document. However, note that the IGP (here ISIS) must be aware of the UE addresses so that they are accessible from the Internet.

```

*A:WLAN-GW# /configure router isis
*A:WLAN-GW>config>router>isis# info
-----
    export "WiFi-clients"
-----

```

The following IP filter redirects all HTTP traffic to the web portal. The filter should also allow DNS and potentially other traffic, so the entry that allows TCP port 80 traffic to the web portal address must be placed before the entry that redirects all traffic to that portal; otherwise, there will be a redirect loop.

The HTTP redirect URL also includes a parameter that provides the MAC address of the UE to the web portal. In this configuration, either \$MAC or \$SUB can be used since both variables contain the MAC address of the UE. The web portal can reply with a CoA request specifying this particular UE MAC as the Subscriber ID after successful login. The URL also returns the IP address of the WLAN-GW to the web portal, so that the portal knows which WLAN-GW to send the CoA request to.

```

*A:WLAN-GW# /configure filter
*A:WLAN-GW>config>filter# info
-----
    ip-filter 2005 create
        default-action forward
        entry 70 create
            match protocol udp
            dst-port eq 53
        exit
        action forward
    exit
    entry 80 create
        match protocol icmp
    exit
    action forward
    exit
    entry 90 create
        match protocol tcp
        dst-ip 10.43.186.2/32
        dst-port eq 80
    exit
    action forward
    exit
    entry 100 create
        match protocol tcp
        dst-port eq 80
    exit
    action http-redirect "http://portal2.3ls.net/portal-no-login.php?gw=
10.10.10.165&mac=$SUB"
    exit
exit

```


The following is a subscriber management configuration, with the RADIUS authentication policy used to authenticate DHCP requests, including the accept-authorization-change option to allow for SLA profile change after portal authentication. This DHCP authentication request also sends the NAS ID attribute that allows the RADIUS server to match on the configuration for this particular SSID. All UEs will be authenticated with their MAC address as user name, and *alcatel* as their password (any DHCP request will result in a successful authentication).

Two SLA profiles are required: profile SLAP_5_portal is initially used for each UE and refers to the portal redirect filter, while profile SLAP is applied using a CoA request after the user successfully authenticates on the web portal. A subscriber identity policy is also required.

```
configure subscriber-mgmt
  authentication-policy "WiFi-165-5-auth-policy" create
    password alcatel
    accept-authorization-change
    include-radius-attribute
      nas-identifier
    exit
    radius-server-policy "RS_5"
  exit
  sla-profile "SLAP" create
  exit
  sla-profile "SLAP_5_portal" create
    ingress
      ip-filter 2005
    exit
  exit
  sub-profile "SUBP" create
  exit
  sub-ident-policy "SIP" create
    sub-profile-map
      use-direct-map-as-default
    exit
    sla-profile-map
      use-direct-map-as-default
    exit
  exit
exit
```

The following VPRN 1000 configuration contains the interface to the AP, and has GRT lookup with export-grt configured to allow APs to be managed from the Base routing instance.

```
*A:WLAN-GW# /configure service vprn 1000
*A:WLAN-GW>config>service>vprn# info
-----
    route-distinguisher 65400:1000
    interface "toAP3" create
      address 10.1.3.1/24
      sap 1/1/10 create
    exit
  exit
  grt-lookup
    enable-grt
      static-route 0.0.0.0/0 grt
    exit
    export-grt "WiFi-APs"
  exit
  no shutdown
-----
```

VPRN 2005 is used for UE termination and contains:

- A local DHCP server with a single pool of addresses that are assigned to UEs.
- A loopback interface used by the DHCP server.
- A subscriber interface and group interface of type **wlangw** (called softgre prior to Release 12.0).
- Subscriber parameters.
- The authentication policy, which will run each time a UE requests a DHCP address.
- The host-connectivity-verify function, which periodically checks the presence of UEs and quickly removes disconnected UEs even before their DHCP lease expires; the WLAN-GW has no other way of knowing when a UE has disconnected from the AP.
- The wlan-gw CLI-node (called soft-gre prior to Release 12.0), including the wlan-gw GRE tunnel end-point address, and the routing instance where AP traffic is terminated, the ISA WLAN-GW group, and mobility parameters, which allow the UE state to be kept if the UE moves between two APs broadcasting the same SSID.
- GRT lookup with export-grt configured to allow UE traffic to be routed to the Internet.

```
*A:WLAN-GW# /configure service vprn 2005
*A:WLAN-GW>config>service>vprn# info
-----
description "WiFi-165-5 Open SSID"
dhcp
  local-dhcp-server "local_dhcp_2005" create
  use-pool-from-client
  pool "pool1" create
  max-lease-time hrs 1
  options
    dns-server 10.43.186.43
  exit
  subnet 10.165.5.0/25 create
  options
    subnet-mask 255.255.255.128
    default-router 10.165.5.1
  exit
  address-range 10.165.5.2 10.165.5.99
  exit
exit
no shutdown
exit
route-distinguisher 65400:2005
interface "dhcp-server" create
  address 10.0.0.1/24
  local-dhcp-server "local_dhcp_2005"
  loopback
exit
subscriber-interface "SI5" create
  address 10.165.5.1/24
  group-interface "GI5" wlangw create
  sap-parameters
    sub-sla-mgmt
      def-sla-profile "SLAP_5_portal"
      def-sub-profile "SUBP"
      sub-ident-policy "SIP"
    exit
  exit
dhcp
  option
    action replace
    circuit-id
```

```

        no remote-id
        vendor-specific-option
        pool-name
    exit
exit
server 10.0.0.1
trusted
lease-populate 10000
gi-address 10.165.5.1
no shutdown
exit
authentication-policy "WiFi-165-5-auth-policy"
host-connectivity-verify interval 5 action remove
wlan-gw
    gw-address 192.5.5.5
    mobility
        trigger data iapp
    exit
    router 1000
    wlan-gw-group 1
    no shutdown
exit
exit
exit
grt-lookup
    enable-grt
    exit
    export-grt "WiFi-clients"
exit
no shutdown

```

Freeradius

This simple default configuration section matches on any host. During the DHCP authentication phase, RADIUS returns the DHCP pool name *pool1* informing the WLAN-GW DHCP server which pool to assign the UE IP address from:

```

/etc/freeradius/users
DEFAULT      Auth-Type := Local, User-Password := "alcatel", user-name=~"
              Alc-Subsc-ID-Str = "%{User-Name}",
              Framed-Pool = "pool1",

```

In */etc/freeradius/clients.conf*, the secret must match the secret configured in the WLAN-GW RADIUS server configuration.

```

client 10.10.10.165 {
    secret      = alcatel
    shortname   = WLAN-GW
}

```

The RADIUS CoA sent during successful portal login allows this UE full access, by applying SLA profile SLAP which does not have an http-redirect filter.

```

echo "ALC-Subsc-Id-Str='68:7f:74:8b:3d:d7',ALC-Subsc-Prof-Str='SUBP_5',ALC-SLA-Prof-Str=
'SLAP',Alc-Primary-Dns = '10.43.186.43'" | /usr/bin/radclient -x -r 1 -t 2 '10.10.10.165' coa
'alcatel'

```

Access Points

At a minimum, the following must be configured on the Access Point:

- IP address 10.1.3.10/24
- Default route to 10.1.3.1
- Open SSID WiFi-165-5 mapped to VLAN 50
- Soft-GRE tunnel with destination 192.5.5.5, with VLAN 50 mapped to this tunnel

Show Commands

The following show commands reflect the status of the router after the UE has connected and obtained an IP address using DHCP.

The following output displays the UEs presently connected.

```
*A:WLAN-GW# show subscriber-mgmt wlan-gw ue
=====
User Equipments
=====
MAC address           : 68:7f:74:8b:3d:d7
-----
VLAN Q-tag            : 50
MPLS label            : (Not Specified)
Tunnel router         : 1000
Tunnel remote IP address : 10.1.3.10
Tunnel local IP address  : 192.5.5.5
Retail service        : N/A
SSID                  : (Not Specified)
Previous Access Point IP : (Not Specified)
IMSI                  : (Not Specified)
Last move time        : 2014/09/22 10:47:58
-----
No. of UE: 1
=====
```

The DHCP lease information indicates that the address was assigned by the local DHCP server.

```
*A:WLAN-GW# show service id 2005 dhcp lease-state
=====
DHCP lease state table, service 2005
=====
IP Address      Mac Address      Sap/Sdp Id      Remaining Lease   MC
                LeaseTime    Origin          Stdby
-----
10.165.5.2      68:7f:74:8b:3d:d7 [2/1/nat-out-ip:20* 00h59m27s  DHCP
-----
Number of lease states : 1
=====
* indicates that the corresponding row element may have been truncated.
```

DHCP statistics can be displayed using following command.

```
*A:WLAN-GW# show service id 2005 dhcp statistics
=====
```

DHCP Global Statistics, service 2005

```
=====
Rx Packets                : 2
Tx Packets                : 2
Rx Malformed Packets     : 0
Rx Untrusted Packets     : 0
Client Packets Discarded  : 0
Client Packets Relayed    : 2
Client Packets Snooped    : 0
Client Packets Proxied (RADIUS) : 0
Client Packets Proxied (User-Db) : 0
Client Packets Proxied (Lease-Split) : 0
Server Packets Discarded  : 0
Server Packets Relayed    : 2
Server Packets Snooped    : 0
DHCP RELEASEs Spoofed    : 0
DHCP FORCERENEWs Spoofed : 0
=====
```

The route table for the routing instance where UEs are terminated shows an entry for the UE.

*A:WLAN-GW# show router 2005 route-table

```
=====
Route Table (Service: 2005)
=====
Dest Prefix[Flags]      Type  Proto  Age      Pref
Next Hop[Interface Name] Metric
-----
10.0.0.0/24             Local  Local  00h32m12s  0
                        dhcp-server
                        0
10.165.5.0/24           Local  Local  00h27m20s  0
                        SI5
                        0
10.165.5.2/32           Remote Sub Mgmt 00h00m33s  0
                        [GI5]
                        0
-----
No. of Routes: 3
Flags: n = Number of times nexthop is repeated
      B = BGP backup route available
      L = LFA nexthop available
      S = Sticky ECMP requested
=====
```

The active subscribers view shows the initial limited SLA profile SLAP_5_portal before the user has logged in to the portal.

*A:WLAN-GW>config>service>vprn# show service active-subscribers

```
=====
Active Subscribers
=====
Subscriber 68:7f:74:8b:3d:d7 (SUBP)
-----
(1) SLA Profile Instance sap:[2/1/nat-out-ip:2049.3] - sla:SLAP_5_portal
-----
IP Address      MAC Address      PPPoE-SID Origin
-----
10.165.5.2      68:7f:74:8b:3d:d7 N/A      DHCP
-----
```

```
Number of active subscribers : 1
-----
```

The following output shows the active subscribers view after the user has logged in and the SLA profile has been updated with the unrestricted SLA profile SLAP.

```
*A:WLAN-GW# show service active-subscribers
=====
Active Subscribers
=====
Subscriber 68:7f:74:8b:3d:d7 (SUBP_5)
-----
(1) SLA Profile Instance sap:[2/1/nat-out-ip:2049.3] - sla:SLAP
-----
IP Address          MAC Address      PPPoE-SID Origin
-----
10.165.5.2          68:7f:74:8b:3d:d7 N/A      DHCP
-----
Number of active subscribers : 1
-----
```

The following output shows the RADIUS statistics for the DHCP authentication.

```
*A:WLAN-GW# show aaa radius-server-policy "RS_5" statistics
=====
RADIUS server policy "RS_5" statistics
=====
Tx transaction requests          : 1
Rx transaction responses         : 1
Transaction requests timed out   : 0
Transaction requests send failed : 0
Packet retries                   : 0
Transaction requests send rejected : 0
Authentication requests failed   : 0
Accounting requests failed       : 0
Ratio of access-reject over auth responses : 0%
Transaction success ratio        : 100%
Transaction failure ratio        : 0%
Statistics last reset at         : n/a

Server 1 "Server2" address 10.43.186.2 auth-port 1812 acct-port 1813
-----
Tx request packets              : 1
Rx response packets             : 1
Request packets timed out       : 0
Request packets send failed     : 0
Request packets send failed (overload) : 0
Request packets waiting for reply : 0
Response packets with invalid authenticator : 0
Response packets with invalid msg authenticator : 0
Authentication packets failed   : 0
Accounting packets failed       : 0
Avg auth response delay (10 100 1K 10K) in ms : 7.24 7.24 7.24 7.24
Avg acct response delay (10 100 1K 10K) in ms : n/a
Statistics last reset at         : n/a
```

The following output shows the CoA statistics after portal authentication.

```
*A:WLAN-GW# show subscriber-mgmt authentication coa-statistics
=====
Radius Notify Statistics      Change-Of-Authorization      Disconnect-Messages
=====
Requests Received            1                             0
Requests Accepted            1                             0
Requests Rejected            0                             0
Requests Dropped             0                             0
  No Auth Policy found       0                             0
  Invalid message            0                             0
  Out of resources           0                             0
  Authentication failure     0                             0
=====
```

Debug

The following is a complete debug of a UE connecting and logging in to the portal. Shortly after logging in, the UE disconnects from the SSID and the subscriber is removed by host-connectivity-verify.

The following debug configuration applies:

```
debug
  router "Base"
    radius
      packet-type authentication accounting coa
      detail-level medium
    exit
  exit
  router "2005"
    ip
      dhcp
        detail-level medium
        mode egr-ingr-and-dropped
      exit
    exit
    local-dhcp-server "local_dhcp_2005"
      detail-level medium
      mode dropped-only
    exit
  exit
  service
    id 2005
      host-connectivity-verify
      mac 68:7f:74:8b:3d:d7
    exit
  exit
exit
```

The WLAN-GW is notified of the UE after receiving the first DHCP packet.

```
1 2014/09/22 10:47:58.46 EDT MINOR: DEBUG #2001 vprn2005 PIP
"PIP: DHCP
instance 5 (2005), interface index 3 (GI5),
received DHCP Boot Request on Interface GI5 (2/1/nat-out-ip:2049.3) Port 67
```

```
H/W Type: Ethernet(10Mb)  H/W Address Length: 6
ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
siaddr: 0.0.0.0           giaddr: 0.0.0.0
chaddr: 68:7f:74:8b:3d:d7  xid: 0x8c0fc642

DHCP options:
[53] Message type: Discover
[61] Client id: (hex) 01 68 7f 74 8b 3d d7
[50] Requested IP addr: 10.165.5.2
[12] Host name: W81VM
[60] Class id: MSFT 5.0
[55] Param request list: len = 13
      1 Subnet mask
     15 Domain name
      3 Router
      6 Domain name server
     44 NETBIOS name server
     46 NETBIOS type
     47 NETBIOS scope
     31 Router discovery
     33 Static route
    121 Unknown option
    249 Unknown option
    252 Unknown option
     43 Vendor specific
[255] End
"
```

DHCP triggers sending the RADIUS Access-Request.

```
2 2014/09/22 10:47:58.48 EDT MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Access-Request(1) 10.43.186.2:1812 id 1 len 79 vrid 1 pol RS_5
    USER NAME [1] 17 68:7f:74:8b:3d:d7
    PASSWORD [2] 16 IyDg9t17sGTbfr/6h0Bs1U
    NAS IP ADDRESS [4] 4 10.10.10.165
    NAS IDENTIFIER [32] 14 WLAN-GW
"
```

The UE authentication request is always accepted and the Access-Accept message contains the required subscriber management and IP parameters, in this case, at least the subscriber ID string as well as the pool name to be used by the local DHCP server.

```
3 2014/09/22 10:47:58.50 EDT MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Access-Accept(2) id 1 len 52 from 10.43.186.2:1812 vrid 1 pol RS_5
    VSA [26] 19 Alcatel(6527)
    SUBSC ID STR [11] 17 68:7f:74:8b:3d:d7
    FRAMED POOL [88] 5 pool1
"
```

The DHCP request is transmitted to the local DHCP server, which assigns the IP address to the UE.

```
4 2014/09/22 10:47:58.50 EDT MINOR: DEBUG #2001 vprn2005 PIP
"PIP: DHCP
instance 5 (2005),
  transmitted DHCP Boot Request to 10.0.0.1 Port 67

H/W Type: Ethernet(10Mb)  H/W Address Length: 6
ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
```



```
siaddr: 0.0.0.0      giaddr: 10.165.5.1
chaddr: 68:7f:74:8b:3d:d7  xid: 0x8c0fc642

DHCP options:
[82] Relay agent information: len = 54
    [1] Circuit-id: WLAN-GW|2005|GI5|2/1/nat-out-ip:2049.3
    [9] Vendor-Specific info: len = 12
        Enterprise [6527] : len = 7
    [13] dhcpPool: pool1
[53] Message type: Discover
[61] Client id: (hex) 01 68 7f 74 8b 3d d7
[50] Requested IP addr: 10.165.5.2
[12] Host name: W81VM
[60] Class id: MSFT 5.0
[55] Param request list: len = 13
    1 Subnet mask
    15 Domain name
    3 Router
    6 Domain name server
    44 NETBIOS name server
    46 NETBIOS type
    47 NETBIOS scope
    31 Router discovery
    33 Static route
    121 Unknown option
    249 Unknown option
    252 Unknown option
    43 Vendor specific
[255] End
"

5 2014/09/22 10:47:58.50 EDT MINOR: DEBUG #2001 vprn2005 PIP
"PIP: DHCP
instance 5 (2005),
received DHCP Boot Reply on 10.0.0.1 Port 67

H/W Type: Ethernet(10Mb) H/W Address Length: 6
ciaddr: 0.0.0.0      yiaddr: 10.165.5.2
siaddr: 10.0.0.1      giaddr: 10.165.5.1
chaddr: 68:7f:74:8b:3d:d7  xid: 0x8c0fc642

DHCP options:
[82] Relay agent information: len = 54
    [1] Circuit-id: WLAN-GW|2005|GI5|2/1/nat-out-ip:2049.3
    [9] Vendor-Specific info: len = 12
        Enterprise [6527] : len = 7
    [13] dhcpPool: pool1
[53] Message type: Offer
[54] DHCP server addr: 10.0.0.1
[51] Lease time: 3600
[1] Subnet mask: 255.255.255.128
[3] Router: 10.165.5.1
[6] Domain name server: 10.43.186.43
[12] Host name: W81VM
[60] Class id: MSFT 5.0
[255] End
"

6 2014/09/22 10:47:58.52 EDT MINOR: DEBUG #2001 vprn2005 PIP
"PIP: DHCP
instance 5 (2005), interface index 3 (GI5),
transmitted DHCP Boot Reply to Interface GI5 (2/1/nat-out-ip:2049.3) Port 68

H/W Type: Ethernet(10Mb) H/W Address Length: 6
```

```
ciaddr: 0.0.0.0          yiaddr: 10.165.5.2
siaddr: 10.0.0.1        giaddr: 10.165.5.1
chaddr: 68:7f:74:8b:3d:d7  xid: 0x8c0fc642

DHCP options:
[53] Message type: Offer
[54] DHCP server addr: 10.0.0.1
[51] Lease time: 3600
[1] Subnet mask: 255.255.255.128
[3] Router: 10.165.5.1
[6] Domain name server: 10.43.186.43
[12] Host name: W81VM
[60] Class id: MSFT 5.0
[255] End
"

7 2014/09/22 10:47:58.69 EDT MINOR: DEBUG #2001 vprn2005 PIP
"PIP: DHCP
instance 5 (2005), interface index 3 (GI5),
received DHCP Boot Request on Interface GI5 (2/1/nat-out-ip:2049.3) Port 67

H/W Type: Ethernet(10Mb) H/W Address Length: 6
ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
siaddr: 0.0.0.0          giaddr: 0.0.0.0
chaddr: 68:7f:74:8b:3d:d7  xid: 0x8c0fc642

DHCP options:
[53] Message type: Request
[61] Client id: (hex) 01 68 7f 74 8b 3d d7
[50] Requested IP addr: 10.165.5.2
[54] DHCP server addr: 10.0.0.1
[12] Host name: W81VM
[81] client FQDN: rcode1: 0, rcode2: 0, domain name = (hex) 00 57 38 31 56
4d
[60] Class id: MSFT 5.0
[55] Param request list: len = 13
    1 Subnet mask
    15 Domain name
    3 Router
    6 Domain name server
    44 NETBIOS name server
    46 NETBIOS type
    47 NETBIOS scope
    31 Router discovery
    33 Static route
    121 Unknown option
    249 Unknown option
    252 Unknown option
    43 Vendor specific
[255] End
"

8 2014/09/22 10:47:58.69 EDT MINOR: DEBUG #2001 vprn2005 PIP
"PIP: DHCP
instance 5 (2005),
transmitted DHCP Boot Request to 10.0.0.1 Port 67

H/W Type: Ethernet(10Mb) H/W Address Length: 6
ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
siaddr: 0.0.0.0          giaddr: 10.165.5.1
chaddr: 68:7f:74:8b:3d:d7  xid: 0x8c0fc642

DHCP options:
[82] Relay agent information: len = 54
```

```
[1] Circuit-id: WLAN-GW|2005|GI5|2/1/nat-out-ip:2049.3
[9] Vendor-Specific info: len = 12
    Enterprise [6527] : len = 7
    [13] dhcpPool: pool1
[53] Message type: Request
[61] Client id: (hex) 01 68 7f 74 8b 3d d7
[50] Requested IP addr: 10.165.5.2
[54] DHCP server addr: 10.0.0.1
[12] Host name: W81VM
[81] client FQDN: rcode1: 0, rcode2: 0, domain name = (hex) 00 57 38 31 56
4d
[60] Class id: MSFT 5.0
[55] Param request list: len = 13
    1 Subnet mask
    15 Domain name
    3 Router
    6 Domain name server
    44 NETBIOS name server
    46 NETBIOS type
    47 NETBIOS scope
    31 Router discovery
    33 Static route
    121 Unknown option
    249 Unknown option
    252 Unknown option
    43 Vendor specific
[255] End
"

9 2014/09/22 10:47:58.69 EDT MINOR: DEBUG #2001 vprn2005 PIP
"PIP: DHCP
instance 5 (2005),
received DHCP Boot Reply on 10.0.0.1 Port 67

H/W Type: Ethernet(10Mb) H/W Address Length: 6
ciaddr: 0.0.0.0 yiaddr: 10.165.5.2
siaddr: 10.0.0.1 giaddr: 10.165.5.1
chaddr: 68:7f:74:8b:3d:d7 xid: 0x8c0fc642

DHCP options:
[82] Relay agent information: len = 54
    [1] Circuit-id: WLAN-GW|2005|GI5|2/1/nat-out-ip:2049.3
    [9] Vendor-Specific info: len = 12
        Enterprise [6527] : len = 7
        [13] dhcpPool: pool1
[53] Message type: Ack
[54] DHCP server addr: 10.0.0.1
[51] Lease time: 3600
[1] Subnet mask: 255.255.255.128
[3] Router: 10.165.5.1
[6] Domain name server: 10.43.186.43
[12] Host name: W81VM
[81] client FQDN: rcode1: 0, rcode2: 0, domain name = (hex) 00 57 38 31 56
4d
[60] Class id: MSFT 5.0
[255] End
"

10 2014/09/22 10:47:58.69 EDT MINOR: DEBUG #2001 vprn2005 PIP
"PIP: DHCP
instance 5 (2005), interface index 3 (GI5),
transmitted DHCP Boot Reply to Interface GI5 (2/1/nat-out-ip:2049.3) Port 68

H/W Type: Ethernet(10Mb) H/W Address Length: 6
```

```

ciaddr: 0.0.0.0          yiaddr: 10.165.5.2
siaddr: 10.0.0.1        giaddr: 10.165.5.1
chaddr: 68:7f:74:8b:3d:d7  xid: 0x8c0fc642

DHCP options:
[53] Message type: Ack
[54] DHCP server addr: 10.0.0.1
[51] Lease time: 3600
[1] Subnet mask: 255.255.255.128
[3] Router: 10.165.5.1
[6] Domain name server: 10.43.186.43
[12] Host name: W81VM
[81] client FQDN: rcode1: 0, rcode2: 0, domain name = (hex) 00 57 38 31 56
4d
[60] Class id: MSFT 5.0
[255] End
"

```

At this point in the configuration, the UE has network connectivity but all HTTP traffic is redirected to the web portal, as configured in the IP filter included in the initial SLA profile.

After web portal authentication, the WLAN-GW receives a RADIUS CoA for this subscriber, which includes the new unrestricted SLA profile SLAP.

```

11 2014/09/22 10:48:12.54 EDT MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Change of Authorization(43) id 162 len 71 from 10.43.186.2:55255 vrid 1
  VSA [26] 19 Alcatel(6527)
    SUBSC ID STR [11] 17 68:7f:74:8b:3d:d7
  VSA [26] 8 Alcatel(6527)
    SUBSC PROF STR [12] 6 SUBP_5
  VSA [26] 6 Alcatel(6527)
    SLA PROF STR [13] 4 SLAP
"

12 2014/09/22 10:48:12.54 EDT MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Change of Authorization Ack(44) 10.43.186.2:55255 id 162 len 20 vrid 1
"

```

The host has accessed a few web sites, then disconnected from the SSID, which is not known by the WLAN-GW. After 5 minutes of inactivity, host-connectivity-verify removes the subscriber and the DHCP lease is cleared.

```

13 2014/09/22 10:48:58.90 EDT MINOR: DEBUG #2001 vprn2005 SHCV
"SHCV: Periodic Check
  2/1/nat-out-ip:2049.3
  DHCP lease state 10.165.5.2 68:7f:74:8b:3d:d7"

14 2014/09/22 10:49:08.90 EDT MINOR: DEBUG #2001 vprn2005 SHCV
"SHCV: Periodic Check
  2/1/nat-out-ip:2049.3
  DHCP lease state 10.165.5.2 68:7f:74:8b:3d:d7"

15 2014/09/22 10:49:18.90 EDT MINOR: DEBUG #2001 vprn2005 SHCV
"SHCV: Periodic Check
  2/1/nat-out-ip:2049.3
  DHCP lease state 10.165.5.2 68:7f:74:8b:3d:d7"

16 2014/09/22 10:49:28.90 EDT MINOR: DEBUG #2001 vprn2005 SHCV
"SHCV: Connectivity Lost
  2/1/nat-out-ip:2049.3
"

```

```
DHCP lease state 10.165.5.2 68:7f:74:8b:3d:d7"
17 2014/09/22 10:49:30.00 EDT MINOR: DEBUG #2001 vprn2005 PIP
"PIP: DHCP
instance 5 (2005),
  transmitted DHCP Boot Request to 10.0.0.1 Port 68

H/W Type: Ethernet(10Mb)  H/W Address Length: 6
ciaddr: 10.165.5.2        yiaddr: 0.0.0.0
siaddr: 0.0.0.0          giaddr: 0.0.0.0
chaddr: 68:7f:74:8b:3d:d7  xid: 0x0

DHCP options:
[53] Message type: Release
[54] DHCP server addr: 10.0.0.1
[255] End
"
```

Conclusion

The 7750 SR WLAN-GW can support many WiFi Offload architectures, including open SSID with portal authentication. WiFi Offload functions such as terminating GRE tunnels or subscribers can be performed in separate routing instances, if required. IP addresses can be assigned from an external or local source and routing can be performed using NAT, by connecting the UE routing instance directly to the Internet, or by leaking routes to other routing instances. Using http-redirect, a web portal can be used to allow users to log in to a paid service or to accept the terms of service for a free WiFi service. Several show commands and debug options are available to help the operator monitor and troubleshoot the solution.

WiFi Aggregation and Offload — Basic Secure SSID with Distributed RADIUS Proxy

This chapter provides information about WiFi Aggregation and Offload — Basic Secure SSID with Distributed RADIUS Proxy.

Topics in this chapter include:

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

Applicability

The information and configuration in this chapter are based on SR OS Release 13.0.R3.

WiFi Aggregation and Offload functionality has been supported in SR OS 10.0.R3, and later. This includes a RADIUS proxy server and support for soft-GRE tunnels.

Overview

WLAN-GW subscribers can be implemented using Enhanced Subscriber Management (ESM) on the CPM in order to benefit from the extensive ESM features available on the SR OS nodes. Many different WiFi Offload configurations are possible, with two versatile categories being open and secure SSID.

Starting from SR OS Release 12.0.R4, distributed RADIUS-proxy functionality (DRP) has been added to the MS-ISA. This feature allows running a high-performance proxy over multiple MS-ISA cards instead of being limited to a single CPM, greatly increasing scalability.

This chapter can be used as a starting point for operators who wish to configure a secure SSID scenario using DRP and ESM. In a secure SSID scenario, the Access Point (AP) uses 802.1x and Extensible Authentication Protocol (EAP) to authenticate the UE. The EAP method used is transparent to the WLAN-GW. In this chapter, PEAP/EAP-MSCHAPv2 is used to associate with the SSID by entering a user name and password, but other methods such as EAP-Subscriber Identity Module (EAP-SIM) can also be used without any configuration change on the WLAN-GW.

IP address assignment and Internet connectivity can be achieved by using various methods on the SR OS node. In this scenario, the RADIUS server provides the IP addresses to the User Equipment (UEs). The same IP private address is assigned to every UE and L2-aware Network Address Translation (NAT) is used to provide a public IP address on the Internet.

For a WiFi Offload solution where the service provider has a record of their users, that is, where users have login accounts, the provider may consider offering a secure SSID as a more convenient and secure alternative to an open SSID with a web portal. In a secure SSID scenario, all user traffic is encrypted between the UE and the AP, and UEs are only granted access if they authenticate successfully. This

makes attacks more difficult and blocks non-paying users who only connect to test if they can get free access.

Authentication in this case requires a centralized Authentication, Authorization and Accounting (AAA) which keeps track of the user accounts. The user is granted full access immediately after connecting to the secure SSID. One drawback is that WiFi clients may need some configuring by the user before they are able to connect to the SSID using the correct EAP method. In the case of EAP-SIM, users do not need to know their user name and password because authentication is done based on credentials contained in the SIM card, but the SSID configuration may need to be preloaded by the operator on the mobile device or provided to the user ahead of time. For other EAP methods such as PEAP/EAP-MSCHAPv2, users need to supply the correct user name and password, without the help of a portal or any instructions to guide them.

An operator offering Internet access to a large number of users while only a limited number of public IP addresses are available will likely use Network Address Translation (NAT) in order to conserve public IP addresses. NAT typically maps a few public IP addresses and ports to a large number of inside (private) IP addresses and ports. The WLAN-GW supports several NAT configurations including L2-aware NAT, where the MAC address of the UE is also used when creating the mapping between the inside IP/port and the outside IP/port. Therefore with L2-aware NAT, the same private IP address can be assigned to all UEs because the unique MAC address for each UE allows the WLAN-GW to distinguish between each UE. This greatly simplifies IP address assignment; the RADIUS server can assign the same private IP address to all UEs and there is no DHCP server required. Using a RADIUS server for IP address assignment means DHCP proxy needs to be configured on the WLAN-GW.


The SR OS platform is flexible in terms of allowing the operator to separate the various WiFi Offload functions between different routing instances. All functions can be configured in the same routing instance, or as shown in the following configuration, the connectivity to the APs (and soft-GRE tunnels) can be provided in one Virtual Private Routed Network (VPRN), the users can be instantiated in another VPRN, and AAA access can be provided in yet another routing instance (in this example the Base router). This provides a clear separation of functions and can enhance security, by separating user traffic from authentication and management traffic.

Configuration

The WiFi offload scenario with SSID and L2-aware NAT shown in [Figure 174: WiFi Offload Scenario with Secure SSID and L2-Aware NAT](#) has following characteristics:

- Secure SSID with EAP authentication
- Local breakout to Internet, routing through the Base routing instance
- Same private IP address assigned to all UEs by RADIUS
- L2-aware NAT
- AP access in VPRN 1000
- UEs terminated in VPRN 2001

Figure 174: WiFi Offload Scenario with Secure SSID and L2-Aware NAT



WLAN-GW

Note that configuring the uplink interface, Interior Gateway Protocol (IGP) and system configuration is outside the scope of this chapter and only partial configuration is provided.

Card and Media Dependent Adapter (MDA) configuration showing only the WLAN-IOM. An IOM card containing two MS-ISA cards provides the WLAN-GW functionality. The MDA type for the ISA cards is isa-bb, the same type that is used for NAT.

```
*A:WLAN-GW# /configure card 2
*A:WLAN-GW>config>card# info
```

```
-----
card-type iom3-xp-b
mda 1
  mda-type isa-bb
  no shutdown
exit
mda 2
  mda-type isa-bb
  no shutdown
exit
no shutdown
-----
```

The ISA configuration contains a wlan-gw-group referencing the IOM in slot 2, which hosts the two MS-ISA cards providing the WLAN-GW functions.

```
A:WLAN-GW# /configure isa
A:WLAN-GW>config>isa# info
```

```
-----
wlan-gw-group 1 create
  active-iom-limit 1
  iom 2
  no shutdown
exit
-----
```


The AAA configuration contains an ISA RADIUS policy used for authentication requests. The source address range configures the source address of the first MS-ISA in the wlan-gw-group. The second MS-ISA card gets the next consecutive IP address and so on. All the IP addresses assigned this way to MS-ISA cards must be configured as clients on the RADIUS server. The secret here must match the secret configured on the RADIUS server.

```
A:WLAN-GW# /configure aaa
A:WLAN-GW>config>aaa# info
-----
    isa-radius-policy "IRS_1" create
        servers
            router "Base"
            source-address-range 10.10.165.1
            server 1 create
                authentication
                ip-address 10.93.186.2
                secret "7USmr6f7JkxaGnDDq1uqwEAJKGbZr5i" hash2
                no shutdown
            exit
        exit
    exit
exit
-----
```

The following policy shows the two routes that must be exported for this scenario to work: UE NAT outside routes (NAT is configured in the next step), to make UEs reachable on the Internet, and MS-ISA RADIUS source address routes, in order for the MS-ISAs to be reachable from the RADIUS server. This policy should be used for export in the IGP configuration (not shown).

```
A:WLAN-GW# /configure router policy-options
A:WLAN-GW>config>router>policy-options# info
-----
    prefix-list "WiFi"
        prefix 10.10.165.0/24 longer
        prefix 10.165.0.0/16 longer
    exit
    policy-statement "WiFi"
        entry 10
            from
                prefix-list "WiFi"
            exit
            action accept
            exit
        exit
    exit
exit
-----
```

The following configures L2-aware NAT by creating an outside NAT pool with a public IP address range. The private inside address used by the UE will be mapped to an outside IP address routable on the Internet. NAT port mapping parameters can be set in this configuration, controlling how many outside ports can be used by each UE. Details of NAT configuration are outside the scope of this document.

```
A:WLAN-GW# /configure router nat
A:WLAN-GW>config>router>nat# info
-----
    outside
        pool "WiFi-165-1" nat-group 1 type l2-aware create
        address-range 10.165.1.0 10.165.1.255 create
        exit
        no shutdown
    exit
-----
```

```
exit
-----
```

The following configures a NAT policy under services, linking this policy with the outside NAT pool. When the NAT policy is invoked for a subscriber, this associates the subscriber with the correct outside pool.

```
A:WLAN-GW# /configure service nat
A:WLAN-GW>config>service>nat# info
-----
    nat-policy "WiFi-165-1" create
        pool "WiFi-165-1" router Base
    exit
-----
```

The following subscriber management configuration includes an SLA profile, a subscriber identity policy, and the subscriber profile that makes use of the previously defined NAT policy. This allows subscriber traffic to be forwarded to the Internet through the Base routing instance where the outside NAT pool exists.

A dummy authentication-policy is required for the CPM to handle the DHCP Discover messages forwarded by the MS-ISA cards.

```
A:WLAN-GW# /configure subscriber-mgmt
A:WLAN-GW>config>subscr-mgmt# info
-----
    authentication-policy "dummy-auth-policy" create
    exit
    sla-profile "SLAP_1" create
    exit
    sub-profile "SUBP_1" create
        nat-policy "WiFi-165-1"
    exit
    sub-ident-policy "SIP" create
        sub-profile-map
            use-direct-map-as-default
        exit
        sla-profile-map
            use-direct-map-as-default
        exit
    exit
exit
-----
```

VPRN 1000 contains the interface to the AP as well as the distributed RADIUS proxy server RP_1. The RADIUS proxy wlan-gw address configures a special NAT route in VPRN 1000 that forwards RADIUS packets from the AP to the correct MS-ISA. That address is known to the AP as the RADIUS server address it uses for EAP authentication for this SSID. The secret configured here has to match the RADIUS secret configured on the AP.

The RADIUS proxy is configured to create cache entries based on attribute 31 in RADIUS access-request packets (Calling-Station-ID), which contains the MAC address of the UE. These cache entries will be stored temporarily and used to authenticate DHCP packets from the UE. The **track-accounting start** parameter allows mobility to be triggered for a UE upon receiving an accounting-start message. The UE's associated tunnel will be moved to the IP address indicated by the NAS-IP-Address. The **track-accounting stop** parameter allows the UE session to be terminated immediately when the AP sends a RADIUS accounting-stop for the UE, when this UE disconnects from the SSID.

The default-authentication-server-policy links the RADIUS proxy with the isa-radius-policy that authenticates the UEs. If accounting is required, the accounting policy can be specified in this configuration and can be the same as or different from the isa-radius-policy. The send-accounting-response option

makes the WLAN-GW acknowledge (and then discard) the RADIUS accounting messages from the AP, instead of proxying the accounting messages to the external RADIUS server.

```
A:WLAN-GW# /configure service vprn 1000
A:WLAN-GW>config>service>vprn# info
-----
    route-distinguisher 65400:1000
    interface "toAP3" create
        address 10.1.3.1/24
        sap 1/1/10 create
        exit
    exit
    radius-proxy
        server "DRP_1" purpose accounting authentication wlan-gw-group 1 create
            cache
                key packet-type request attribute-type 31
                track-accounting start stop
                no shutdown
            exit
            default-authentication-server-policy "IRS_1"
            secret "nUeorYjgFZtuAqIwoU0LODFxF43rhSf/" hash2
            send-accounting-response
            wlan-gw
                address 192.11.11.11
            exit
            no shutdown
        exit
    exit
    no shutdown
-----
```

VRPN 2001 is used for UE termination and contains:

- A subscriber interface and group interface of type **wlangw** (**soft-gre** prior to Release 12.0).
- Default subscriber parameters assigned to every UE.
- DHCP proxy, which allows the RADIUS-assigned IP address parameters stored in the DRP cache during authentication to be passed to the UE.
- A dummy authentication policy which allows the CPM to handle the DHCP Discover passed on by the MS-ISA.
- The **wlan-gw** node (**soft-gre** prior to Release 12.0), which includes:
 - The gw-address that is the end-point of the GRE tunnel
 - The routing instance where AP traffic is terminated
 - The ISA wlan-gw-group, which associates this WLAN-GW configuration with a set of IOMs
 - Mobility parameters, which allow the UE state to be kept if the UE moves between two APs broadcasting the same SSID
 - The authenticate-on-dhcp option required for the CPM to instantiate ESM UEs when using DRP
 - The L2-aware address/subnet used for L2-aware NAT. This address matches the default gateway assigned to the UEs.

```
A:WLAN-GW# /configure service vprn 2001
A:WLAN-GW>config>service>vprn# info
-----
    description "WiFi-165-1 Secure SSID"
```

```

route-distinguisher 65400:2001
subscriber-interface "SII" create
  address 192.168.1.1/24 populate-host-routes
  group-interface "GI1" wlangw create
    sap-parameters
      sub-sla-mgmt
        def-sla-profile "SLAP_1"
        def-sub-profile "SUBP_1"
        sub-ident-policy "SIP"
      exit
    exit
  dhcp
    proxy-server
      emulated-server 192.168.1.1
      no shutdown
    exit
    lease-populate 10000
    gi-address 192.168.1.1
    no shutdown
  exit
  authentication-policy "dummy-auth-policy"
  wlan-gw
    gw-address 192.1.1.1
    mobility
      trigger data iapp
    exit
    router 1000
    wlan-gw-group 1
    vlan-tag-ranges
      range default
      authenticate-on-dhcp
    exit
    exit
    no shutdown
  exit
exit
nat
  inside
    l2-aware
    address 192.168.1.1/24
  exit
exit
exit
no shutdown
-----

```

Freeradius

This part of the user's configuration file matches on the user name entered by the UE while connecting to this secure SSID. If the password entered is correct, RADIUS returns the IP addressing parameters configured as follows. The same IP address 192.168.1.9 is assigned to every user on this SSID, but the L2-aware NAT on the WLAN-GW can distinguish between all the UEs based on their L2 MAC address.

```

/etc/freeradius/users:
"user1"      User-Password := "alcatel"
             Alc-Subsc-ID-Str = "%{User-Name}",
             Alc-Default-Router = 192.168.1.1,
             Alc-Primary-Dns = 10.43.186.43,
             Framed-IP-Address = 192.168.1.9,

```

```
Framed-IP-Netmask = 255.255.255.0,
```

In `/etc/freeradius/clients.conf` the secret matches the one configured in the WLAN-GW `isa-radius-policy` configuration. Since there are only two MS-ISA cards in the `wlan-group` used in this example, two clients are configured.

```
client 10.10.165.1 {
    secret      = alcatel
    shortname   = WLAN-GW-ISA1
}
client 10.10.165.2 {
    secret      = alcatel
    shortname   = WLAN-GW-ISA2
}
```

Access Points

The following must be configured on the Access Point as a minimum:

- IP address 10.1.3.10/24
- Default route to 10.1.3.1
- Secure SSID WiFi-165-1 mapped to VLAN 10, using WPA2 with EAP/802.1x authenticating against RADIUS server 192.11.11.11, with RADIUS accounting enabled
- Soft-GRE tunnel with destination 192.1.1.1, with VLAN 10 mapped to this tunnel

Show Commands

The following show commands reflect the status of the WLAN-GW after the UE has connected and obtained an IP address using DHCP.

The following output displays the connected UEs:

```
A:WLAN-GW# show subscriber-mgmt wlan-gw ue
=====
User Equipments
=====
MAC address           : 68:7f:74:8b:3d:d7
-----
VLAN Q-tag            : 10
MPLS label            : (Not Specified)
Tunnel router         : 1000
Tunnel remote IP address : 10.1.3.10
Tunnel local IP address  : 192.1.1.1
Retail service        : N/A
SSID                  : "WiFi-165-1"
Previous Access Point IP : (Not Specified)
IMSI                  : (Not Specified)
Subscriber host service : 2001
Subscriber host SAP     : 2/1/nat-out-ip:2049.1
Last move time        : 2015/09/15 16:20:01
-----
No. of UE: 1
=====
```

```
A:WLAN-GW# tools dump wlan-gw ue

=====
Matched 1 session on Slot #2 MDA #1
=====
UE-Mac       : 68:7f:74:8b:3d:d7    UE-vlan      : 10
UE IP Addr   : N/A                 UE Timeout   : N/A
DHCPv6 Timeout : N/A               SLAAC Timeout : N/A
DHCPv6 IA-NA ID : N/A              RA Timeout   : N/A
DHCPv6 Addr   : N/A
SLAAC Prefix  : N/A
Description   : ESM-user
Auth/CoA-time : 09/16/2015 10:47:38 Retail Service : N/A
Tunnel MDA    : 2/2                 Tunnel Router : 1000
MPLS label    : N/A                 Shaper        : 1
Tunnel Src IP : 10.1.3.10            Tunnel Dst IP  : 192.1.1.1
Tunnel L2 Svc : N/A                 Tunnel L2 Vlan : N/A
Tunnel Type   : GRE
Anchor SAP    : 2/1/nat-out-ip:2049.2
AP-Mac        : 00:0d:67:39:0b:65    AP-RSSI       : Unknown
AP-SSID       : "WiFi-165-1"
Last-forward  : 09/16/2015 15:59:26 Last-move     : 09/16/2015 10:47:38
Session Timeout : None               Idle Timeout   : N/A
Acct Update    : None               Acct Interval  : N/A
Acct Session-Id : N/A
Acct Policy    : N/A
NAT Policy     : N/A
Redirect Policy : N/A
IP Filter      : N/A
App-profile    : N/A
Rx Oper PIR    : N/A                 Rx Oper CIR    : N/A
Tx Oper PIR    : N/A                 Tx Oper CIR    : N/A
Rx Frames      : N/A                 Rx Octets      : N/A
Tx Frames      : N/A                 Tx Octets      : N/A
-----
No sessions on Slot #2 MDA #2 match the query
=====
```

The DHCP lease information indicates that the address was assigned by RADIUS.

```
A:WLAN-GW# show service id 2005 dhcp lease-state

=====
DHCP lease state table, service 2001
=====
IP Address      Mac Address      Sap/Sdp Id      Remaining Lease   MC
                  LeaseTime   Origin   Stdby
-----
192.168.1.9     68:7f:74:8b:3d:d7 [2/1/nat-out-ip:20* 06d23h59m Radius
-----
Number of lease states : 1
=====
* indicates that the corresponding row element may have been truncated.
```

When troubleshooting DHCP issues, displaying DHCP statistics is useful.

```
A:WLAN-GW# show service id 2005 dhcp statistics

=====
DHCP Global Statistics, service 2001
=====
Rx Packets      : 2
Tx Packets      : 2
Rx Malformed Packets : 0
Rx Untrusted Packets : 0
```

```
Client Packets Discarded      : 0
Client Packets Relayed       : 0
Client Packets Snooped       : 0
Client Packets Proxied (RADIUS) : 2
Client Packets Proxied (Diameter) : 0
Client Packets Proxied (User-Db) : 0
Client Packets Proxied (Lease-Split) : 0
Server Packets Discarded     : 0
Server Packets Relayed      : 0
Server Packets Snooped      : 0
DHCP RELEASEs Spoofed       : 0
DHCP FORCERENEWs Spoofed    : 0
=====
```

The following output lists the active subscribers, showing each UE SLA profile, MAC address and IP address.

```
A:WLAN-GW# show service active-subscribers
=====
Active Subscribers
=====
-----
Subscriber DUACBU2ZLE (SUBP_1)
-----
NAT Policy: WiFi-165-1
Outside IP: 10.165.1.0
Ports      : 1024-65535
-----
(1) SLA Profile Instance sap:[2/1/nat-out-ip:2049.1] - sla:SLAP_1
-----
IP Address      MAC Address      PPPoE-SID Origin
-----
192.168.1.9     68:7f:74:8b:3d:d7 N/A      DHCP
-----
Number of active subscribers : 1
-----
```

The following output displays distributed RADIUS proxy server statistics after the UE has authenticated, showing all the EAP messages exchanged between the AP and RADIUS proxy:

```
A:WLAN-GW# show router 1000 radius-proxy-server "DRP_1" statistics
=====
ISA RADIUS Proxy server statistics for "DRP_1"
=====
Group 1 member 1
-----
Rx packet                : 12
Rx Access-Request        : 11
Rx Accounting-Request    : 1
Rx dropped                : 0
  Retransmit             : 0
  Wrong purpose          : 0
  No UE MAC to cache     : 0
  Client context limit reached : 0
  No ISA RADIUS policy configured : 0
  Invalid attribute encoding : 0
  Invalid password       : 0
  Accounting-Request with invalid Acct-Status-Type : 0
```

```

Accounting-Request with no Acct-Status-Type      : 0
Invalid accounting Authenticator                 : 0
Invalid Message-Authenticator                   : 0
Management core overload                         : 0

Tx Access-Accept                                : 1
Tx Access-Reject                                : 0
Tx Access-Challenge                             : 10
Tx Accounting-Response                           : 1
Tx dropped                                       : 0
  Server timeout                               : 0
  Invalid response Authenticator                : 0
  Invalid Message-Authenticator                : 0
  Invalid attribute encoding                   : 0
  RADIUS server send failure                   : 0

Group 1 member 2
-----
Rx packet                                         : 0
Rx Access-Request                               : 0
Rx Accounting-Request                           : 0
Rx dropped                                       : 0
  Retransmit                                    : 0
  Wrong purpose                                : 0
  No UE MAC to cache                           : 0
  Client context limit reached                  : 0
  No ISA RADIUS policy configured               : 0
  Invalid attribute encoding                   : 0
  Invalid password                             : 0
  Accounting-Request with invalid Acct-Status-Type : 0
  Accounting-Request with no Acct-Status-Type   : 0
  Invalid accounting Authenticator              : 0
  Invalid Message-Authenticator                : 0
  Management core overload                     : 0

Tx Access-Accept                                : 0
Tx Access-Reject                                : 0
Tx Access-Challenge                             : 0
Tx Accounting-Response                           : 0
Tx dropped                                       : 0
  Server timeout                               : 0
  Invalid response Authenticator                : 0
  Invalid Message-Authenticator                : 0
  Invalid attribute encoding                   : 0
  RADIUS server send failure                   : 0
=====

```

The following output shows the ISA RADIUS policy statistics after the UE has connected, showing the transactions between the WLAN-GW and the RADIUS server.

```

A:WLAN-GW# show aaa isa-radius-policy "IRS_1"
=====
ISA RADIUS policy "IRS_1"
=====
Description          : (Not Specified)
Include attributes acct : N/A
Include attributes auth : nas-ip-address
User name format      : mac
User name MAC format   : alu
NAS-IP-Address         : system-ip
-----
RADIUS server settings
-----

```



```

Router          : "Base"
Source address start : 10.10.165.1
Source address end   : 10.10.165.2
Access algorithm    : direct
Retry               : 3
Timeout (s)         : 5
Last management change : 09/15/2015 15:05:02
=====
Servers for "IRS_1"
=====
Index Address      Acct-port Auth-port CoA-port
-----
1    10.93.186.2    0        1812    0
=====
Status for ISA RADIUS server policy "IRS_1"
=====
Server 1, group 1, member 1
-----
Purposes Up           : authentication
Source IP address      : 10.10.165.1
Acct Tx Requests       : 0
Acct Tx Retries        : 0
Acct Tx Timeouts       : 0
Acct Rx Replies        : 0
Auth Tx Requests       : 11
Auth Tx Retries        : 0
Auth Tx Timeouts       : 0
Auth Rx Replies        : 11
CoA Rx Requests        : 0

Server 1, group 1, member 2
-----
Purposes Up           : (None)
Source IP address      : 10.10.165.2
Acct Tx Requests       : 0
Acct Tx Retries        : 0
Acct Tx Timeouts       : 0
Acct Rx Replies        : 0
Auth Tx Requests       : 0
Auth Tx Retries        : 0
Auth Tx Timeouts       : 0
Auth Rx Replies        : 0
CoA Rx Requests        : 0
=====

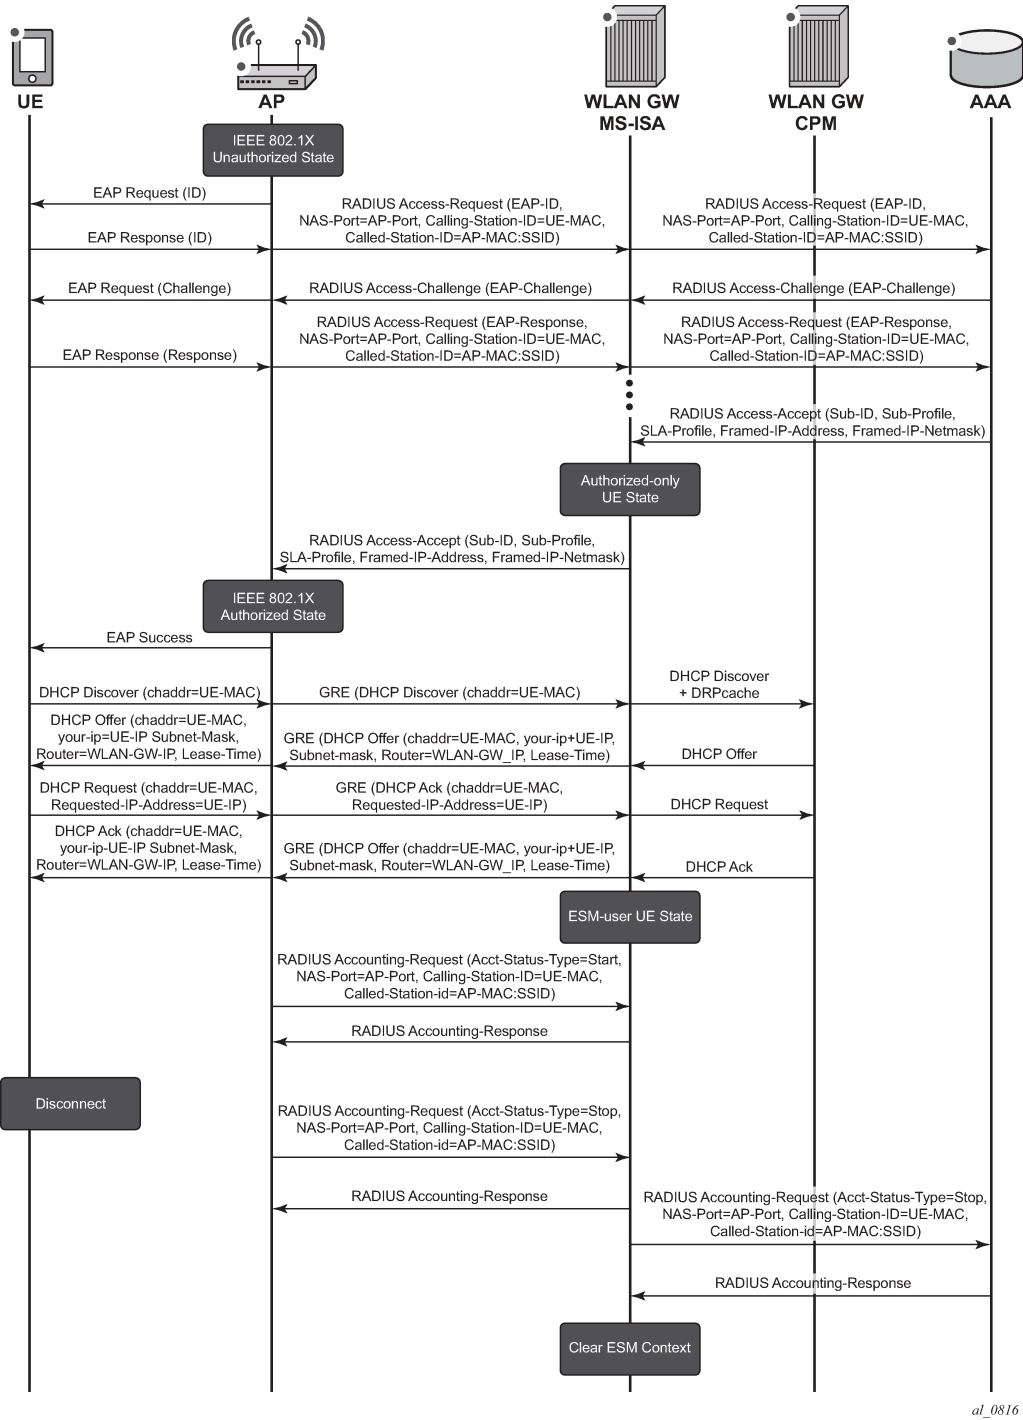
```

Call Flow

Figure 175: Call Flow for Secure SSID with DSM shows the call flow for a secure SSID with DRP. The main steps are:

- EAP authentication using DRP on MS-ISA placing UE authorized-only state
- UE sends DHCP Discover which is forwarded by the MS-ISA to the CPM
- CPM places UE in ESM-user state
- Upon disconnect, the AP sends a RADIUS accounting-stop which clears the UE context on the WLAN-GW

Figure 175: Call Flow for Secure SSID with DSM



al_0816

Debug

In this example, the following debug configuration is used (note that some default options are automatically added and do not need to be entered manually, e.g. mode under dhcp). For DRP, only a limited number of UEs can be debugged at a time and their MAC address have to be specified.

```
debug
  router "2001"
    ip
      dhcp
        detail-level medium
        mode egr-ingr-and-dropped
      exit
    exit
  exit
wlan-gw
  group 1
    ue 68:7f:74:8b:3d:d7 packet radius dhcp
  exit
exit
```

The following is a partial debug of a UE connecting to the SSID and authenticating with the RADIUS server. Shortly after logging in the UE disconnects from the SSID and the subscriber is removed on reception of the RADIUS accounting-stop message.

As soon as the UE attempts to connect to the secure SSID, the WLAN-GW distributed RADIUS proxy in VPRN 1000 receives the first Access-Request packet from the AP. Note that the CALLING STATION ID [31] attribute contains the MAC address of the UE, and that the AP sends the SSID name in the NAS IDENTIFIER [32] attribute.

```
1464 2015/09/15 16:19:52.93 EDT MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 2/1, SeqNo 3291
  Info:      anchor ingressing frame
          radius upstream from client

IP/UDP:   from 10.1.3.10 (port 51235) to 192.11.11.11 (port 1812)

RADIUS:   Access-Request (1) id 122 len 190
  USER NAME [1] 5 user1
  NAS IP ADDRESS [4] 4 10.1.3.10
  FRAMED IP ADDRESS [8] 4 255.255.255.255
  NAS IDENTIFIER [32] 10 WiFi-165-1
  CALLED STATION ID [30] 28 00-0D-67-39-0B-65:WiFi-165-1
  NAS PORT TYPE [61] 4 Wireless - IEEE 802.11(19)
  NAS PORT [5] 4 0
  CALLING STATION ID [31] 17 68-7F-74-8B-3D-D7
  CONNECT INFO [77] 21 CONNECT 0Mbps 802.11b
  SESSION ID [44] 17 556F2789-0000008D
  FRAMED MTU [12] 4 1400
  EAP MESSAGE [79] 10 0x02e2000a017573657231
  MESSAGE AUTHENTICATOR [80] 16 0xbc3a66d7f9d4e02465797f2018914ed7
"
```

The WLAN-GW MS-ISA forwards the Access-Request to the RADIUS server in the Base router.

```
1465 2015/09/15 16:19:52.94 EDT MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 2/1, SeqNo 3292
  Info:      anchor egressing frame
```

```

radius upstream to server

IP/UDP:   from 10.10.165.1 (port 1024) to 10.93.186.2 (port 1812)

RADIUS:   Access-Request (1) id 20 len 190
USER NAME [1] 5 user1
NAS IP ADDRESS [4] 4 10.1.3.10
FRAMED IP ADDRESS [8] 4 255.255.255.255
NAS IDENTIFIER [32] 10 WiFi-165-1
CALLED STATION ID [30] 28 00-0D-67-39-0B-65:WiFi-165-1
NAS PORT TYPE [61] 4 Wireless - IEEE 802.11(19)
NAS PORT [5] 4 0
CALLING STATION ID [31] 17 68-7F-74-8B-3D-D7
CONNECT INFO [77] 21 CONNECT 0Mbps 802.11b
SESSION ID [44] 17 556F2789-0000008D
FRAMED MTU [12] 4 1400
EAP MESSAGE [79] 10 0x02e2000a017573657231
MESSAGE AUTHENTICATOR [80] 16 0xbf48919833584995109b8387efc03b21
"

```

Many RADIUS Access-Request and Access Challenge messages are exchanged, which encapsulate the EAP authentication between the UE and the RADIUS server. At the end of the exchange, for a successful authentication, the WLAN-GW receives an Access-Accept message (for a failed authentication it would receive an Access-Reject).

```

1506 2015/09/15 16:20:01.69 EDT MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 2/1, SeqNo 3333
Info:      anchor ingressing frame
           radius downstream from server

IP/UDP:    from 10.93.186.2 (port 1812) to 10.10.165.1 (port 1024)

RADIUS:    Access-Accept (2) id 25 len 203
VSA [26] 6 Alcatel(6527)
  DEFAULT ROUTER [18] 4 192.168.1.1
VSA [26] 6 Alcatel(6527)
  PRIMARY DNS [9] 4 10.92.186.92
FRAMED IP ADDRESS [8] 4 192.168.1.9
FRAMED IP NETMASK [9] 4 255.255.255.0
USER NAME [1] 5 user1
VSA [26] 52 Microsoft(311)
  MS MPPE RECV KEY [17] 50 0xc1af6befb148f03d5bd9bb8863500dd0a1ffcf57392dcda
8db5529be6e2de52fc239d3595212ee1b181e50c064e292595db8
VSA [26] 52 Microsoft(311)
  MS MPPE SEND KEY [16] 50 0xcbe708a0751bc3c9ef43bb58e2b103cca0a6373b6800279
148a0f1934176f000e1540e5078eeba9d43af5f42d4799b16a79d
EAP MESSAGE [79] 4 0x03ec0004
MESSAGE AUTHENTICATOR [80] 16 0xb3d2459f830217fd455b26e7767012c3
"

```

The Access-Accept contains the IP addressing parameters for the UE such as the IP address, netmask, and default gateway, as well as the subscriber ID string. The IP addressing information is used by the WLAN-GW, but the Access-Accept message is also forwarded by the RADIUS proxy to the AP to tell it that the UE authenticated successfully so it can associate with the SSID.

```

1507 2015/09/15 16:20:01.69 EDT MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 2/1, SeqNo 3334
Info:      anchor egressing frame
           radius downstream to client

IP/UDP:    from 192.11.11.11 (port 1812) to 10.1.3.10 (port 51235)

```

```
RADIUS: Access-Accept (2) id 132 len 203
VSA [26] 6 Alcatel(6527)
  DEFAULT ROUTER [18] 4 192.168.1.1
VSA [26] 6 Alcatel(6527)
  PRIMARY DNS [9] 4 10.92.186.92
FRAMED IP ADDRESS [8] 4 192.168.1.9
FRAMED IP NETMASK [9] 4 255.255.255.0
USER NAME [1] 5 user1
VSA [26] 52 Microsoft(311)
  MS MPPE RECV KEY [17] 50 0xc1afa8a2e9f23dbe5c0d41410a8bcc7fc42406813a3bff6
a61c957fbad58b7af6de0447898603980aeebe5cc2d5db54b8ca7
VSA [26] 52 Microsoft(311)
  MS MPPE SEND KEY [16] 50 0xcbe7fb9182312534ea50ecdfc8ed59874401515968ae276
7826fa664e3871d0b13e2946b01750825dbb95b3fe6ee615afa1a
EAP MESSAGE [79] 4 0x03ec0004
MESSAGE AUTHENTICATOR [80] 16 0x66b269e340328cee108dfc1d27f46fca
"
```

After the AP allows the UE to connect to the secure SSID, establishing L2 connectivity to the WLAN-GW across the soft-GRE tunnel, the UE can obtain an IP address through DHCP. The WLAN-GW receives a DHCP Discover from the UE on MS-ISA MDA 2/1:

```
1508 2015/09/15 16:20:01.83 EDT MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 2/1, SeqNo 3335
  Info:      anchor ingressing frame
          received upstream from tunnel

Ethernet: from 68:7f:74:8b:3d:d7 to ff:ff:ff:ff:ff:ff (ethertype: 0x0800)

IP/UDP:   from 0.0.0.0 (port 68) to 255.255.255.255 (port 67)

DHCP:
ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
siaddr: 0.0.0.0          giaddr: 0.0.0.0
chaddr: 68:7f:74:8b:3d:d7  xid: 0xfb4fb37

DHCP options:
[53] Message type: Discover
[61] Client id: (hex) 01 68 7f 74 8b 3d d7
[12] Host name: VMS11
[60] Class id: MSFT 5.0
[55] Param request list: len = 12
      1 Subnet mask
     15 Domain name
      3 Router
      6 Domain name server
     44 NETBIOS name server
     46 NETBIOS type
     47 NETBIOS scope
     31 Router discovery
     33 Static route
    121 Unknown option
    249 Unknown option
     43 Vendor specific
[255] End
"
```

The MS-ISA forwards the DHCP Discover to the CPM and it arrives on group interface GI1 in VPRN 2001.

```
1509 2015/09/15 16:20:01.83 EDT MINOR: DEBUG #2001 vprn2001 PIP
"PIP: DHCP"
```

```
instance 6 (2001), interface index 3 (GI1),
  received DHCP Boot Request on Interface GI1 (2/1/nat-out-ip:2049.1) Port 67

H/W Type: Ethernet(10Mb)  H/W Address Length: 6
ciaddr: 0.0.0.0            yiaddr: 0.0.0.0
siaddr: 0.0.0.0            giaddr: 0.0.0.0
chaddr: 68:7f:74:8b:3d:d7  xid: 0xfb4fb37

DHCP options:
[53] Message type: Discover
[61] Client id: (hex) 01 68 7f 74 8b 3d d7
[12] Host name: VMS11
[60] Class id: MSFT 5.0
[55] Param request list: len = 12
      1 Subnet mask
     15 Domain name
      3 Router
      6 Domain name server
     44 NETBIOS name server
     46 NETBIOS type
     47 NETBIOS scope
     31 Router discovery
     33 Static route
    121 Unknown option
    249 Unknown option
     43 Vendor specific
[255] End
"
```

The WLAN-GW sends a DHCP Offer to the UE with the IP address information retrieved from the RADIUS Access-Accept message.

```
1510 2015/09/15 16:20:01.85 EDT MINOR: DEBUG #2001 vprn2001 PIP
"PIP: DHCP
instance 6 (2001), interface index 3 (GI1),
  transmitted DHCP Boot Reply to Interface GI1 (2/1/nat-out-ip:2049.1) Port 68

H/W Type: Ethernet(10Mb)  H/W Address Length: 6
ciaddr: 0.0.0.0            yiaddr: 192.168.1.9
siaddr: 192.168.1.1        giaddr: 192.168.1.1
chaddr: 68:7f:74:8b:3d:d7  xid: 0xfb4fb37

DHCP options:
[53] Message type: Offer
[54] DHCP server addr: 192.168.1.1
[51] Lease time: 604800
[1] Subnet mask: 255.255.255.0
[3] Router: 192.168.1.1
[6] Domain name server: 10.92.186.92
[28] Broadcast addr: 192.168.1.255
[61] Client id: (hex) 01 68 7f 74 8b 3d d7
[12] Host name: VMS11
[255] End
"
```

The Offer message is sent to the MS-ISA and towards the UE (not shown). The UE then sends a DHCP Request and the WLAN-GW responds with an Ack.

```
1513 2015/09/15 16:20:01.86 EDT MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 2/1, SeqNo 3337
  Info:      anchor ingressing frame
         received upstream from tunnel
```

```
Ethernet: from 68:7f:74:8b:3d:d7 to ff:ff:ff:ff:ff:ff (ethertype: 0x0800)

IP/UDP:   from 0.0.0.0 (port 68) to 255.255.255.255 (port 67)

DHCP:
ciaddr: 0.0.0.0          yiaddr: 0.0.0.0
siaddr: 0.0.0.0          giaddr: 0.0.0.0
chaddr: 68:7f:74:8b:3d:d7  xid: 0xfb4fb37

DHCP options:
[53] Message type: Request
[61] Client id: (hex) 01 68 7f 74 8b 3d d7
[50] Requested IP addr: 192.168.1.9
[54] DHCP server addr: 192.168.1.1
[12] Host name: VMS11
[81] client FQDN: rcode1: 0, rcode2: 0, domain name = (hex) 00 56 4d 53 31
31
[60] Class id: MSFT 5.0
[55] Param request list: len = 12
      1 Subnet mask
     15 Domain name
      3 Router
      6 Domain name server
     44 NETBIOS name server
     46 NETBIOS type
     47 NETBIOS scope
     31 Router discovery
     33 Static route
    121 Unknown option
    249 Unknown option
     43 Vendor specific
[255] End
"

1515 2015/09/15 16:20:01.86 EDT MINOR: DEBUG #2001 vprn2001 PIP
"PIP: DHCP
instance 6 (2001), interface index 3 (GI1),
  transmitted DHCP Boot Reply to Interface GI1 (2/1/nat-out-ip:2049.1) Port 68

H/W Type: Ethernet(10Mb)  H/W Address Length: 6
ciaddr: 0.0.0.0          yiaddr: 192.168.1.9
siaddr: 192.168.1.1      giaddr: 192.168.1.1
chaddr: 68:7f:74:8b:3d:d7  xid: 0xfb4fb37

DHCP options:
[53] Message type: Ack
[54] DHCP server addr: 192.168.1.1
[51] Lease time: 604800
[1] Subnet mask: 255.255.255.0
[3] Router: 192.168.1.1
[6] Domain name server: 10.92.186.92
[28] Broadcast addr: 192.168.1.255
[61] Client id: (hex) 01 68 7f 74 8b 3d d7
[12] Host name: VMS11
[81] client FQDN: rcode1: 0, rcode2: 0, domain name = (hex) 00 56 4d 53 31
31
[255] End
"
```

The AP sends a RADIUS accounting Start to the WLAN-GW as a result of the UE successfully associating with the SSID.

```
1518 2015/09/15 16:20:01.88 EDT MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 2/1, SeqNo 3339
  Info:      anchor ingressing frame
           radius upstream from client

IP/UDP:   from 10.1.3.10 (port 51236) to 192.11.11.11 (port 1813)

RADIUS:   Accounting-Request (4) id 133 len 197
SESSION ID [44] 17 556F2789-0000008D
EVENT TIMESTAMP [55] 4 1442292269
STATUS TYPE [40] 4 Start(1)
AUTHENTIC [45] 4 RADIUS(1)
USER NAME [1] 5 user1
NAS IP ADDRESS [4] 4 10.1.3.10
FRAMED IP ADDRESS [8] 4 192.168.1.9
NAS IDENTIFIER [32] 10 WiFi-165-1
CALLED STATION ID [30] 28 00-0D-67-39-0B-65:WiFi-165-1
NAS PORT TYPE [61] 4 Wireless - IEEE 802.11(19)
NAS PORT [5] 4 0
CALLING STATION ID [31] 17 68-7F-74-8B-3D-D7
CONNECT INFO [77] 21 CONNECT 0Mbps 802.11b
SESSION ID [44] 17 556F2789-0000008D
DELAY TIME [41] 4 0
"
```

At the end of the session, the UE disconnects from the SSID, and the AP sends a RADIUS accounting Stop to the WLAN-GW.

```
1520 2015/09/15 16:20:37.45 EDT MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 2/1, SeqNo 3401
  Info:      anchor ingressing frame
           radius upstream from client

IP/UDP:   from 10.1.3.10 (port 51237) to 192.11.11.11 (port 1813)

RADIUS:   Accounting-Request (4) id 134 len 233
SESSION ID [44] 17 556F2789-0000008D
EVENT TIMESTAMP [55] 4 1442292304
STATUS TYPE [40] 4 Stop(2)
AUTHENTIC [45] 4 RADIUS(1)
USER NAME [1] 5 user1
NAS IP ADDRESS [4] 4 10.1.3.10
FRAMED IP ADDRESS [8] 4 192.168.1.9
NAS IDENTIFIER [32] 10 WiFi-165-1
CALLED STATION ID [30] 28 00-0D-67-39-0B-65:WiFi-165-1
NAS PORT TYPE [61] 4 Wireless - IEEE 802.11(19)
NAS PORT [5] 4 0
CALLING STATION ID [31] 17 68-7F-74-8B-3D-D7
CONNECT INFO [77] 21 CONNECT 0Mbps 802.11b
SESSION ID [44] 17 556F2789-0000008D
DELAY TIME [41] 4 0
SESSION TIME [46] 4 35
INPUT PACKETS [47] 4 57
OUTPUT PACKETS [48] 4 36
INPUT OCTETS [42] 4 5228
OUTPUT OCTETS [43] 4 5538
TERMINATE CAUSE [49] 4 User Request(1)
"
```


This removes the subscriber from the WLAN-GW, clears the DHCP state, and also removes the GRE tunnel if this UE is the last one on the tunnel.

Conclusion

The WLAN-GW can support many WiFi Offload architectures including secure SSID with various types of EAP authentication. WiFi Offload functions such as terminating GRE tunnels, NAT, and RADIUS server connectivity can be performed in separate routing instances if required. UE IP addresses can be assigned locally or from an external source such as RADIUS, and routing to the Internet can be performed in various ways, including NAT. Several show commands and debug options are available to help the operator monitor and troubleshoot the solution.

WiFi Aggregation and Offload — IPv4/v6 Dual-Stack UEs

This chapter provides information about WiFi aggregation and offload IPv4/v6 dual-stack UEs.

Topics in this chapter include:

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

Applicability

The information and configuration in this chapter are based on SR OS Release 13.0.R3.

WiFi Aggregation and Offload functionality is supported on SR OS 10.0.R3 or later. The functionality includes enhanced subscriber management (ESM) for user equipment (UE) that gains network access via a WiFi service.

Overview

This chapter provides a functional description of the WLAN-GW features related to IPv4/v6 dual-stack UEs, as well as the related configuration.

Because IP address demand is mainly due to mobile devices, the support of IPv6 on mobile devices is a major requirement to manage IPv4 address depletion.

However, IPv6 on mobile devices is currently considered as an add-on rather than a replacement of IPv4, so the demand is for IPv4/v6 dual-stack UEs.

The basic concepts of ESMv6 for IPoE dual-stack hosts also apply to dual-stack UEs. However, WLAN-GW operates in a bridged environment where the Access Point (AP) performs L2 forwarding of Ethernet frames between the IEEE 802.11 air interface and the soft-GRE, soft-L2TPv3, or VLAN tunnel. Therefore, a WLAN-GW treats each UE as an individual subscriber who connects to the WiFi service. This contrasts with ESMv6 IPoE hosts behind a routed residential gateway (RG), where multiple hosts connect via the RG and the BNG treats the RG as the subscriber.

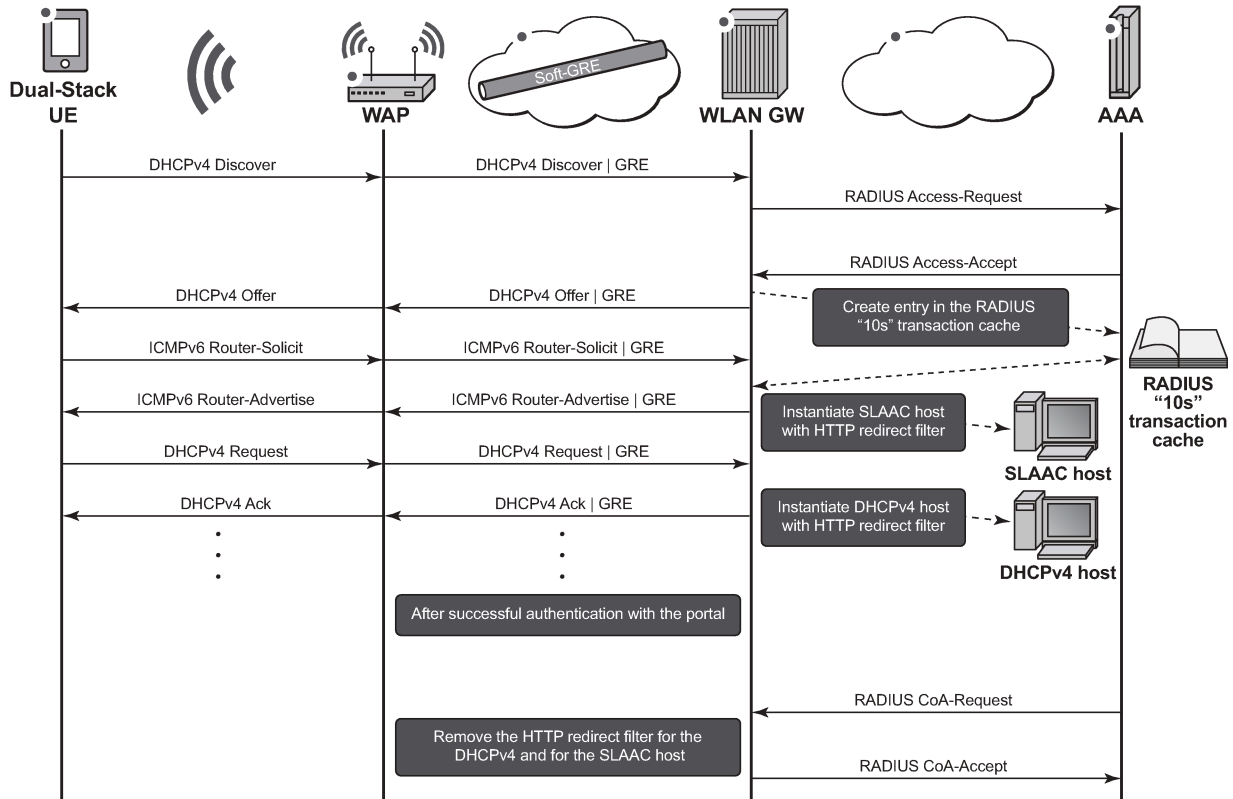
Depending on the type of UE, it may be allocated an IPv4 address through DHCPv4 and an IPv6 address through Stateless Address Auto-Configuration (SLAAC), or DHCPv6, or both (not all UEs have support for DHCPv6). Therefore, a UE can instantiate up to three IPoE hosts: a DHCPv4 host, a SLAAC host, and a DHCPv6 host.

Authentication and authorization depend on whether the UE connects to a WiFi with open or closed SSID. With an open SSID, authentication and authorization are performed when the first packet is received from the UE (typically a DHCPv4 Discover, an ICMPv6 Router-Solicit, or a DHCPv6 Solicit), similar to the routed RG model. Upon successful authentication, the Access-Accept is stored for 10 s on the WLAN-GW, so

for a dual-stack IPv4/v6 UE, two or three authentication rounds can be avoided if DHCPv4, SLAAC, and DHCPv6 are started within this 10 s interval.

When the UE has successfully authenticated with the portal, a CoA-Request may lift the HTTP redirect filter by changing the SLA profile and, optionally, the subscriber profile. If the CoA-Request contains the subscriber ID, the CoA-Request applies to both the DHCPv4 host and the SLAAC and/or DHCPv6 host. See the RADIUS attributes reference guide for more information about alternative subscriber host identification in RADIUS CoA-Request messages.

Figure 176: DHCPv4 + SLAAC/64 — Open SSID



al_0818

With a closed SSID, there is a separation between the authentication and authorization phases. When a UE connects to a WiFi with closed SSID in WPA-Enterprise mode, also known as WPA-802.1X mode, the UE initiates authentication before it obtains an IP address. The WLAN-GW is aware of the successful authentication when it receives the DHCPv4 Discover, the ICMPv6 Router-Solicit, or the DHCPv6 Solicit.

As with ESMv6, the WLAN-GW supports SLAAC/64 and DHCPv6/128 Identity Association for Non-temporary Addresses (IA_NA). DHCPv6 Identity Association for Prefix Delegation (IA_PD) is not supported because the UEs are considered as individual hosts that have direct Layer 2 connectivity with the WLAN-GW. Devices that use the UE as an IPv6 gateway are currently not supported.

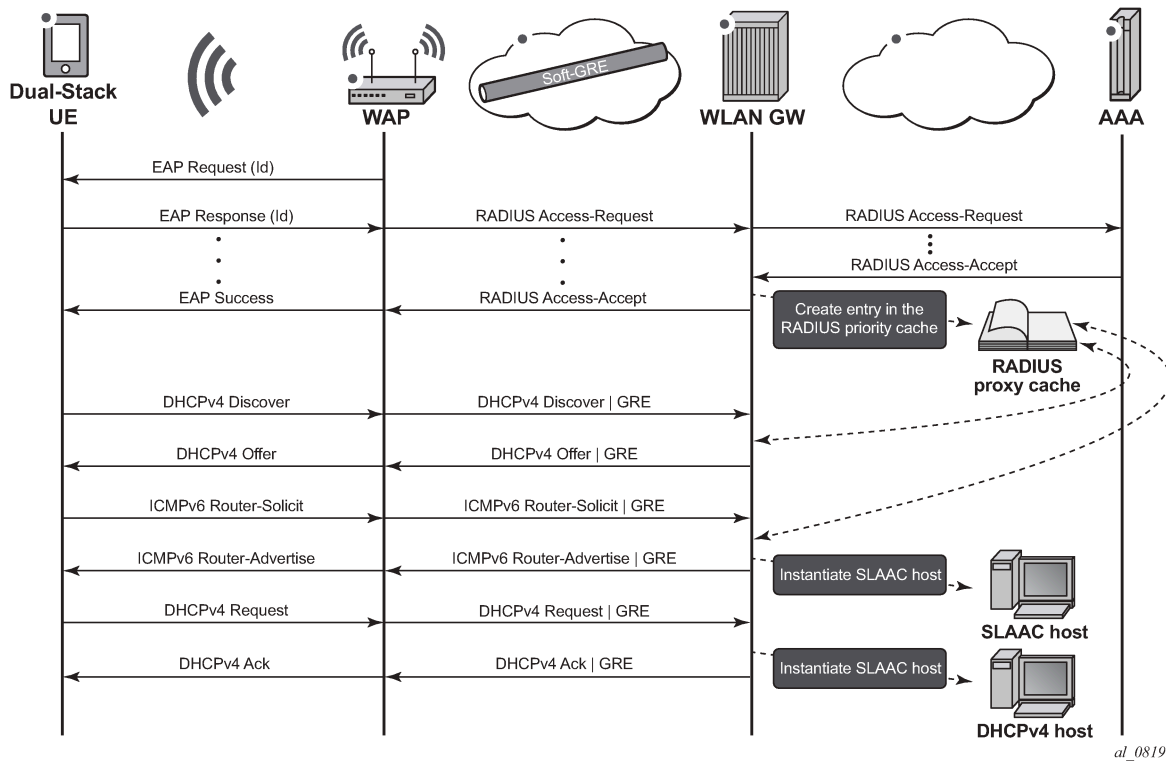
For SLAAC/64 hosts, the DNS information can be advertised with the recursive DNS server (RDNSS) option [RFC 6106] via SLAAC or via stateless DHCPv6 [RFC 3736]. For DHCPv6/128 hosts, the DNS information is advertised via DNS options for DHCPv6 [RFC3646]. If the AP supports a Lightweight DHCPv6 Relay Agent (LDRA), the WLAN-GW can learn the AP MAC address and the SSID that the UE connects to if the DHCPv6 Interface-Id option is in the format `<ap-mac>:<ssid>:{o (open) | s (secure)}`. This information can then be used in subsequent accounting messages.

The following three IPv4/v6 dual-stack UE IP address assignment models are available:

- DHCPv4 + SLAAC/64
- DHCPv4 + SLAAC/64 with DHCPv4 linking
- DHCPv4 + DHCPv6/128 IA_NA

In the DHCPv4 + SLAAC/64 model, DHCPv4 DORA and SLAAC/64 are processed independently of each other. If successful, two IPoE hosts are instantiated on the WLAN-GW for a particular UE: a DHCPv4 IPoE host and an IPv6 SLAAC/64 host.

Figure 177: DHCPv4 + SLAAC/64 model — closed SSID

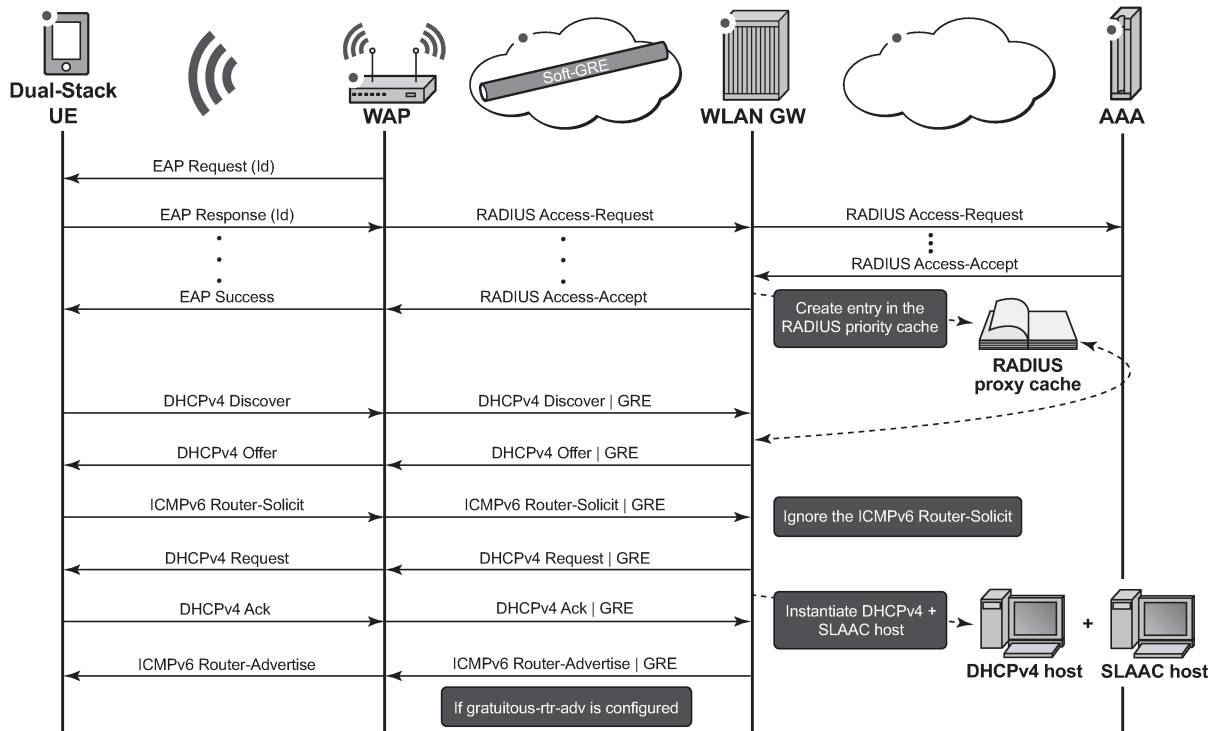


When the AP sends a RADIUS Accounting-Stop for a particular UE while track-accounting is enabled for Accounting-Stop messages, both the DHCPv4 IPoE host and the IPv6 SLAAC/64 host will be removed.

However, it is not always possible for the AP to send RADIUS accounting messages (for example, in the case of an open SSID). Because SLAAC has no renew or release mechanism, the only way to delete a SLAAC host is to determine which UE was stopped using the SLAAC prefix; for example, by using idle-timeout and/or by periodic Subscriber Host Connectivity Verification (SHCV).

In the DHCPv4 + SLAAC/64 with DHCPv4 linking model, a SLAAC/64 host is instantiated when a DHCPv4 host is instantiated. The state of the SLAAC/64 host is linked to the state of the DHCPv4 host. This is useful to speed up the removal of the SLAAC host in cases where the AP does not send RADIUS accounting messages. With DHCPv4 linking, when the DHCPv4 host is removed, also the SLAAC/64 host is removed.

Figure 178: DHCPv4 + SLAAC/64 with DHCPv4 linking model — closed SSID

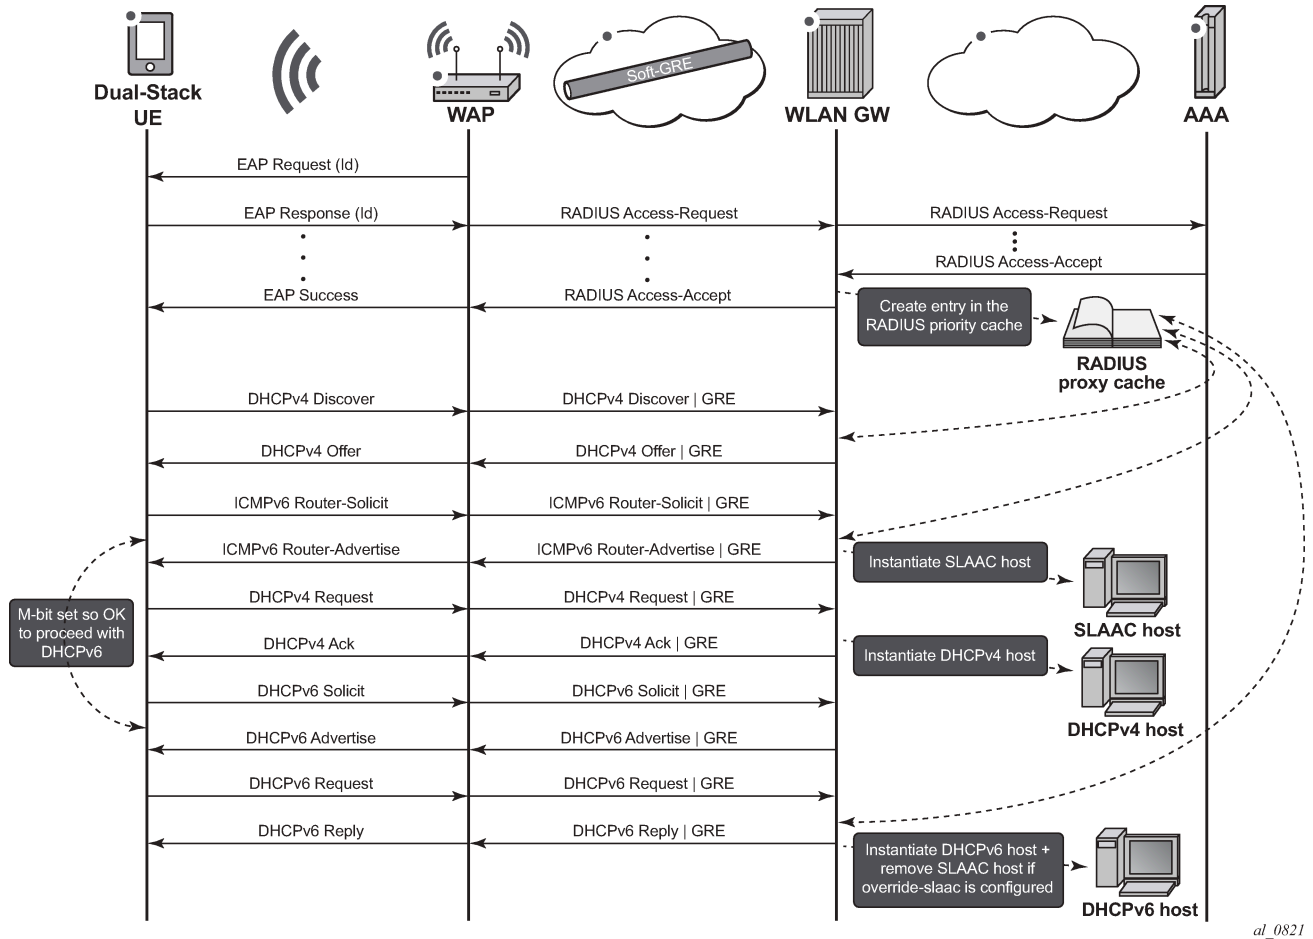


al_0820

In the DHCPv4 + DHCPv6/128 IA_NA model, similar to the DHCPv4 + SLAAC/64 model, DHCPv4 DORA and DHCPv6/128 IA_NA are processed independently of each other. SLAAC/64 is optional in this model although it is typically enabled because some UEs do not support DHCPv6.

UEs that do support stateful address auto-configuration only initiate DHCPv6 when they receive an ICMPv6 Router-Advertisement with the M-bit set (RFC 2462, *IPv6 Stateless Address Autoconfiguration*). Because the WLAN-GW does not know whether the UE supports DHCPv6, the WLAN-GW must include a SLAAC/64 prefix in the ICMPv6 Router-Advertisement, also for UEs that do support DHCPv6. Therefore, for a UE that does support DHCPv6, three IPoE hosts are instantiated in the WLAN-GW. To avoid this, the WLAN-GW can be configured to flush the SLAAC/64 host when a DHCPv6/128 IA_NA host is established. The UE should always prefer the DHCPv6/128 IA_NA address for sending data traffic above the IPv6 address derived from the SLAAC/64 prefix.

Figure 179: DHCPv4 + DHCPv6/128 IA_NA model — closed SSID



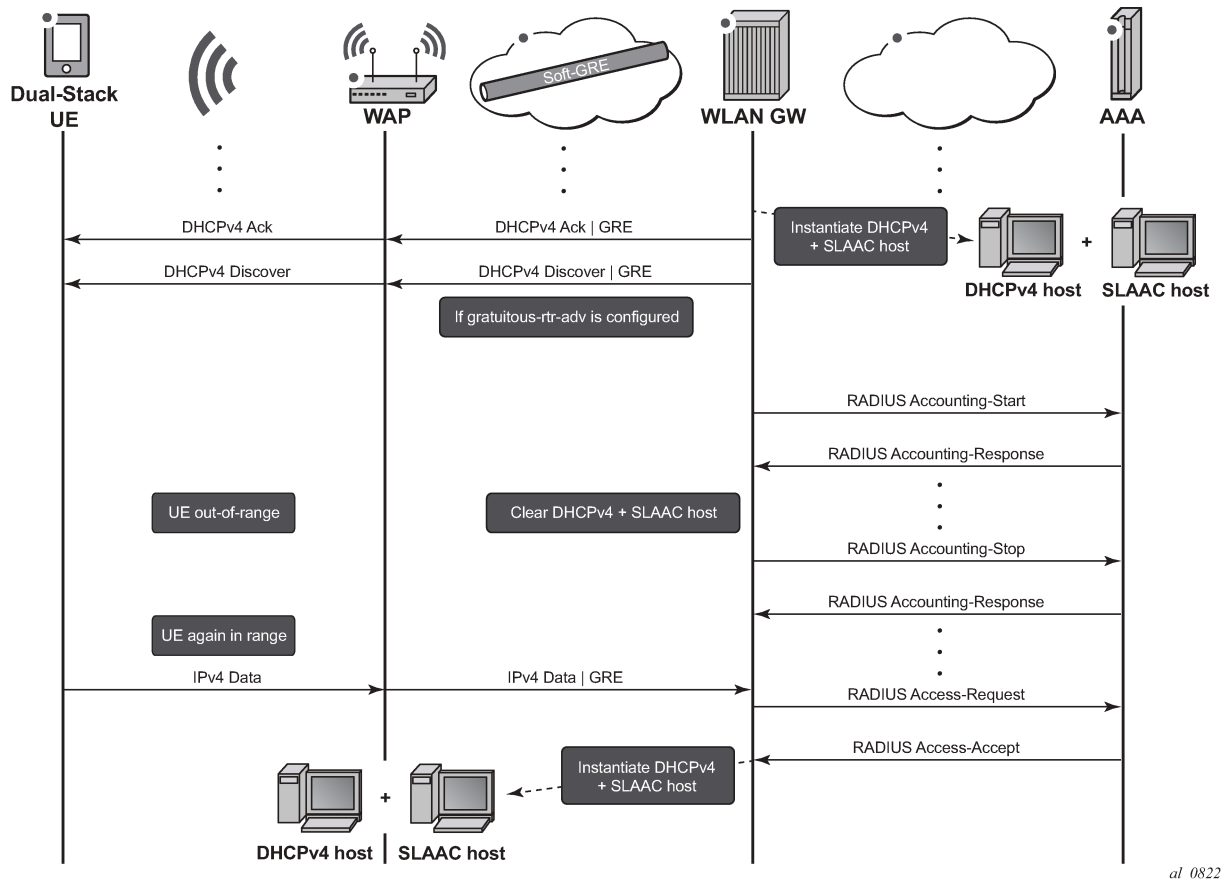
As with ESMv6, the SLAAC/64 prefix could come from the Local User Database (LUDB); this is typically not used because it requires configuring individual UE MAC addresses with their associated SLAAC/64 prefix. Alternatively, the SLAAC/64 prefix could come from RADIUS via the Framed-IPv6-Prefix attribute, or from a local SLAAC prefix pool that is referenced in the LUDB or from RADIUS via the Alc-SLAAC-IPv6-Pool attribute.

The DHCPv6/128 prefix comes from a DHCPv6 server that could be local (collocated with the WLAN-GW) or external, or from RADIUS via the Alc-IPv6-Address attribute. When a DHCPv6 server is used, the WLAN-GW relays the DHCPv6 messages between the UE and the local or external DHCPv6 server. If the DHCPv6/128 prefix comes from RADIUS/LUDB, the WLAN-GW must be configured as a DHCPv6 proxy server.

Note that IPv6 for WLAN-GW UEs is not supported in combination with certain other features, which include GPRS Tunneling Protocol (GTP(v2)) offload, migrant UEs, and data-triggered authentication (DTA).

Data-triggered authentication is not supported for IPv6 hosts, which means that an IPv6 packet from a UE for which no ESM context exists will not trigger RADIUS authentication. However, by using SLAAC/64 with DHCPv4 linking, the SLAAC host will be created together with the DHCPv4 host by successful completion of IPv4 data-triggered authentication. This requires the RADIUS Access-Accept to contain the necessary DHCPv4 and SLAAC/64 attributes.

Figure 180: DHCPv4 + SLAAC/64 with DHCPv4 linking model — DTA



al_0822

IPv6 is also not supported for migrant UEs, which means that ICMPv6 Router-Solicitation and DHCPv6 Solicit messages will be dropped by the WLAN-GW as long as the UE is in a migrant state. However, by using SLAAC/64 with DHCPv4 linking, when the UE becomes an ESM subscriber and a DHCPv4 host is created, a SLAAC/64 host is also created.

Configuration

Open versus closed SSID

The configuration examples in this chapter always refer to a closed SSID scenario. With an open SSID, the lookup in the RADIUS proxy cache is typically not configured. Instead, an authentication policy is directly referenced.

```
configure service vprn 2 customer 1 create
  subscriber-interface "sub-int-1" create
  group-interface "group-int-1" wlangw create
  authentication-policy "auth-pol-1"
  dhcp
  no user-db
```

```
        exit
    exit
exit
```

DHCPv4 + SLAAC/64 model

In this model, DHCPv4 and SLAAC/64 are enabled independently of each other. The autonomous flag tells the UE that the IPv6 prefix in the ICMPv6 Router-Advertisement can be used for SLAAC. The no on-link configuration commands the UE to always perform neighbor discovery for the WLAN-GW, even for destinations within the IPv6 prefix.

```
configure service vprn 2 customer 1 create
  subscriber-interface "sub-int-1" create
    address 10.255.255.254/8
    ipv6
      subscriber-prefixes
        prefix 2001:db8:ffff::/48 wan-host
      exit
    exit
  group-interface "group-int-1" wlangw create
    ipv6
      router-advertisements
        no managed-configuration
        no other-stateful-configuration
        dns-options
          include-dns
        exit
        prefix-options
          autonomous
          no on-link
        exit
        no shutdown
      exit
      router-solicit
        user-db "ludb-1"
        no shutdown
      exit
    exit
  ipoe-linking
    shutdown
  exit
  sap-parameters
    sub-sla-mgmt
      def-sub-id use-auto-id
      sub-ident-policy "policy-sub-ident-1"
    exit
  exit
  dhcp
    proxy-server
      emulated-server 172.16.0.1
      no shutdown
    exit
    lease-populate 10000
    user-db "ludb-1"
    no shutdown
  exit
  ip-mtu 1454
  wlan-gw
    gw-address 172.16.74.244
```



```

        gw-ipv6-address 2001:db8::1:1
        router 1
        tcp-mss-adjust 1400
        wlan-gw-group 1
        no shutdown
    exit
exit
exit
    no shutdown
exit

```

The SLAAC/64 prefix can come from the RADIUS server, as in the following RADIUS users file:

```

"user-1" Cleartext-Password := "pass-1"
    Alc-Subsc-ID-Str := "user-1",
    Alc-Subsc-Prof-Str := "sub-profile-1",
    Alc-SLA-Prof-Str := "sla-profile-1",
    Framed-IP-Address := 10.255.0.1,
    Alc-Primary-DNS := 67.138.54.100,
    Framed-IPv6-Prefix := 2001:db8:ffff::/64,
    Alc-IPv6-Primary-Dns := 2001:db8::8:8:8:8,
    Alc-IPv6-Secondary-Dns := 2001:db8::8:8:4:4

```

If the UE is successfully connected, two IPoE hosts will exist on the WLAN-GW.

```

*A:WLAN-GW # show service active-subscribers
=====
Active Subscribers
=====
-----
Subscriber user-1 (sub-profile-1)
-----
(1) SLA Profile Instance sap:[4/2/nat-out-ip:2049.4] - sla:sla-profile-1
-----
IP Address
-----
MAC Address      PPPoE-SID Origin
-----
10.255.0.1
b0:9f:ba:b9:40:f8 N/A      DHCP
2001:db8:ffff::/64
b0:9f:ba:b9:40:f8 N/A      IPoE-SLAAC
-----
Number of active subscribers : 1
-----

```

The trigger that created the SLAAC host and the origin is shown by issuing:

```

*A:WLAN-GW # show service id 2 slaac host detail
=====
SLAAC hosts for service 2
=====
Service ID      : 2
Prefix          : 2001:db8:ffff::/64
Interface Id    : N/A
Mac Address     : b0:9f:ba:b9:40:f8
Subscriber-interface : sub-int-1
Group-interface : group-int-1
SAP             : [4/2/nat-out-ip:2049.4]
Creation Time   : 2015/07/09 11:24:19
Persistence Key : N/A

```

```
IPoE Trigger      : rtr-solicit
Last Auth Time    : 2015/07/09 11:24:19
Inactivity Timer   : 0d 00:03:59
```

```
Sub-Ident         : "user-1"
Sub-Profile-String : "sub-profile-1"
SLA-Profile-String : "sla-profile-1"
App-Profile-String : ""
ANCP-String       : ""
Int Dest Id       : ""
Category-Map-Name : ""
```

```
Info origin       : radius
Pool              : ""
```

```
Primary-Dns       : 2001:db8::8:8:8:8
Secondary-Dns     : 2001:db8::8:8:4:4
```

```
Circuit Id        : N/A
Remote Id         : N/A
```

```
-----
Number of hosts : 1
=====
```

The SLAAC/64 prefix can also come from a local SLAAC prefix pool:

```
configure service vprn 2 customer 1 create
  dhcp6
    local-dhcp-server "local-dhcp-server-1" create
    use-pool-from-client
    pool "slaac-prefix-pool-1" create
    prefix 2001:db8:ffff:ffff::/64 wan-host create
    options
      dns-server 2001:db8::8:8:8:8
    exit
  exit
exit
no shutdown
exit
exit
exit
```

The subscriber interface must then be configured with local-address-assignment enabled:

```
configure service vprn 2 customer 1 create
  subscriber-interface "sub-int-1" create
  group-interface "group-int-1" wlangw create
  local-address-assignment
  ipv6
    client-application ipoe-slaac
    server "local-dhcp-server-1"
  exit
  no shutdown
exit
exit
exit
exit
exit
```

The origin of the SLAAC host then changes to:

```
*A:WLAN-GW # show service id 2 slaac host detail | match origin
Info origin      : localPool
```

DHCPv4 + SLAAC/64 with DHCPv4 linking model

In this model, DHCPv4 linking instantiates a SLAAC/64 host when a DHCPv4 host is instantiated. This requires **ipoe-linking** to be configured:

```
configure service vprn 2 customer 1 create
  subscriber-interface "sub-int-1" create
    address 10.255.255.254/8
    ipv6
      subscriber-prefixes
        prefix 2001:db8:ffff::/48 wan-host
      exit
    exit
  group-interface "group-int-1" wlangw create
    ipv6
      router-advertisements
        no managed-configuration
        no other-stateful-configuration
        dns-options
          include-dns
        exit
        prefix-options
          autonomous
          no on-link
        exit
        no shutdown
      exit
      router-solicit
        shutdown
      exit
    ipoe-linking
      gratuitous-rtr-adv
        no shutdown
    exit
    sap-parameters
      sub-sla-mgmt
        def-sub-id use-auto-id
        sub-ident-policy "policy-sub-ident-1"
      exit
    exit
    dhcp
      proxy-server
        emulated-server 172.16.0.1
        no shutdown
      exit
      lease-populate 10000
      user-db "ludb-1"
      no shutdown
    exit
    ip-mtu 1454
    wlan-gw
      gw-address 172.16.74.244
      gw-ipv6-address 2001:db8::1:1
      router 1
      tcp-mss-adjust 1400
      wlan-gw-group 1
      no shutdown
    exit
  exit
exit
no shutdown
```

```
exit
```

Note that DHCPv4 linking is mutually exclusive with ICMPv6 Router-Solicit handling. Configuring DHCPv4 linking while ICMPv6 Router-Solicit handling is still enabled results in the following error:

```
*A:WLAN-GW # configure service vprn 2 subscriber-interface "sub-int-1" group-interface "group-
int-1" ipoe-linking no shutdown
MINOR: SVCMMGR #1543 Can't enable linking if router solicit authentication is enabled
```

Similarly, enabling ICMPv6 Router-Solicit handling while DHCPv4 linking is still enabled, results in the following error:

```
*A:WLAN-GW # configure service vprn 2 subscriber-interface "sub-int-1" group-interface "group-
int-1" ipv6 router-solicit no shutdown
MINOR: SVCMMGR #1544 Can't enable router solicit authentication if linking is enabled
```

As with the DHCPv4 + SLAAC/64 model without DHCPv4 linking, if the UE is successfully connected, two IPoE hosts will exist on the WLAN-GW:

```
*A:WLAN-GW # show service active-subscribers
=====
Active Subscribers
=====
-----
Subscriber user-1 (sub-profile-1)
-----
(1) SLA Profile Instance sap:[4/2/nat-out-ip:2049.4] - sla:sla-profile-1
-----
IP Address          MAC Address          PPPoE-SID Origin
-----
10.255.0.1          b0:9f:ba:b9:40:f8 N/A          DHCP
2001:db8:ffff::/64  b0:9f:ba:b9:40:f8 N/A          IPoE-SLAAC
-----
Number of active subscribers : 1
-----
```

The trigger that created the SLAAC host and the origin is shown by issuing:

```
*A:WLAN-GW # show service id 2 slaac host detail
=====
SLAAC hosts for service 2
=====
Service ID          : 2
Prefix              : 2001:db8:ffff::/64
Interface Id        : N/A
Mac Address         : b0:9f:ba:b9:40:f8
Subscriber-interface : sub-int-1
Group-interface     : group-int-1
SAP                 : [4/2/nat-out-ip:2049.4]
Creation Time       : 2015/07/09 11:49:42
Persistence Key     : N/A

IPoE Trigger        : linking
Last Auth Time      : N/A
Inactivity Timer    : N/A
```

```
Sub-Ident      : "user-1"
Sub-Profile-String : "sub-profile-1"
SLA-Profile-String : "sla-profile-1"
App-Profile-String : ""
ANCP-String    : ""
Int Dest Id    : ""
Category-Map-Name : ""
```

```
Info origin    : radius
Pool           : ""
```

```
Primary-Dns    : 2001:db8::8:8:8:8
Secondary-Dns  : 2001:db8::8:8:4:4
```

```
Circuit Id     : N/A
Remote Id      : N/A
```

```
-----
Number of hosts : 1
=====
```

Clearing the DHCPv4 host results in both the DHCPv4 host and the SLAAC host being deleted.

```
*A:WLAN-GW # clear service id 2 dhcp lease-state mac b0:9f:ba:b9:40:f8
```

```
*A:WLAN-GW # show service active-subscribers
```

```
=====
Active Subscribers
=====
```

```
-----
No active subscribers found
-----
```

DHCPv4 + DHCPv6/128 IA_NA model

Because some UEs do not support DHCPv6, this model configures DHCPv4 + DHCPv6/128 IA_NA with SLAAC/64 enabled. To avoid having two IPv6oE hosts set up for the UEs that do support DHCPv6, the **allow-multiple-wan-addresses** and **override-slaac** parameters are both configured. The **allow-multiple-wan-addresses** allows handling of DHCPv6 when a SLAAC host exists already, and the **override-slaac** parameter removes the SLAAC host after successful assignment of an IPv6 address via DHCPv6:

```
configure service vprn 2 customer 1 create
  subscriber-interface "sub-int-1" create
    address 10.255.255.254/8
    ipv6
      subscriber-prefixes
        prefix 2001:db8:ffff::/48 wan-host
      exit
    exit
  group-interface "group-int-1" wlangw create
    ipv6
      allow-multiple-wan-addresses
      router-advertisements
        managed-configuration
        other-stateful-configuration
        dns-options
          include-dns
        exit
      prefix-options
        autonomous
```

```

        no on-link
        exit
        no shutdown
    exit
    router-solicit
        user-db "ludb-1"
        no shutdown
    exit
    dhcp6
        user-db "ludb-1"
        proxy-server
        no shutdown
    exit
    override-slaac
    exit
exit
ipoe-linking
shutdown
exit
sap-parameters
    sub-sla-mgmt
        def-sub-id use-auto-id
        sub-ident-policy "policy-sub-ident-1"
    exit
exit
dhcp
    proxy-server
        emulated-server 172.16.0.1
        no shutdown
    exit
    lease-populate 10000
    user-db "ludb-1"
    no shutdown
exit
ip-mtu 1454
wlan-gw
    gw-address 172.16.74.244
    gw-ipv6-address 2001:db8::1:1
    router 1
    tcp-mss-adjust 1400
    wlan-gw-group 1
    no shutdown
exit
exit
exit
no shutdown
exit

```

If the UE is successfully connected, two IPoE hosts will exist on the WLAN-GW:

```

*A:WLAN-GW # show service active-subscribers
=====
Active Subscribers
=====
-----
Subscriber user-1 (sub-profile-1)
-----
-----
(1) SLA Profile Instance sap:[4/2/nat-out-ip:2049.4] - sla:sla-profile-1
-----
IP Address
-----
MAC Address      PPPoE-SID Origin
-----

```

```
10.255.0.1          b0:9f:ba:b9:40:f8 N/A      DHCP
2001:db8:ffff::1/128 b0:9f:ba:b9:40:f8 N/A      IPoE-DHCP6
-----
Number of active subscribers : 1
-----
```

The origin of the DHCPv6 lease is shown by issuing:

```
*A:WLAN-GW # show service id 2 dhcp6 lease-state
=====
DHCP lease state table, service 2
=====
IP Address      Mac Address      Sap/Sdp Id      Remaining   Lease   MC
                  LeaseTime   Origin   Stdby
-----
2001:db8:ffff::1/128
                  b0:9f:ba:b9:40:f8 [4/2/nat-out-ip:20* 23h59m29s  Radius
-----
Number of lease states : 1
=====
* indicates that the corresponding row element may have been truncated.
```

Conclusion

The WLAN-GW supports IPv4/v6 dual-stack UEs. Although the IPv6 support for UEs can handle single-stack IPv6-only UEs, the UEs only have IPv6 support as an add-on to IPv4.

WiFi Aggregation and Offload — Migrant User Support

This chapter provides information about WiFi aggregation and offload for migrant user support configurations.

Topics in this chapter include:

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

Applicability

The information and configuration in this chapter is based on SR OS Release 11.0.R4.

Overview

The term “Migrant user” refers to user equipment (UEs) that connects to a WiFi network service set identification (SSID) but moves out of the range of the access point before initiating or completing authentication. For open-SSIDs, a migrant user may stay in the range of the access point just long enough to get a DHCP lease from the WLAN-GW. In actual WiFi deployments with portal authentication, it has been observed that a large percentage of users are migrant such that they get a DHCP lease but do not initiate or complete authentication.

Prior to this feature, an Enhanced Subscriber Management (ESM) host is created when the DHCP process completes. This results in the consumption of resources on both the CPM and IOM, limiting the ESM scale and performance for fully authenticated active users. This feature adds support to create an ESM host only after a user has been fully authenticated, either via a web portal or with an AAA server based on completing EAP exchange. In addition, with this feature L2-aware NAT is required, such that each UE gets the same shared configured inside IP address from the ISA via DHCP. Until a user is authenticated, forwarding of user traffic is constrained (via policy) to DNS and portal server access only.

Each user is allocated a small number of configured NAT outside ports to minimize public IP address consumption for unauthenticated users. Once the user is successfully authenticated, as indicated via a RADIUS Change of Authorization (COA) on successful portal authentication, an ESM host is created, and the L2-aware NAT is applied via a normal per-subscriber NAT policy. The inside IP address of the user does not change. The outside IP pool used is as per the NAT policy, and the L2-aware NAT could be 1:1 or NAT with larger number of outside ports than in the un-authenticated phase. If a user is already pre-authenticated (for example if the RADIUS server remembers the MAC address of the UE from a previous successful portal authentication) then the initial access-accept from RADIUS will trigger the creation of the ESM host.

Migrant user support for open SSID based on portal authentication

Sequence of events

1. DHCP is received from UE on ISA

Based on the DHCP and L2-aware NAT configuration on the ISA, an IP address is assigned to the user via DHCP. The DHCP and L2-aware NAT configuration is under the soft-gre node under the group-interface, or under vlan-tag range under the soft-gre node on the group-interface.

A different DHCP lease-time can be configured for an un-authenticated user (initial-lease-time) and an authenticated user (active-lease-time) for which an ESM host has been created. It is suggested that the initial lease be configured to a smaller value while the UE is migrant so that resources can be reclaimed quickly for a truly migrant user that will not complete authentication.

In addition to lease-times, DHCP return options, for example primary and secondary DNS and NBNS server addresses, that can be configured. This configuration can be per soft-GRE group interface or per VLAN range (where a VLAN tag corresponds to an SSID).

Up to 512 bytes of received DHCP options from clients are stored on the ISA. Once the DHCP ACK is sent back to the UE from the ISA, the UE will be created on the ISA in "migrant (or unauthenticated) state".

A configured L2-aware IP address is returned to each UE and a temporary L2-aware host is created on the anchor ISA for the UE. The NAT policy applicable to this L2-aware NAT for UE in migrant state is also configured under the group-interface (under soft-gre node or under vlan-tag range).

ARP requests coming from the UE in migrant state will be responded to from the ISA. The authentication to RADIUS is triggered on receiving the first Layer 3 data packet as opposed to on a DHCP DISCOVER.

2. Layer 3 data packet received on the ISA

The first Layer 3 packet (other than DHCP) will trigger RADIUS authentication from the ISA based on configured **isa-radius-policy** in the **configure>aaa** context. The user-name in the access-request is as per the user-name-format configured in the isa-radius-policy. By default it is the MAC address of the UE. The isa-radius-policy can be configured as the authentication policy under the soft-gre group-interface, or under specific VLAN tag ranges on the soft-gre group-interface. The latter allows for the use of a different authentication policy per SSID.

The RADIUS packets from the ISA are sourced with the IP address owned by the ISA. Each ISA in the WLAN-GW group gets an IP address from a set of contiguous addresses, the start of which is configurable in isa-radius-policy. The nas-ip-address sent in access-request message is configurable in the isa-radius-policy as the ISA's local IP address or the system IP address. In case the RADIUS server is behind a load-balancer which updates the source IP address of the RADIUS messages, the RADIUS server may use nas-ip-address to route the RADIUS response back. In this case the nas-ip-address should be configured as the ISA's IP address otherwise the response would incorrectly be routed to the CPM instead of the ISA.

The following debug output shows a RADIUS accept-request being sent to the RADIUS server on reception of first Layer 3 packet. The debug can be enabled by issuing:

```
debug router "management" radius packet-type authentication | accounting | coa

253 2013/08/07 20:58:35.53 UTC MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 2/1, SeqNo 11830"
```

```

Info:      anchor egressing frame
          radius-auth-req

IP/UDP:    from 192.168.0.2:1142 to 192.0.2.3:1812

RADIUS:    Access-Request (continued)
"

254 2013/08/07 20:58:35.53 UTC MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Access-Request(1) 192.168.0.2:1142 id 40 len 158 vrid 1
    NAS IP ADDRESS [4] 4 192.0.2.3
    NAS PORT TYPE [61] 4 Virtual(5)
    NAS PORT ID [87] 43 GRE rtr-3#lip-192.168.0.1#rip-192.0.2.1
    USER NAME [1] 17 00:0a:0a:00:01:00
    PASSWORD [2] 16 rCmhFboYeM2M8h0uBYJXJk
    CALLING STATION ID [31] 17 00:0a:0a:00:01:00
    VSA [26] 19 Alcatel(6527)
      CHADDR [27] 17 00:0a:0a:00:01:00
"

```

Received Layer 3 packets from the UE are handled as per the redirect-policy configured under the soft-gre group-interface or under applicable VLAN tag range on the soft-gre interface.

The redirect-policy is an IP ACL that should contain one more "forward rules" for traffic that should be forwarded while the UE is pending portal authentication. This typically should include traffic to and from DNS and web portal and is subjected to temporary L2-aware NAT. The redirect-policy also specifies the URL for redirecting triggered by http packets. The redirect-policy and/or the redirect URL can also be overridden via the RADIUS access-accept. Any other non-http traffic that does not match the forward rules is dropped.

While a UE is pending portal authentication no accounting messages are sent to the AAA server. Disconnect-Message from AAA server is supported while the UE is pending authentication.

3. Access-accept from RADIUS

The access-accept is received on the ISA from which the access-request was generated. The initial access-accept from RADIUS can indicate if a user needs to be authenticated by the portal or is a pre-authenticated user. The indication is based on inclusion of a "redirect policy" applicable to the user in a vendor specific attribute (VSA) (Alc-Wlan-Portal-Redirect, type = string) received from the RADIUS server. The access-accept can also include a redirect URL VSA (Alc-Wlan-Portal-Url, type = string) for the user. An empty Alc-Wlan-Portal_redirect VSA forces the use of the redirect policy that is locally specified under the soft-gre interface or under vlan-tag ranges on soft-gre interface. The redirect-policy is created under sub-mgmt node.

The UE state is changed to "portal" to indicate the UE is pending portal authentication and has limited access.

The following debug output shows the RADIUS accept-request being received from the RADIUS server and being processed by the WLAN-GW.

```

255 2013/08/07 20:58:35.61 UTC MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 2/1, SeqNo 11831
  Info:      anchor ingressing frame
          portal auth-accept

IP/UDP:    from 192.0.2.3:1812 to 192.168.0.2:1142

RADIUS:    Access-Accept (continued)
"

```

```
256 2013/08/07 20:58:35.62 UTC MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Access-Accept(2) id 40 len 64 from 192.0.2.3:1812 vrid 1
    VSA [26] 14 Alcatel(6527)
      SUBSC ID STR [11] 12 migrant_user
    VSA [26] 18 Alcatel(6527)
      WLAN PORTAL REDIRECT [172] 16 redirect-policy-1
"
```

The following command is used to display UE information on the ISA, including the state of the UE and the GRE tunnel to the AP through which the UE is connected.

```
*A:PE-1# tools dump wlan-gw ue
=====
Matched 1 session on Slot #2 MDA #1
=====
UE-Mac      : 00:0a:0a:00:01:00    UE-vlan     : N/A
UE IP Addr  : 10.0.0.10            Description  : Portal
UE timeout  : 288 sec              Auth-time   : 08/07/13 20:58:35
Tunnel MDA  : 2/2                  Tunnel Router : 10
MPLS label  : 3000                 Shaper      : Default
GRE Src IP Addr : 192.0.2.2        GRE Dst IP Addr : 192.168.0.1
Anchor SAP  : 2/1/nat-out-ip:2049.1
Last-forward : None                Last-move    : None
Rx Frames   : 0                    Rx Octets    : 0
Tx Frames   : 0                    Tx Octets    : 0
-----
No sessions on Slot #2 MDA #2 match the query
=====
```

If neither of the two redirect related VSAs are included in access-accept, then this indicates a "pre-authenticated user", and an ESM host is created for the subscriber with a subscriber-profile and other subscriber configuration from access-accept; from here normal ESM based forwarding occurs for the subscriber.

If a user is determined as a "pre-authenticated user", a message is generated to the CPM to create an ESM host. The information received from RADIUS in the access-accept message (for example subscriber-profile, app-profile etc) and the information from DHCP (for example the DHCP options) are passed in this message.

- COA from RADIUS

When user's credentials entered on the portal are successfully verified, the portal triggers the AAA server to generate COA to WLAN-GW. The COA serves as a trigger to create an ESM host. The COA MUST contain the subscriber-id and user-name, which are used as a key to identify the UE pending portal authentication.

The following shows an example debug of a COA being received from the AAA server.

```
248 2013/08/07 19:12:38.29 UTC MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Change of Authorization(43) 192.0.2.3:36776 id 124 len 96 vrid 1
    VSA [26] 19 Alcatel(6527)
      SUBSC ID STR [11] 17 00:0a:0a:00:01:00
    USER NAME [1] 17 00:0a:0a:00:01:00
    VSA [26] 10 Alcatel(6527)
      SLA PROF STR [13] 8 sla-profile-1
    VSA [26] 10 Alcatel(6527)
      SUBSC PROF STR [12] 8 sub-profile-1
"
```

When the COA is received and successfully processed, a COA-ACK is sent back to the AAA server. The COA message is passed to the CPM to create an ESM host. The information received in the COA, as well as stored information from DHCP (for example the DHCP options) are passed in this message. Once the ESM host is successfully created, the state of the UE on the ISA is changed accordingly to "ESM-user", and can be seen in the output of **tools dump WLAN-GW UE** command, as shown below.

The UE now has full access (and is not restricted by the original redirect-policy). The COA provides a reference to a subscriber profile that contains the NAT policy for an authenticated UE. The UE continues to keep the same inside L2-aware IP address that was provided originally via DHCP on the ISA. However, the NAT for an authenticated user could be an L2-aware 1:1 NAT or NAPT with a different outside pool and outside ports than the UE in migrant state. The ESM host that is created as described above will also result in the creation of a normal L2-aware host. The original temporary L2-aware host is retained for 10 seconds (and then deleted) to ensure the http response from the portal can be successfully routed back to the UE on the existing connection.

```
A:PE-1# tools dump wlan-gw ue
=====
Matched 1 session on Slot #2 MDA #1
=====
UE-Mac       : 00:0a:0a:00:01:00   UE-vlan      : N/A
UE IP Addr   : N/A                 Description   : ESM-user
UE timeout   : N/A                 Auth-time    : 08/07/13 19:12:38
Tunnel MDA   : 2/2                 Tunnel Router : 10
MPLS label   : 3000                Shaper       : 1
GRE Src IP Addr : 192.0.2.2        GRE Dst IP Addr : 192.168.0.1
Anchor SAP   : 2/1/nat-out-ip:2049.1
Last-forward  : 08/07/13 19:12:25   Last-move    : None
Rx Frames     : 1                   Rx Octets    : 88
Tx Frames     : 1                   Tx Octets    : 222
-----
No sessions on Slot #2 MDA #2 match the query
```

If UE goes out of range such that the idle timeout expires, the ESM host is deleted and an accounting-stop is sent to the AAA server. If a UE then comes back, and still has a valid DHCP lease, it may not send DHCP DISCOVER or REQUEST and continue to send data. The **data-triggered-ue-creation** command can be configured under soft-gre node on the group-interface (or under vlan-tag ranges on the group-interface) to trigger authentication and recreation of the ESM host for this UE.

The overall sequence of events to take a UE from migrant to authenticated state, where the forwarding of UE traffic is not restricted, is shown in [Figure 181: Sequence of events to establish and authenticate a migrant User](#).

Figure 181: Sequence of events to establish and authenticate a migrant User

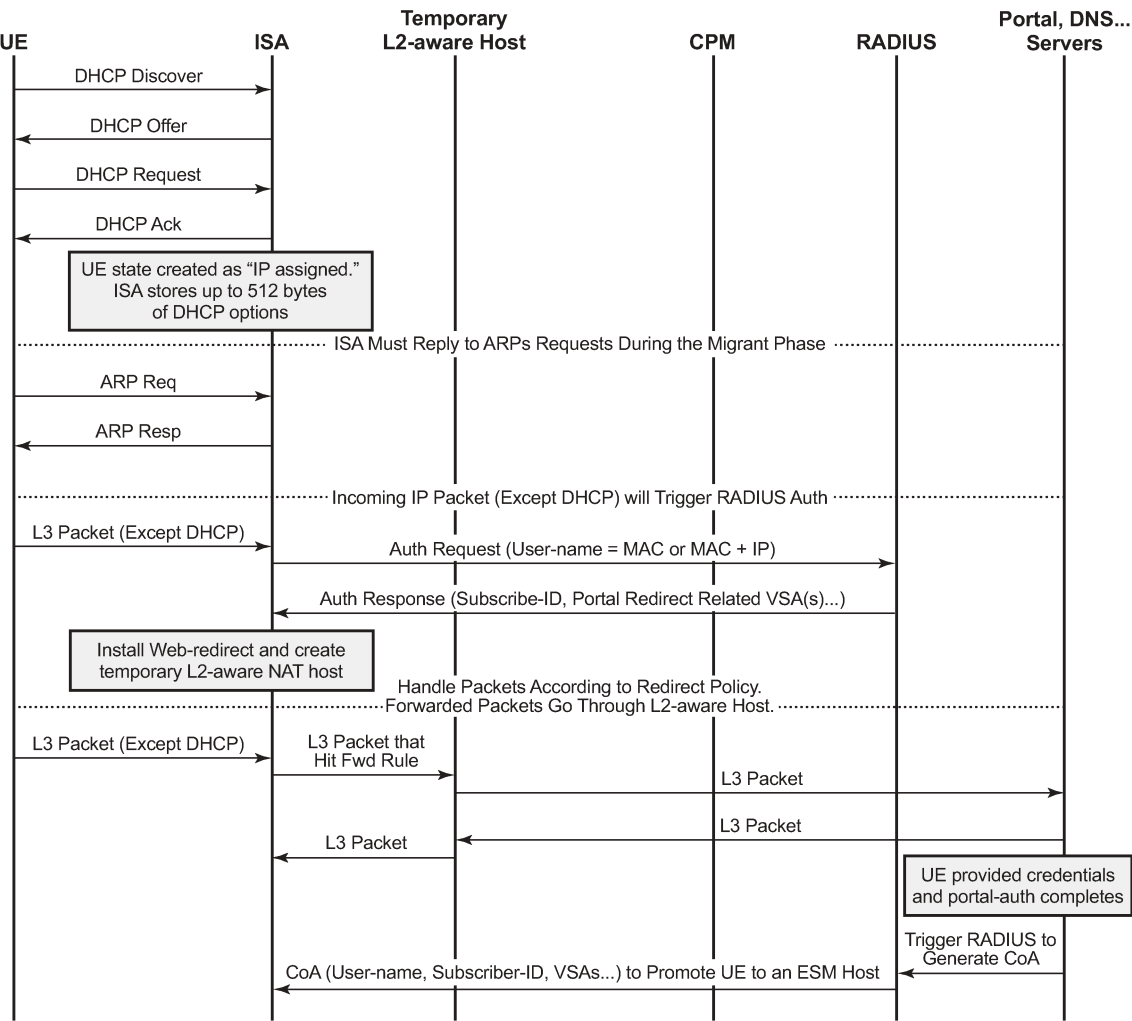
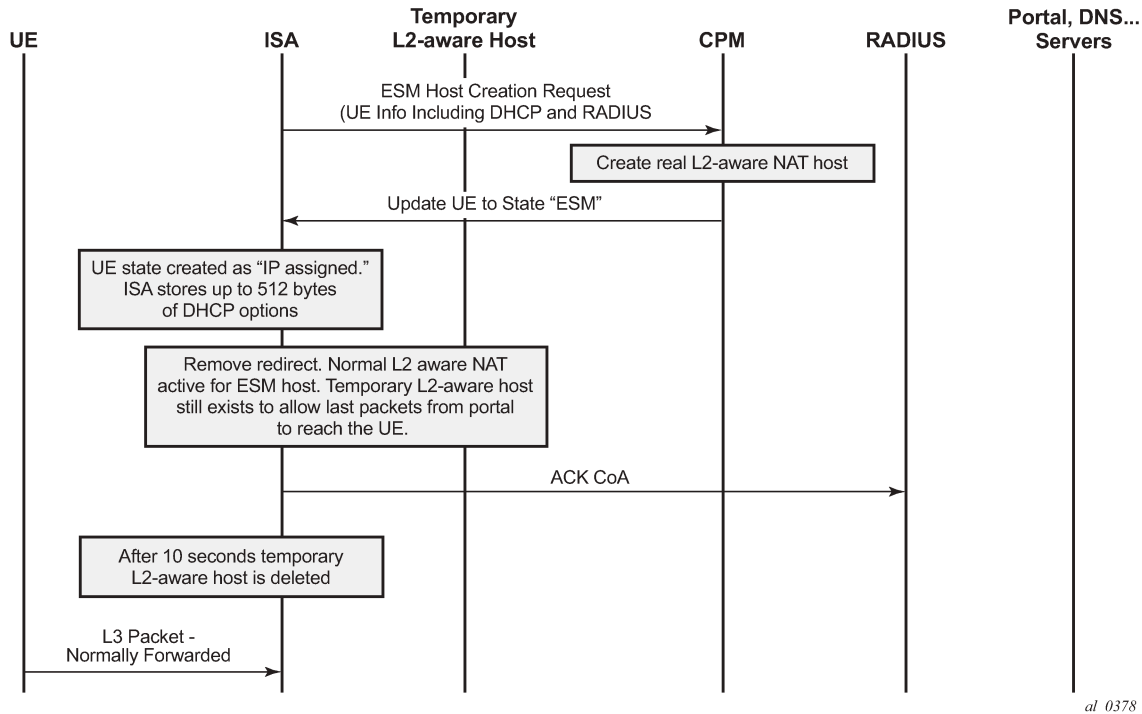


Figure 182: Sequence of events to establish and authenticate a migrant user (continued)



Configuration

The authentication-policy as shown below is used to configure a RADIUS server, and is applicable to the UEs in authenticated state.

```

subscriber-mgmt
  authentication-policy "authentication-1" create
    password "E40PedK6aqrEIpr2DEoJyVR8PQ3XkFF7" hash2
    radius-authentication-server
      source-address 192.0.2.1
      router "management"
      server 1 address 192.0.2.3 secret "6uuGli25Vtl49q0." hash2
    exit
  accept-authorization-change
  include-radius-attribute
    acct-session-id
    circuit-id
    remote-id
    nas-port-id
    nas-identifier
    nas-port-type
    pppoe-service-name
    dhcp-options
    dhcp-vendor-class-id
    access-loop-options
    mac-address
    called-station-id
    calling-station-id sap-string
  
```

tunnel-server-attrs

An isa-radius-policy is required for authentication from the ISA, as below – this contains the attributes to be sent in the access request message to the RADIUS server, which is also configured in this policy.

```
aaa
  isa-radius-policy "isa-policy-1" create
    nas-ip-address-origin isa-ip
    password "CA06ALDnhyBJERE4xnXoW15MQ/hu74x5nDE7F.0JxHM" hash2
    auth-include-attributes
      called-station-id
      calling-station-id
      circuit-id
      dhcp-options
      dhcp-vendor-class-id
      mac-address
      nas-identifier
      nas-port-id
      nas-port-type
      remote-id
    exit
  servers
    router 1
      source-address-range 192.168.0.2
      server 1 create
        authentication
        coa
        ip-address 192.0.2.3
        secret "CA06ALDnhyBJERE4xnXoW15MQ/hu74x5nDE7F.0JxHM" hash2
        no shutdown
      exit
    exit
  exit
exit
exit
```

The HTTP redirect policy is shown below, this is enforced on ISA while a UE is migrant and contains the configurations defining the forwarding of traffic in this state.

```
subscriber-mgmt
  http-redirect-policy "redirect-policy-1" create
  url "http://66.185.84.163"
  forward-entries
    dst-ip 192.168.1.1 protocol udp dst-port 53
    dst-ip 192.168.1.2 protocol udp dst-port 53
    dst-ip 66.185.84.163 protocol tcp dst-port 80
    dst-ip 10.0.0.1 protocol udp dst-port 67
    dst-ip 10.0.0.1 protocol udp dst-port 68
  exit
  exit
exit
```

The NAT pool configuration for migrant and authenticated UEs is shown below.

```
vprn 10 customer 1 create
  nat
    inside
      l2-aware
      address 10.0.0.1/24
    exit
  exit
  outside
```

```

        pool "migrant-pool-1" nat-group 1 type wlan-gw-anchor create
        address-range 192.168.2.0 192.168.2.255 create
        exit
        no shutdown
    exit
    pool "auth-pool-1" nat-group 1 type l2-aware create
    address-range 192.168.3.0 192.168.3.255 create
    exit
    no shutdown
exit
exit
exit
exit

```

The NAT policy for migrant UEs is as follows.

```

service
  nat
    nat-policy "migrant-policy" create
    pool "migrant-pool-1" router 1
    timeouts
      tcp-established min 1
    exit
  exit
exit
exit

```

Below is the NAT policy for authenticated UEs.

```

service
  nat
    nat-policy "nat-auth-policy-1" create
    pool "auth-pool-1" router 10
  exit
exit
exit

```

The migrant user configuration under the soft-gre group-interface within the VPRN service is shown below. This includes configuration for authentication, DHCP, and forwarding from the ISA, as defined in the sections above. The migrant user related configuration can be specified per VLAN tag (or range) under soft-gre interface, where each VLAN tag represents an SSID.

```

vprn 1 customer 1 create
  subscriber-interface "sub-int-1" create
  address 10.0.0.1/24
  group-interface "soft-gre-1" softgre create
    sap-parameters
      sub-sla-mgmt
        def-sla-profile "sla-profile-1"
        def-sub-id use-auto-id
        def-sub-profile "sub-profile-1"
        sub-ident-policy "sub_ident"
      exit
    exit
  dhcp
    proxy-server
      emulated-server 10.0.0.1
      lease-time hrs 1
      no shutdown
    exit
  trusted

```



```
        lease-populate 32767
        gi-address 10.0.0.1
        no shutdown
    exit

    authentication-policy "authentication-1"
    host-connectivity-verify

    soft-gre
        authentication
            authentication-policy "isa-policy-1"
        exit
        gw-address 192.168.0.1
        mobility
            hold-time 0
            trigger data iapp
        exit
        router 1
        wlan-gw-group 1
        vlan-tag-ranges
            range start 100 end 100
            authentication
                authentication-policy "isa-policy-1"
            exit
            data-triggered-ue-creation
            dhcp
                active-lease-time min 12
                initial-lease-time min 5
                l2-aware-ip-address 10.0.0.10
                primary-dns 192.168.1.1
                secondary-dns 192.168.1.2
                no shutdown
            exit
            http-redirect-policy "redirect-policy-1"
            nat-policy "migrant-policy"
        exit
    exit
    no shutdown
exit
exit
exit
exit
exit
```

Conclusion

Migrant user support is a useful feature that optimizes system resources (public IP addresses, ESM hosts, CPU processing, etc.) to provide the scale and performance required in live hot-spot and home-spot WiFi deployments at peak times.

WiFi Aggregation and Offload — Open SSID with DSM and Lawful Intercept

This chapter provides information about WiFi Aggregation and Offload — Open SSID with DSM and Lawful Intercept.

Topics in this chapter include:

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

Applicability

The information and configuration in this chapter are based on SR OS Release 12.0.R4.

Overview

WiFi Aggregation and Offload functionality for the 7750 SR has been supported in SR OS 10.0.R1 and later. This includes a RADIUS proxy server with RADIUS proxy cache and support for soft-GRE tunnels.

Initially, WLAN-GW subscribers were implemented using Enhanced Subscriber Management (ESM) on the Control Processing Module (CPM). To achieve higher scalability, subscribers can be implemented using Distributed Subscriber Management (DSM) on the Multi-Service Integrated Service Adapter (MS-ISA) cards, as described in this chapter.

Law enforcement agencies often require operators to provide a method of intercepting traffic from specific User Equipment (UE). This chapter describes a method of configuring Lawful Intercept (LI) for a DSM UE.

Starting with Release 12.0.R4, DSM can be used for higher scalability by instantiating subscribers on the MS-ISA cards, even after authentication, instead of creating them on the CPM as when using ESM. Therefore, the maximum number of UEs per WLAN-GW, and other performance factors such as setup rate, are higher. When using DSM, commands that are different from those used with ESM are used to monitor the UEs. These commands are similar to those used by the previously available migrant users feature, which only instantiated the users on the MS-ISA cards prior to their authentication.


Lawful intercept can also be configured to intercept traffic to and from a UE. For security reasons, the configuration for LI can be kept separate and invisible to the regular admin user of the 7750 SR, even if this admin user has full admin access. In this situation, any information about the intercept is only available to the LI user. An example LI configuration is shown in this document, with the intercept configured using CLI. It is also possible to trigger DSM LI using RADIUS (in Access-Accept or Change of Authorization (CoA) messages). The RADIUS attributes are encrypted so that knowledge of the intercept cannot be gained by examining a packet capture.

Configuration

The WiFi offload scenario with open SSID, DSM and LI shown in [Figure 183: WiFi Offload Scenario with Open SSID, DSM and LI](#) has following characteristics:

- Open SSID with web portal authentication
- DSM with fixed IPv4 address for all UEs, with L2-aware NAT
- Access Point (AP) access in Virtual Private Routed Network (VPRN) 1000
- UEs terminated in VPRN 2004
- Lawful Intercept with separate user account for monitoring or configuring intercepts using the CLI

Figure 183: WiFi Offload Scenario with Open SSID, DSM and LI



WLAN-GW

Note that the uplink interface, Interior Gateway Protocol (IGP), and system configuration is outside the scope of this document.

The following Card and Media Dependent Adapter (MDA) configuration only shows the WLAN-Input/Output Module (IOM). An IOM card containing two MS-ISA cards provides the WLAN-GW functionality. The MDA type for the ISA cards is isa-bb, the same type that is used for NAT.

```
*A:WLAN-GW# /configure card 3
*A:WLAN-GW>config>card# info
-----
card-type iom3-xp
mda 1
mda-type isa-bb
no shutdown
exit
mda 2
mda-type isa-bb
```

```

        no shutdown
    exit
no shutdown
-----

```

The following ISA configuration applies.

```

*A:WLAN-GW# /configure isa
*A:WLAN-GW>config>isa# info
-----
    wlan-gw-group 1 create
        active-iom-limit 1
        iom 3
        no shutdown
    exit
-----

```

The following NAT configuration provides an outside pool of type wlan-gw-anchor required to support DSM.

```

*A:WLAN-GW# /configure router nat
*A:WLAN-GW>config>router>nat# info
-----
    outside
        pool "WiFi-4-dsm" nat-group 1 type wlan-gw-anchor create
        port-reservation ports 15
        address-range 10.0.40.0 10.0.40.255 create
        exit
        no shutdown
    exit
exit
-----

*A:WLAN-GW# /configure service nat
*A:WLAN-GW>config>service>nat# info
-----
    nat-policy "WiFi-4-dsm" create
        pool "WiFi-4-dsm" router Base
    exit
-----

```

The following Authentication, Authorization and Accounting (AAA) configuration contains an isa-radius-policy. The nas-ip-address-origin parameter selects the IP address sent as the Network Access Server (NAS) IP Address attribute in the ISA RADIUS requests. The source-address-range configures the IP address used by the first MS-ISA card on this WLAN-GW to send and receive RADIUS messages. Other WLAN IOM MS-ISA cards will get consecutive IP addresses in order of slot number. The password and secret configured here have to match the RADIUS server configuration.

```

*A:WLAN-GW# /configure aaa
*A:WLAN-GW>config>aaa# info
-----
    isa-radius-policy "IRS_4" create
        nas-ip-address-origin isa-ip
        password "7USmr6f7JkxD5zb3MeEZnjf1BSqaZkcH" hash2
        servers
            access-algorithm round-robin
            router "Base"
            source-address-range 10.10.186.1
            server 1 create
                authentication
                coa
                ip-address 10.43.186.1

```

```

                secret "7USmr6f7JkxD5zb3MeEZnjf1BSqaZkcH" hash2
                no shutdown
            exit
        exit
    exit
-----

```

The following subscriber management configuration contains the http-redirect-policy for redirecting newly connected UEs to the web portal, and allowing only traffic to the web portal IP address and the Domain Name Server (DNS) server.

```

*A:WLAN-GW# /configure subscriber-mgmt
*A:WLAN-GW>config>subscr-mgmt# info
-----
    http-redirect-policy "WiFi-4-dsm-redirect" create
        url "http://portal1.3ls.net/portal4.php?mac=$MAC"
        forward-entries
            dst-ip 10.43.186.1 protocol tcp dst-port 80 prefix-length 32
            dst-ip 10.43.186.43 protocol udp dst-port 53 prefix-length 32
        exit
    exit
-----

```

The following policy configuration is used for exporting required routes, including the address of the APs, the outside NAT prefixes, and the ISA RADIUS source IP addresses.

```

*A:WLAN-GW# /configure router policy-options
*A:WLAN-GW>config>router>policy-options# info
-----
    prefix-list "WiFi"
        prefix 10.0.0.0/16 longer
        prefix 10.10.186.0/24 longer
    exit
    prefix-list "WiFi-APs"
        prefix 10.1.0.0/16 longer
    exit
    policy-statement "toisis"
        entry 10
            from
                prefix-list "WiFi" "WiFi-APs"
            exit
            action accept
            exit
        exit
    exit
    policy-statement "WiFi-APs"
        entry 10
            from
                prefix-list "WiFi-APs"
            exit
            action accept
            exit
        exit
    exit
-----

```

The following configuration is used for exporting the outside NAT prefixes and the ISA RADIUS source IP addresses to ISIS.

```

*A:WLAN-GW# /configure router isis
*A:WLAN-GW>config>router>isis# info

```

```
-----  
export "toisis"  
-----
```

The following configures VPRN 1000 for AP connectivity, where the AP prefix is exported to the Global Route Table (GRT) so that it can be managed from servers reachable through the Base router.

```
*A:WLAN-GW# /configure service vprn 1000  
*A:WLAN-GW>config>service>vprn# info  
-----  
route-distinguisher 65400:1000  
interface "toAP" create  
address 10.1.0.1/24  
sap 1/1/7 create  
exit  
exit  
grt-lookup  
enable-grt  
static-route 0.0.0.0/0 grt  
exit  
export-grt "WiFi-APs"  
exit  
-----
```

The following configures VPRN 2004 for UE termination, with distributed-sub-mgmt enabled and the ISA RADIUS policy configured under vlan-tag-ranges.

```
*A:WLAN-GW# /configure service vprn 2004  
*A:WLAN-GW>config>service>vprn# info  
-----  
description "Open WiFi with DSM"  
route-distinguisher 65400:2004  
subscriber-interface "SI4" create  
address 10.0.4.1/24  
group-interface "GI4" wlangw create  
wlan-gw  
gw-address 192.4.4.4  
mobility  
trigger data iapp  
exit  
router 1000  
wlan-gw-group 1  
vlan-tag-ranges  
range default  
authentication  
authentication-policy "IRS_4"  
exit  
dhcp  
active-lease-time min 5  
initial-lease-time min 5  
l2-aware-ip-address 10.1.4.2  
primary-dns 10.43.186.43  
secondary-dns 10.44.186.44  
no shutdown  
exit  
distributed-sub-mgmt  
no shutdown  
exit  
nat-policy "WiFi-4-dsm"  
exit  
exit  
no shutdown
```

```
        exit
    exit
exit
nat
    inside
        l2-aware
        address 10.1.4.1/24
    exit
exit
exit
wlan-gw
exit
no shutdown
-----
```

The following LI user configuration allows user Lladmin to configure and view the Lawful Intercept configuration.

```
*A:WLAN-GW# /configure system security
*A:WLAN-GW>config>system>security# info
-----
    profile "li"
        default-action deny-all
        li
        entry 1
            match "back"
            action permit
        exit
        entry 2
        exit
        entry 10
            match "configure system security"
            action permit
        exit
        entry 20
            match "configure li"
            action permit
        exit
        entry 30
            match "show li"
            action permit
        exit
        entry 40
            match "file"
            action permit
        exit
        entry 50
            match "info"
            action permit
        exit
        entry 60
            match "admin display-config"
            action permit
        exit
        entry 70
            match "tools perform security"
            action permit
        exit
        entry 80
            match "tools dump li wlan-gw ue"
            action permit
        exit
        entry 100
```

```

        match "exit"
        action permit
    exit
exit
user "LIadmin"
password "$2y$10$Yp3sQZpG1bg6K3CeQoCHi.wyB0j7ts5/tsY/nqb0bbFjuFZ9G5wsi"
access console li
console
    no member "default"
    member "li"
exit
exit
-----

```

The following mirror configuration of type ip-only forwards intercepted traffic to a server.

```

*A:WLAN-GW# /configure mirror
*A:WLAN-GW>config>mirror# info
-----
    mirror-dest 199 type ip-only create
    encap
        layer-3-encap ip-udp-shim create
        gateway create
            ip src 10.10.10.186 dest 10.43.186.43
            udp src 3199 dest 3199
        exit
    exit
exit
no shutdown
exit
-----

```

The following configures a BOF, with li-local-save, a local LI config file, and li-separate, ensuring that only the LI user can view or modify LI parameters.

```

*A:WLAN-GW# show bof
=====
BOF (Memory)
=====
    li-local-save
    li-separate

```

The following LI source configuration to intercept the UE can only be configured or viewed by user LIadmin. The configuration is saved in cf3:\li.cfg, encrypted. The intercept-id and session-id will appear in the LI packets, which can be decoded in Wireshark using Decode As, Jmirror.

```

*A:WLAN-GW# /configure li
*A:WLAN-GW>config>li# info
-----
#-----
echo "LI Log Configuration"
#-----
    log
    exit
#-----
echo "LI Filter Lock State Configuration"
#-----
    li-filter-lock-state locked
#-----
echo "LI Mirror Source Configuration"
#-----

```



```

    li-source 199
    wlan-gw
        dsm-subscriber mac 68:7f:74:8b:3d:d7
        intercept-id 1
        session-id 199
    exit
    exit
    no shutdown
    exit
-----

```

Freeradius

This default configuration section sets the VSA Alc-Wlan-Ue-Creation-Type with value 1, which triggers the creation of a DSM host (value 0 is ESM). The Nas-Ip-Address is returned in the Alc-Wlan-Portal-URL to tell the web portal which MS-ISA address should receive the RADIUS CoA request:

```

/etc/freeradius/users:

DEFAULT      Auth-Type := Local, User-Password := "alcatel", user-name=~"^.*$"
              Alc-Subsc-ID-Str = "%{User-Name}",
              Alc-Wlan-Ue-Creation-Type = 1,
              Alc-Wlan-Portal-Redirect = "WiFi-4-dsm-redirect",
              Alc-Wlan-Portal-URL = "http://portal1.3ls.net/portal4.php?nas=%{Nas-Ip-
Address}&mac=%{User-Name}&ssid=WiFi-4",

```

As an alternative to configuring the LI for the UE in the CLI, the following RADIUS attributes can be sent in the Access-Accept and CoA.

```

    Alc-LI-Action = "enable",
    Alc-LI-Destination = "199",
    Alc-LI-Intercept-Id = 1,

```

In /etc/freeradius/clients.conf, each ISA is a client.

```

client 10.10.186.1 {
    secret      = alcatel
    shortname   = WLAN-GW-ISA1
}
client 10.10.186.2 {
    secret      = alcatel
    shortname   = WLAN-GW-ISA2
}

```

A RADIUS CoA sent during a successful portal login makes the following UE a DSM subscriber with full access.

```

echo "User-Name='.$mac.',Alc-Wlan-Ue-Creation-Type=1,Alc-Subsc-Prof-Str="SUBP_4",Alc-SLA-
Prof-Str=SLAP_4,Alc-Primary-Dns = 10.43.186.43," | /usr/bin/radclient -x -r 1 -t 2 '.$nas.' coa
alcatel

```

Access Points

The following must be configured on the Access Point as a minimum:

- IP address 10.1.1.10/24
- Default route to 10.1.1.1
- Open SSID WiFi-4 mapped to VLAN 40
- Soft-GRE tunnel with destination 192.4.4.4, with VLAN 40 mapped to this tunnel

Show Commands

The following commands show the status of the UEs. For DSM users, the UEs are displayed using a tools command. Before portal authentication, the UE is in Portal state.

```
*A:WLAN-GW# /tools dump wlan-gw ue
=====
Matched 1 session on Slot #3 MDA #1
=====
UE-Mac       : 68:7f:74:8b:3d:d7    UE-vlan      : 40
UE IP Addr   : 10.1.4.2             UE timeout   : 293 sec
Description  : Portal
Auth-time    : 09/08/2014 11:30:34
Tunnel MDA   : 3/2                 Tunnel Router : 1000
MPLS label   : 40                  Shaper       : Default
Tunnel Src IP : 10.1.0.10           Tunnel Dst IP : 192.4.4.4
Tunnel Type  : GRE
Anchor SAP   : 3/1/nat-out-ip:2049.3
AP-Mac       : Unknown             AP-RSSI      : Unknown
AP-SSID      : Unknown
Last-forward  : 09/08/2014 11:30:39 Last-move    : None
Session Timeout : None             Idle Timeout  : N/A
Acct Update   : None               Acct Interval : N/A
Acct Session-Id : N/A
Acct Policy   : N/A
NAT Policy    : WiFi-4-dsm
Redirect Policy : WiFi-4-dsm-redirect
IP Filter     : N/A
App-profile   : N/A
Rx Oper PIR   : N/A                Rx Oper CIR   : N/A
Tx Oper PIR   : N/A                Tx Oper CIR   : N/A
Rx Frames     : 204                 Rx Octets     : 17381
Tx Frames     : 78                  Tx Octets     : 67793
-----
No sessions on Slot #3 MDA #2 match the query
```

After login to the web portal, the UE transitions to a DSM-user and the Redirect Policy is removed,

```
*A:WLAN-GW# /tools dump wlan-gw ue
=====
Matched 1 session on Slot #3 MDA #1
=====
UE-Mac       : 68:7f:74:8b:3d:d7    UE-vlan      : 40
UE IP Addr   : 10.1.4.2             UE timeout   : 284 sec
Description  : DSM-user
Auth-time    : 09/08/2014 11:30:45
Tunnel MDA   : 3/2                 Tunnel Router : 1000
MPLS label   : 40                  Shaper       : Default
Tunnel Src IP : 10.1.0.10           Tunnel Dst IP : 192.4.4.4
Tunnel Type  : GRE
Anchor SAP   : 3/1/nat-out-ip:2049.3
AP-Mac       : Unknown             AP-RSSI      : Unknown
```

```

AP-SSID      : Unknown
Last-forward  : 09/08/2014 11:30:47  Last-move     : None
Session Timeout : None                Idle Timeout   : N/A
Acct Update    : None                Acct Interval  : N/A
Acct Session-Id : N/A
Acct Policy    : N/A
NAT Policy     : WiFi-4-dsm
Redirect Policy : N/A
IP Filter      : N/A
App-profile    : N/A
Rx Oper PIR    : N/A                Rx Oper CIR    : N/A
Tx Oper PIR    : N/A                Tx Oper CIR    : N/A
Rx Frames      : 273                Rx Octets      : 23186
Tx Frames      : 122                Tx Octets      : 108083

```

=====

No sessions on Slot #3 MDA #2 match the query

User Lladmin can view the configured intercept,

```

*A:WLAN-GW>config>li# /tools dump li wlan-gw ue
=====
Matched 1 session on Slot #3 MDA #1
=====
UE-Mac       : 68:7f:74:8b:3d:d7    Mirror Service : 199
LI Intercept-Id : 1                LI Session-Id  : 199
-----
=====
No sessions on Slot #3 MDA #2 match the query

```

Debug

In this example, the following debug configuration applies.

```

debug
  mirror-source 99
    port 1/1/7 egress ingress
    port 1/1/9 egress ingress
    no shutdown
  exit
  wlan-gw
    group 1
      ue 68:7f:74:8b:3d:d7 packet dhcp radius
    exit
  exit
exit

```

The debug trace starts with DHCP.

```

150 2014/09/08 11:30:31.93 EDT MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 3/1, SeqNo 3528
  Info:      anchor ingressing frame
           received upstream from tunnel

Ethernet: from 68:7f:74:8b:3d:d7 to ff:ff:ff:ff:ff:ff (ethertype: 0x0800)

IP/UDP:   from 0.0.0.0 (port 68) to 255.255.255.255 (port 67)

DHCP:
ciaddr: 0.0.0.0          yiaddr: 0.0.0.0

```

```
siaddr: 0.0.0.0      giaddr: 0.0.0.0
chaddr: 68:7f:74:8b:3d:d7  xid: 0xca073331

DHCP options:
[53] Message type: Discover
[61] Client id: (hex) 01 68 7f 74 8b 3d d7
[12] Host name: W81VM
[60] Class id: MSFT 5.0
[55] Param request list: len = 13
      1 Subnet mask
     15 Domain name
      3 Router
      6 Domain name server
     44 NETBIOS name server
     46 NETBIOS type
     47 NETBIOS scope
     31 Router discovery
     33 Static route
    121 Unknown option
    249 Unknown option
    252 Unknown option
     43 Vendor specific
[255] End
"

151 2014/09/08 11:30:31.93 EDT MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 3/2, SeqNo 706
  Info:      tunnel ingressing frame
          received downstream from anchor

Ethernet: from 00:00:00:02:02:02 to 68:7f:74:8b:3d:d7 (ethertype: 0x0800)

IP/UDP:   from 10.1.4.1 (port 67) to 10.1.4.2 (port 68)

DHCP:
ciaddr: 0.0.0.0      yiaddr: 10.1.4.2
siaddr: 10.1.4.1      giaddr: 0.0.0.0
chaddr: 68:7f:74:8b:3d:d7  xid: 0xca073331

DHCP options:
[53] Message type: Offer
[54] DHCP server addr: 10.1.4.1
[1] Subnet mask: 255.255.255.0
[3] Router: 10.1.4.1
[51] Lease time: 300
[6] Domain name server: length = 8
    10.43.186.43
    10.44.186.44
[255] End
"

152 2014/09/08 11:30:32.09 EDT MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 3/1, SeqNo 3529
  Info:      anchor ingressing frame
          received upstream from tunnel

Ethernet: from 68:7f:74:8b:3d:d7 to ff:ff:ff:ff:ff:ff (ethertype: 0x0800)

IP/UDP:   from 0.0.0.0 (port 68) to 255.255.255.255 (port 67)

DHCP:
ciaddr: 0.0.0.0      yiaddr: 0.0.0.0
siaddr: 0.0.0.0      giaddr: 0.0.0.0
chaddr: 68:7f:74:8b:3d:d7  xid: 0xca073331
```

```
DHCP options:
[53] Message type: Request
[61] Client id: (hex) 01 68 7f 74 8b 3d d7
[50] Requested IP addr: 10.1.4.2
[54] DHCP server addr: 10.1.4.1
[12] Host name: W81VM
[81] client FQDN: rcode1: 0, rcode2: 0, domain name = (hex) 00 57 38 31 56
4d
[60] Class id: MSFT 5.0
[55] Param request list: len = 13
      1 Subnet mask
      15 Domain name
      3 Router
      6 Domain name server
      44 NETBIOS name server
      46 NETBIOS type
      47 NETBIOS scope
      31 Router discovery
      33 Static route
      121 Unknown option
      249 Unknown option
      252 Unknown option
      43 Vendor specific
[255] End
"

153 2014/09/08 11:30:32.09 EDT MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 3/2, SeqNo 707
  Info:      tunnel ingressing frame
         received downstream from anchor

Ethernet: from 00:00:00:02:02:02 to 68:7f:74:8b:3d:d7 (ethertype: 0x0800)

IP/UDP:   from 10.1.4.1 (port 67) to 10.1.4.2 (port 68)

DHCP:
ciaddr: 0.0.0.0          yiaddr: 10.1.4.2
siaddr: 10.1.4.1          giaddr: 0.0.0.0
chaddr: 68:7f:74:8b:3d:d7 xid: 0xca073331

DHCP options:
[53] Message type: Ack
[54] DHCP server addr: 10.1.4.1
[1] Subnet mask: 255.255.255.0
[3] Router: 10.1.4.1
[51] Lease time: 300
[58] Renew timeout: 150
[59] Rebind timeout: 263
[6] Domain name server: length = 8
    10.43.186.43
    10.44.186.44
[255] End
"
```

RADIUS authentication is triggered by the first data packet.

```
154 2014/09/08 11:30:34.76 EDT MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 3/1, SeqNo 3563
  Info:      anchor egressing frame
         radius-auth-req

IP/UDP:   from 10.10.186.1 (port 1082) to 10.43.186.1 (port 1812)
```

```
RADIUS: Access-Request (continued)
"

155 2014/09/08 11:30:34.76 EDT MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Access-Request(1) 10.10.186.1:1082 id 45 len 126 vrid 1
    NAS IP ADDRESS [4] 4 10.10.186.1
    USER NAME [1] 17 68:7f:74:8b:3d:d7
    PASSWORD [2] 16 MvqAtmA0vSeeWgNIGyT/t.
    CALLING STATION ID [31] 17 68:7f:74:8b:3d:d7
    CALLED STATION ID [30] 17 00:00:00:00:00:00
    VSA [26] 19 Alcatel(6527)
    CHADDR [27] 17 68:7f:74:8b:3d:d7
"

156 2014/09/08 11:30:34.76 EDT MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 3/1, SeqNo 3564
  Info: anchor ingressing frame
        portal auth-accept

  IP/UDP: from 10.43.186.1 (port 1812) to 10.10.186.1 (port 1082)

  RADIUS: Access-Accept (continued)
"

157 2014/09/08 11:30:34.76 EDT MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Access-Accept(2) id 45 len 258 from 10.43.186.1:1812 vrid 1
    VSA [26] 19 Alcatel(6527)
    SUBSC ID STR [11] 17 68:7f:74:8b:3d:d7
    VSA [26] 20 Alcatel(6527)
    VSA [26] 20 Alcatel(6527)
    VSA [26] 20 Alcatel(6527)
    VSA [26] 6 Alcatel(6527)
    WLAN UE CREATION TYPE [184] 4 1
    VSA [26] 25 Alcatel(6527)
    WLAN PORTAL REDIRECT [172] 23 WiFi-4-dsm-redirect
    VSA [26] 86 Alcatel(6527)
    WLAN PORTAL URL [173] 84 http://portal1.3ls.net/portal4.php?nas=10.10.186.
1&mac=68:7f:74:8b:3d:d7&ssid=WiFi-4
"
```

RADIUS returns the redirect policy and portal URL, and the UE is then in Portal state. Next, the user logs in and a CoA is sent by the portal.

```
159 2014/09/08 11:30:45.07 EDT MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Change of Authorization(43) id 220 len 91 from 10.43.186.1:53449 vrid 1
    USER NAME [1] 17 68:7f:74:8b:3d:d7
    VSA [26] 6 Alcatel(6527)
    WLAN UE CREATION TYPE [184] 4 1
    VSA [26] 8 Alcatel(6527)
    SUBSC PROF STR [12] 6 SUBP_4
    VSA [26] 8 Alcatel(6527)
    SLA PROF STR [13] 6 SLAP_4
    VSA [26] 6 Alcatel(6527)
    PRIMARY DNS [9] 4 10.43.186.43
"
```

Finally the UE is in DSM state with unrestricted access.

Conclusion

The 7750 SR WLAN-GW, with Open SSID, can support WiFi Offload users as DSM subscribers instantiated on MS-ISA cards. This allows the support of a greater number of UEs on a single system when a full ESM feature set is not required. DSM UEs, just as ESM UEs, can have their traffic intercepted using LI.

Customer document and product support



Customer documentation

[Customer documentation welcome page](#)



Technical support

[Product support portal](#)



Documentation feedback

[Customer documentation feedback](#)