



7450 Ethernet Service Switch
7750 Service Router
7950 Extensible Routing System
Releases Up To 23.7.R2

Advanced Configuration Guide - Part III

3HE 14992 AAJ TQZZA
Edition: 01
September 2023

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

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

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

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

No part of this document may be modified.

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

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

© 2023 Nokia.

Table of contents

| | |
|---|------------|
| List of tables | 5 |
| List of figures | 8 |
| | |
| Preface | 20 |
| About This Guide..... | 20 |
| | |
| Multi-Service Integrated Service Adapter and Extended Services Appliance | 21 |
| Application Assurance — Application Identification and User-Defined Applications..... | 22 |
| Application Assurance — App-Profile, ASO and Control Policies..... | 44 |
| Application Assurance — Asymmetry Removal..... | 65 |
| Application Assurance — Best Practices for ISA and Host IOM Overload Protection..... | 77 |
| Application Assurance — DNS IP Cache..... | 94 |
| Application Assurance — GTP Roaming Firewall..... | 103 |
| Application Assurance — HTTP and HTTPS Redirect..... | 133 |
| Application Assurance — HTTP In Browser Notification..... | 155 |
| Application Assurance — Local URL List Filtering..... | 164 |
| Application Assurance — Security Gateway Stateful Firewall..... | 175 |
| Application Assurance — Stateful Firewall..... | 200 |
| Application Assurance — Usage Monitoring and Policy Control via Diameter Gx Protocol..... | 217 |
| Deterministic Large Scale NAT44..... | 241 |
| IP/GRE Termination..... | 274 |
| L2-aware NAT (with dNAT and MNPs)..... | 300 |
| L2TP Network Server..... | 333 |
| Multi-Chassis IPSec Redundancy..... | 362 |
| NAT Stateless Dual-Homing..... | 396 |
| | |
| Triple Play Service Delivery Architecture | 420 |
| ARP Hosts..... | 422 |
| Bridged CO..... | 442 |
| DHCP Server Failover States..... | 473 |
| DHCPv4 Server Basics..... | 497 |
| Diameter Application NASREQ..... | 528 |
| Diameter Base Protocol: Establishing a Diameter Peer Connection..... | 540 |

| | |
|---|------|
| ESM Basics..... | 552 |
| ESM 128-bit Mode for DHCPv6 IA_NA WAN Hosts..... | 582 |
| ESM IPv4: Multicast in a Wholesale/Retail Scenario..... | 587 |
| ESM IPv4: Multicast with Redirection..... | 600 |
| ESM IPv4: Multicast with SRRP..... | 631 |
| ESM SLAAC Prefix Assignment via Local Address Server..... | 659 |
| ESMv4: PPPoE Hosts..... | 676 |
| ESMv6: IPoE Dual Stack Hosts..... | 717 |
| ESMv6: PPPoE Dual Stack Hosts..... | 750 |
| Flexible Authentication Model in ESM..... | 777 |
| GTP Access..... | 820 |
| High Scale QoS IOM in ESM Context: Expanded SLA Mode..... | 844 |
| High Scale QoS IOM in ESM Context: Single SLA Mode..... | 878 |
| Ingress Multicast Path Management..... | 911 |
| IPoE Sessions..... | 944 |
| IPv4 DHCP Hosts..... | 969 |
| L2TP for Subscriber Access — LAC..... | 1004 |
| Local User Database Basics..... | 1052 |
| Local User Database for DHCPv4 Server..... | 1080 |
| Local User Database for Enhanced Subscriber Management..... | 1097 |
| Managed SAPs with Routed CO..... | 1125 |
| Multi-Chassis Ring Layer 2 with Enhanced Subscriber Management..... | 1147 |
| Python Cache Support for ESM Applications..... | 1171 |
| RADIUS-Triggered Dynamic Data Service Provisioning..... | 1187 |
| Raw Formatting of DHCPv4/v6 Options in ESM..... | 1222 |
| Routed CO..... | 1249 |
| Subscriber Redundancy for Routed CO..... | 1284 |
| Virtual Residential Gateway Authentication Scenarios..... | 1313 |
| Virtual Residential Gateway Home LAN Extension..... | 1337 |
| Virtual Residential Gateway Home Pool Management..... | 1357 |
| WiFi Aggregation and Offload — Basic Open SSID..... | 1372 |
| WiFi Aggregation and Offload — Basic Secure SSID with Distributed RADIUS Proxy..... | 1391 |
| WiFi Aggregation and Offload — IPv4/v6 Dual-Stack UEs..... | 1411 |
| WiFi Aggregation and Offload — Migrant User Support..... | 1425 |
| WiFi Aggregation and Offload — Open SSID with DSM and Lawful Intercept..... | 1435 |

List of tables

| | |
|--|-----|
| Table 1: Customer reserved App-filter ranges..... | 26 |
| Table 2: Classification rules for the ISP ON-NET content services..... | 29 |
| Table 3: Default QoS Policy, Application QoS Policy Table..... | 52 |
| Table 4: AA asymmetry removal topology..... | 66 |
| Table 5: Tracking ISA load in the reporting interval..... | 86 |
| Table 6: Denied GTP message types for roaming interface..... | 106 |
| Table 7: Allowed GTP message types (Cat-1)..... | 107 |
| Table 8: AA redirect errors..... | 144 |
| Table 9: SCTP PPIDs..... | 187 |
| Table 10: GTP Messages..... | 194 |
| Table 11: Traffic Streams..... | 310 |
| Table 12: ARP Host Time-Related Parameters..... | 432 |
| Table 13: Correlation of Hosts and BSA/BSR Services..... | 446 |
| Table 14: BSA/BSR Configuration for Host-1 Operation..... | 450 |
| Table 15: BSA/BSR Configuration for Host-2 Operation..... | 451 |
| Table 16: BSA/BSR Configuration for Host-3 Operation..... | 452 |
| Table 17: 64-bit WAN Mode Versus 128-bit WAN Mode..... | 583 |
| Table 18: Reserved PPPoE Tags..... | 678 |
| Table 19: LCP and IPCP Code..... | 681 |
| Table 20: Applicable Subscriber-Prefix Parameters..... | 727 |
| Table 21: Timer Parameters..... | 728 |

| | |
|---|------|
| Table 22: Router Advertisements Parameters..... | 730 |
| Table 23: RADIUS AVPs..... | 732 |
| Table 24: Local User Database Parameters..... | 734 |
| Table 25: DHCP Lease State Information..... | 734 |
| Table 26: Subscriber Prefix Parameters..... | 755 |
| Table 27: Subscriber Prefix Subnetting for SLAAC..... | 756 |
| Table 28: Subscriber-Prefix Parameters..... | 759 |
| Table 29: Prefix Subnetting for delegated-prefix-length /56..... | 759 |
| Table 30: RADIUS AVPs..... | 760 |
| Table 31: SLAAC-Related Parameters..... | 763 |
| Table 32: IPv6CP Nack Message Format..... | 775 |
| Table 33: Input/Output Rates..... | 872 |
| Table 34: Input and output rates throughout the subscriber QoS hierarchy..... | 903 |
| Table 35: Supported DHCP Option 82 Sub-Options..... | 975 |
| Table 36: Information in DHCP Lease State..... | 980 |
| Table 37: L2TPv2 header fields and descriptions..... | 1007 |
| Table 38: AVP header fields and descriptions..... | 1008 |
| Table 39: Generic L2TP RADIUS attributes..... | 1014 |
| Table 40: Nokia Specific L2TP RADIUS attributes..... | 1014 |
| Table 41: Masking examples..... | 1059 |
| Table 42: Dynamic Service Attribute List for Setup, Modify and Teardown..... | 1194 |
| Table 43: Dynamic Service Actions on Control- and Data-Channel..... | 1195 |
| Table 44: Function and Dictionary Relationship..... | 1198 |

| | |
|--|------|
| Table 45: RADIUS inserted raw options..... | 1223 |
| Table 46: Python Modified DHCP Fields..... | 1224 |
| Table 47: CLI Inserted DHCP Options..... | 1224 |
| Table 48: DHCP options inserted via RADIUS..... | 1235 |
| Table 49: ARP/MAT/BD-MAC-Prefix Possible Combinations..... | 1350 |

List of figures

| | |
|--|-----|
| Figure 1: App-Filters/Applications/AppGroup..... | 23 |
| Figure 2: HTTP persistent connection..... | 27 |
| Figure 3: Wireshark® www.wikipedia.org..... | 29 |
| Figure 4: Wireshark® HTTPS www.whatsapp.com..... | 31 |
| Figure 5: HTTPS SNI..... | 32 |
| Figure 6: SIP Wireshark® capture..... | 33 |
| Figure 7: H323 Wireshark® capture..... | 34 |
| Figure 8: Wireshark® GoGlobal..... | 39 |
| Figure 9: Service Tier Example using ASO, App-Profile and AQP..... | 48 |
| Figure 10: App-Profile, ASO, AQP Workflow Summary..... | 51 |
| Figure 11: Default Downstream Bandwidth Policing..... | 53 |
| Figure 12: AA asymmetry removal topology..... | 66 |
| Figure 13: Network to subscriber traffic flow..... | 75 |
| Figure 14: Subscriber to network traffic flow..... | 76 |
| Figure 15: System packet datapath to AA ISA..... | 77 |
| Figure 16: Basic Message Flow When Accessing Website..... | 95 |
| Figure 17: AA GTP roaming FW deployment..... | 104 |
| Figure 18: CLI command..... | 106 |
| Figure 19: Configuration topology..... | 111 |
| Figure 20: HTTP redirect..... | 135 |
| Figure 21: Example setup..... | 136 |

| | |
|--|-----|
| Figure 22: Connected device without full Internet access..... | 143 |
| Figure 23: HTTP redirect using session filter captive redirect..... | 145 |
| Figure 24: HTTP redirection..... | 147 |
| Figure 25: Access to domain included in the allow-list versus HTTPS redirect..... | 148 |
| Figure 26: Subscriber receives warning..... | 149 |
| Figure 27: HTTP notification – setup..... | 156 |
| Figure 28: Notification Message Example – Quota 80%..... | 156 |
| Figure 29: Local URL-list filtering setup..... | 165 |
| Figure 30: LTE SeGW Firewall Deployment..... | 176 |
| Figure 31: SeGW in Small Cells Architecture..... | 176 |
| Figure 32: Configuration Topology..... | 178 |
| Figure 33: Block Unsolicited Traffic..... | 201 |
| Figure 34: SFW — Allow Gaming..... | 202 |
| Figure 35: ALG Support Example — FTP..... | 203 |
| Figure 36: Configuration Topology..... | 205 |
| Figure 37: Gx reference point..... | 218 |
| Figure 38: Convergence..... | 219 |
| Figure 39: Gx reference point..... | 220 |
| Figure 40: Diameter protocol stack..... | 220 |
| Figure 41: ADC rules and related Nokia-defined AVPs defined for use by AA..... | 221 |
| Figure 42: PCC rules and related Nokia-defined AVPs defined for use by AA..... | 222 |
| Figure 43: ADC rule example of AVPs to install the application profile "gold_level"..... | 222 |
| Figure 44: PCC rule example of AVPs to install the application profile "gold_level"..... | 223 |

| | |
|---|-----|
| Figure 45: Capture of the ADC rule assignment of the "gold_level" appProfile..... | 223 |
| Figure 46: Call flow diagram..... | 225 |
| Figure 47: Example configuration setup..... | 227 |
| Figure 48: PCRF AVPs override call flow diagram..... | 235 |
| Figure 49: RAR containing ASOs and AppProfile override AVPs example..... | 235 |
| Figure 50: RAR containing usage monitoring ADC rules example..... | 238 |
| Figure 51: Deterministic NAT Mapping..... | 242 |
| Figure 52: Deterministic NAT Algorithm..... | 242 |
| Figure 53: Deterministic Mapping: Inside -> Outside Routing Instances..... | 243 |
| Figure 54: Deterministic mapping: outside IP port-blocks/ranges..... | 244 |
| Figure 55: Example topology..... | 244 |
| Figure 56: Case 1..... | 248 |
| Figure 57: Case 1 results..... | 253 |
| Figure 58: Case 1 flows..... | 253 |
| Figure 59: Case 2..... | 257 |
| Figure 60: Case 2: Prefix 10.1.0.0/23 results..... | 263 |
| Figure 61: Case 2: Prefix 10.2.0.0/22 results..... | 263 |
| Figure 62: Case 3..... | 264 |
| Figure 63: Case 3 results..... | 269 |
| Figure 64: Inverse mapping approach..... | 269 |
| Figure 65: Sending flows: deterministic + non-deterministic NAT..... | 273 |
| Figure 66: GRE packet format..... | 274 |
| Figure 67: Implementation..... | 276 |

| | |
|---|-----|
| Figure 68: IP/GRE over IPSec tunnel..... | 276 |
| Figure 69: GRE for remote access to a VPRN service..... | 280 |
| Figure 70: GRE for remote access to a VPRN service..... | 283 |
| Figure 71: IP/GRE tunneling via static route..... | 283 |
| Figure 72: Example GRE over IPSec tunnel..... | 288 |
| Figure 73: NAT Binding..... | 301 |
| Figure 74: Test Setup..... | 303 |
| Figure 75: Logical Mapping of Subscribers to L2-aware Pool..... | 307 |
| Figure 76: Example Topology..... | 334 |
| Figure 77: Ingress/Egress QoS Processing..... | 355 |
| Figure 78: MC-IPSec architecture..... | 363 |
| Figure 79: Example topology..... | 364 |
| Figure 80: Example Topology..... | 397 |
| Figure 81: Redundancy Status..... | 410 |
| Figure 82: Post-Failover Redundancy State..... | 417 |
| Figure 83: Bridged CO and Routed CO Example..... | 423 |
| Figure 84: ARP Hosts in a Bridged CO Environment Example..... | 424 |
| Figure 85: ARP Hosts in a Routed CO Environment Example..... | 426 |
| Figure 86: ARP Host Session Timeout Example..... | 433 |
| Figure 87: Trap Generation Example..... | 435 |
| Figure 88: Throttling Toward RADIUS Example..... | 437 |
| Figure 89: ARP Host Mobility Example..... | 438 |
| Figure 90: Bridged CO Network Topology..... | 442 |

| | |
|--|-----|
| Figure 91: Key Concepts of Bridged CO Model..... | 443 |
| Figure 92: Flow Chart for Subscriber-Profile Identification Algorithm..... | 445 |
| Figure 93: Flowchart for SLA-Profile Identification Algorithm..... | 446 |
| Figure 94: Sample Topology..... | 448 |
| Figure 95: Functionality of Each Node..... | 448 |
| Figure 96: Host-1 Setup Process..... | 457 |
| Figure 97: Host-2 Setup Process..... | 460 |
| Figure 98: Host-3 Setup Process..... | 466 |
| Figure 99: General Redundancy Model..... | 474 |
| Figure 100: DHCP Relay Agent and Server Redundancy Model..... | 475 |
| Figure 101: Access-Driven and Local-Remote Model..... | 476 |
| Figure 102: Local-Remote Model – Active-Standby..... | 477 |
| Figure 103: Local-Remote Model – Subnet-Based Load Sharing..... | 478 |
| Figure 104: Local-Remote Model – Range-Based Load Sharing..... | 478 |
| Figure 105: Failover State Transition Diagram..... | 479 |
| Figure 106: VPRN-1 Service Configuration..... | 482 |
| Figure 107: Accessing a DHCP server..... | 498 |
| Figure 108: Addresses, Subnets, and Pools in a DHCPv4 Server..... | 502 |
| Figure 109: General Address Allocation for DHCP..... | 504 |
| Figure 110: Baseline Service Configuration..... | 506 |
| Figure 111: NASREQ trigger..... | 529 |
| Figure 112: Diameter protocol stack..... | 540 |
| Figure 113: Diameter network topology..... | 541 |

| | |
|---|-----|
| Figure 114: Bridged RGW Scenario..... | 554 |
| Figure 115: Routed RGW Scenario..... | 554 |
| Figure 116: SLA-Profile and Sub-Profile..... | 555 |
| Figure 117: Subscriber Host Identification and Instantiation Process..... | 558 |
| Figure 118: Direct Address Assignment using LUDB/RADIUS..... | 562 |
| Figure 119: Indirect Address Assignment using a DHCP Server..... | 563 |
| Figure 120: Indirect Address Assignment using LAA..... | 565 |
| Figure 121: Wholesale/Retail Model 1..... | 587 |
| Figure 122: Wholesale/Retail Model 2..... | 588 |
| Figure 123: Layer 3 Wholesale/Retail..... | 589 |
| Figure 124: L2TP Wholesale-Retail Multicast..... | 595 |
| Figure 125: Network Topology Overview..... | 601 |
| Figure 126: Single BNG Setup with Multicast Redirection..... | 602 |
| Figure 127: Network Topology with MC-LAG..... | 608 |
| Figure 128: IPoE Multicast Message Flow..... | 617 |
| Figure 129: PPPoE Multicast Flow..... | 621 |
| Figure 130: Network Topology Overview..... | 632 |
| Figure 131: Example Topology..... | 633 |
| Figure 132: IPoE Subscriber Multicast Flow..... | 642 |
| Figure 133: PPPoE Multicast Flow..... | 647 |
| Figure 134: TPSDA Network Topology..... | 659 |
| Figure 135: Routed CO Network Topology..... | 677 |
| Figure 136: Discovery Stage Messages..... | 680 |

| | |
|---|-----|
| Figure 137: LCP Phase Messages..... | 682 |
| Figure 138: CHAP Handshaking Overview Process..... | 683 |
| Figure 139: PAP Overview Process..... | 683 |
| Figure 140: IPCP Phase Messages..... | 684 |
| Figure 141: Keepalive Messages..... | 684 |
| Figure 142: Link Termination Phase..... | 685 |
| Figure 143: Authentication Flow Chart..... | 695 |
| Figure 144: Pado-Delay Scenario..... | 715 |
| Figure 145: Stateful IPoE Dual Stack Subscriber Hosts..... | 718 |
| Figure 146: Dual Stack IPoE Routed Gateway Service..... | 719 |
| Figure 147: DHCPv6 Lease Process (Part A)..... | 720 |
| Figure 148: DHCPv6 Lease Process (Part B)..... | 721 |
| Figure 149: Prefix Delegation..... | 722 |
| Figure 150: IPv6 Address/Prefix Timers..... | 728 |
| Figure 151: PPPoE Dual Stack Hosts..... | 750 |
| Figure 152: Dual Stack PPPoE Bridged Gateway Service Example..... | 751 |
| Figure 153: Dual Stack PPPoE Routed Gateway Service Example..... | 752 |
| Figure 154: Message Flow for a Dual Stack PPPoE Host..... | 754 |
| Figure 155: Dual Stack PPPoE for Routed Gateway..... | 757 |
| Figure 156: DHCPv6 Renewals..... | 765 |
| Figure 157: Topology..... | 778 |
| Figure 158: GTP Access to the BNG..... | 820 |
| Figure 159: EPS Bearer Across the Different Interfaces..... | 821 |

| | |
|--|-----|
| Figure 160: GTP-C and GTP-U Encapsulation..... | 821 |
| Figure 161: GTP-U in Up and Downstream..... | 822 |
| Figure 162: GTP-C Control - GTP Host Creation..... | 824 |
| Figure 163: GTP-C Control - GTP Host Deletion..... | 825 |
| Figure 164: GTP Access Topology..... | 825 |
| Figure 165: GTP Tunnel and Subscriber Termination Configuration Logic..... | 828 |
| Figure 166: Test Environment Example..... | 846 |
| Figure 167: QoS Hierarchy in Expanded SLA Mode..... | 847 |
| Figure 168: Managing Congestion on HSQ in Expanded SLA Mode..... | 871 |
| Figure 169: Test environment example..... | 880 |
| Figure 170: QoS hierarchy in single SLA mode..... | 881 |
| Figure 171: Managing Congestion on HSQ in Single SLA Mode..... | 902 |
| Figure 172: IOM/IMM Paths Connecting to Switch Fabric Planes..... | 912 |
| Figure 173: Dynamic Bandwidth Rate Management..... | 923 |
| Figure 174: Falling-Percent-Reset..... | 924 |
| Figure 175: Admin-Bw Rate Management..... | 925 |
| Figure 176: IPoE Session..... | 944 |
| Figure 177: IPoE Session Key..... | 946 |
| Figure 178: IPoE Session Creation Flow..... | 947 |
| Figure 179: IPoE Session Creation via AAA/RADIUS..... | 948 |
| Figure 180: Configuring IPoE session authentication..... | 953 |
| Figure 181: Baseline configuration..... | 954 |
| Figure 182: Bridged CO Network Topology..... | 970 |

| | |
|---|------|
| Figure 183: Routed CO Network Topology..... | 970 |
| Figure 184: DHCP Lease Process..... | 971 |
| Figure 185: Subscriber Host Connectivity Verification..... | 995 |
| Figure 186: DHCP Proxy Server: Lease Split Operation..... | 1000 |
| Figure 187: DHCP Proxy Server: Lease Split Operation, DHCP Client Disconnected..... | 1001 |
| Figure 188: DHCP Host Mobility..... | 1002 |
| Figure 189: PPP access architectures..... | 1005 |
| Figure 190: Supported L2TP reachability options..... | 1006 |
| Figure 191: RADIUS-triggered tunnel/session setup without LNS renegotiation..... | 1010 |
| Figure 192: RADIUS-triggered tunnel/session setup with LNS renegotiation..... | 1011 |
| Figure 193: Running multiple PPP sessions over a single L2TP tunnel..... | 1012 |
| Figure 194: PPP user-initiated release/terminate..... | 1012 |
| Figure 195: L2TP tunnel and session state diagram..... | 1013 |
| Figure 196: Base router hosted LAC with single endpoint/single tunnel..... | 1016 |
| Figure 197: Base router hosted LAC with multiple endpoints..... | 1017 |
| Figure 198: VRF hosted LAC..... | 1018 |
| Figure 199: RADIUS returns L2TP tunnel group..... | 1020 |
| Figure 200: LUDB returns L2TP tunnel group..... | 1021 |
| Figure 201: L2TP keepalive mechanism..... | 1042 |
| Figure 202: Floating peers accept..... | 1045 |
| Figure 203: Floating peers ignore..... | 1045 |
| Figure 204: Floating peers reject..... | 1046 |
| Figure 205: LUDB applications..... | 1053 |

| | |
|---|------|
| Figure 206: Processing an LUDB lookup request..... | 1054 |
| Figure 207: Creating LUDBs and LUDB entries..... | 1055 |
| Figure 208: Host matching examples..... | 1060 |
| Figure 209: Host matching examples (continued)..... | 1061 |
| Figure 210: LUDB access via a DHCPv4 server..... | 1081 |
| Figure 211: Example configuration..... | 1085 |
| Figure 212: Decoding the ESM user option..... | 1092 |
| Figure 213: LUDB authentication..... | 1098 |
| Figure 214: Direct and indirect LUDB authentication..... | 1099 |
| Figure 215: LUDB parameters for IPoE..... | 1101 |
| Figure 216: LUDB parameters for PPPoE..... | 1102 |
| Figure 217: LUDB authentication for regular SAPs..... | 1103 |
| Figure 218: LUDB authentication for capture and managed SAPs..... | 1104 |
| Figure 219: Baseline setup..... | 1106 |
| Figure 220: Network Topology..... | 1127 |
| Figure 221: MC-Ring Layer 2 CO Dual Homing..... | 1148 |
| Figure 222: Dual homing Under Steady-State Condition..... | 1149 |
| Figure 223: Broken Ring State..... | 1150 |
| Figure 224: Network Topology..... | 1151 |
| Figure 225: Unicast Services — Logical Topology..... | 1158 |
| Figure 226: Multicast Service — Logical Setup..... | 1166 |
| Figure 227: Test topology..... | 1172 |
| Figure 228: Principle Model of Dynamic Data Services..... | 1188 |

| | |
|---|------|
| Figure 229: Test Topology..... | 1190 |
| Figure 230: Building Blocks of Dynamic Data Services..... | 1190 |
| Figure 231: Hierarchy of Snippets..... | 1199 |
| Figure 232: DHCPv4 Lease-Time Inserted by RADIUS and DHCPv4 Server..... | 1226 |
| Figure 233: Python Injected Hint for Lease-Time..... | 1227 |
| Figure 234: Format of the IA-NA Option..... | 1228 |
| Figure 235: Format of the IA Address Option..... | 1228 |
| Figure 236: Topology..... | 1229 |
| Figure 237: Components of the Routed CO Model..... | 1250 |
| Figure 238: Numbered Scenario For IES 1..... | 1257 |
| Figure 239: Unnumbered Scenario for IES 1..... | 1267 |
| Figure 240: Hybrid Configuration..... | 1275 |
| Figure 241: Network Redundancy Components for ESM Routed CO..... | 1285 |
| Figure 242: BRG and home device management..... | 1314 |
| Figure 243: Explicit BRG authentication..... | 1315 |
| Figure 244: Implicit BRG authentication..... | 1316 |
| Figure 245: Example service configuration for explicit and implicit BRG authentication..... | 1318 |
| Figure 246: vRGW-HLE..... | 1338 |
| Figure 247: BD Connections in the Data Plane..... | 1338 |
| Figure 248: ARP Requests Flooded with AAR Disabled..... | 1347 |
| Figure 249: No ARP Request Flooding with AAR Enabled..... | 1348 |
| Figure 250: MAT - Access to Network Direction..... | 1349 |
| Figure 251: MAT - Network to Access Direction..... | 1349 |

| | |
|--|------|
| Figure 252: Virtual Residential Gateway in the Network with Bridged Residential Gateway at Home..... | 1358 |
| Figure 253: Services Configuration Overview..... | 1359 |
| Figure 254: Call Flow for Open SSID..... | 1374 |
| Figure 255: WiFi Offload Scenario with Open SSID and Local DHCP Server..... | 1375 |
| Figure 256: WiFi Offload Scenario with Secure SSID and L2-Aware NAT..... | 1393 |
| Figure 257: Call Flow for Secure SSID with DSM..... | 1403 |
| Figure 258: DHCPv4 + SLAAC/64 — Open SSID..... | 1412 |
| Figure 259: DHCPv4 + SLAAC/64 model — closed SSID..... | 1413 |
| Figure 260: DHCPv4 + SLAAC/64 with DHCPv4 linking model — closed SSID..... | 1414 |
| Figure 261: DHCPv4 + DHCPv6/128 IA_NA model — closed SSID..... | 1415 |
| Figure 262: DHCPv4 + SLAAC/64 with DHCPv4 linking model — DTA..... | 1416 |
| Figure 263: Sequence of events to establish and authenticate a migrant User..... | 1430 |
| Figure 264: Sequence of events to establish and authenticate a migrant user (continued)..... | 1431 |
| Figure 265: WiFi Offload Scenario with Open SSID, DSM and LI..... | 1436 |

Preface

About This Guide

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

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

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

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

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

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

Audience

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

Multi-Service Integrated Service Adapter and Extended Services Appliance

This section provides MS-ISA and ESA configuration information for the following topics:

- [Application Assurance — Application Identification and User-Defined Applications](#)
- [Application Assurance — App-Profile, ASO and Control Policies](#)
- [Application Assurance — Asymmetry Removal](#)
- [Application Assurance — Best Practices for ISA and Host IOM Overload Protection](#)
- [Application Assurance — DNS IP Cache](#)
- [Application Assurance — GTP Roaming Firewall](#)
- [Application Assurance — HTTP and HTTPS Redirect](#)
- [Application Assurance — HTTP In Browser Notification](#)
- [Application Assurance — Local URL List Filtering](#)
- [Application Assurance — Security Gateway Stateful Firewall](#)
- [Application Assurance — Stateful Firewall](#)
- [Application Assurance — Usage Monitoring and Policy Control via Diameter Gx Protocol](#)
- [Deterministic Large Scale NAT44](#)
- [IP/GRE Termination](#)
- [L2-aware NAT \(with dNAT and MNPs\)](#)
- [L2TP Network Server](#)
- [Multi-Chassis IPSec Redundancy](#)
- [NAT Stateless Dual-Homing](#)

Application Assurance — Application Identification and User-Defined Applications

This chapter describes Application Assurance (AA) Application Identification and User-Defined Applications configurations.

Topics in this chapter include:

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

Applicability

The information and configuration in this chapter were initially based on SR OS Release 11.0.R3, but updates were made based on SR OS Release 19.10.R2.

There are no specific prerequisites for this example.

Overview

This chapter is intended for Application Assurance (AA) network architects and engineers. It provides best practice information to customize the AA policy and classify any type traffic to meet the service provider reporting, charging or control requirements.

In addition to the signatures built and supported by Nokia, service providers can create their own application signatures based on various criteria. This customization capability can be used to classify traffic hosted on the provider network (web portal, streaming service) or hosted on the Internet and not yet covered by the default AA signature set.

Basics and terminology

The following main components are used for AA classification:

- **Application Filters** — App-filters are used to define applications based on Layer 3 to Layer 7 criteria. They provide a mapping between one or more protocol signatures or customized traffic patterns into an application of interest.
- **Application** — Such as BitTorrent®, Netflix®. Traffic is classified into applications using app-filters.
- **Application Group** — Such as peer-to-peer, multimedia streaming. For the purpose of reporting and control, applications of similar type/function can be grouped together in Application Groups (App-Group).
- **Charging Group** — Such as zero rating, default. For the purpose of charging or control, applications and app-group can be grouped together in charging groups.

The following table is a high-level example to illustrate how app-filters are used to defined applications and show their logical grouping into app-group and charging group.

Figure 1: App-Filters/Applications/AppGroup

Maximum Flexibility to Identify Standard and Custom Applications of Interest

| Criteria | App-Filter (ordered list of entries, ACL like) | Application | Application Group | Charging Group |
|---|---|-------------|-------------------|--------------------|
| - Protocol - Expression: (HTTP, SIP, H323, TLS, RTSP) - L4 Server Port - IP Server Address - Flow Direction - Custom Protocol | Expression - http: yahoo.com | Yahoo | Web | CG#1 - Default |
| | Expression - http: maps.google.com | Google Maps | | CG#2 - Zero Rating |
| | Expression - http: facebook.com | Facebook | Social Networking | CG#1 - Default |
| | Protocol: ftp_control, ftp_data | FTP | File Transfer | |
| | Protocol: bittorrent, dht, utp | BitTorrent | Peer to Peer | |
| | Protocol: emule | Emule | | |

| | | |
|--|---|--|
| Flexible classification/identification rules (apps-filters) to identify: - Standard applications - Custom defined applications | Flexible applications/app-group creation and mapping for: - Reporting - Control (redirect, enrichment, policing...) | Independent charging group mapping for differentiated billing. |
|--|---|--|

al_0680

- BitTorrent® and Emule® applications are defined using their protocol signature and grouped in the P2P app-group.
- FTP application is defined using both ftp_data and ftp_control protocol signatures, the app is mapped in the file transfer app-group.
- Google Maps® and Yahoo® web sites are defined using http expression and grouped together in the Web app-group.

Configuration

Classification criteria (App-filter)

The operator can take full advantage of the flexible AA policy configuration to classify traffic from any application of interest using various criteria ranging from Layer 3 to Layer 7 expressions.

Expression match criteria allows to further refine traffic classification by identifying traffic from HTTP, HTTPS (SSL/TLS), SIP, H323, RTSP, Citrix protocol signatures.

The different app-filter match criteria are listed below:

- L7 expression
 - HTTP: host, URI, user agent, referer
 - SSL/TLS: certificate org name, common name, SNI
 - H323: product ID
 - SIP: URI, user agent, media type
 - RTSP: host, URI, user agent

- Citrix: application published name
- RTMP: page-host, page-uri, swf-host, swf-uri
- IP protocol number
- IP server address
- TCP/UDP server port
- Custom protocol
- Protocol signature

The following operators are supported to define expression-based app-filters:

- ^ : Expression starts with
- \$: Expression ends with
- * : Wildcard - anything before or after
- \l : Forces case sensitivity
- \d : Any single decimal digit [0-9]
- \. : Any single character
- * : Asterisk character

Examples of expression match combinations:

- ^abcd* : match 'abcd' at beginning, can end with anything
- *abcd* : match 'abcd' anywhere
- *abcd\$: match 'abcd' at the end
- ^abcd\$: exact expression match 'abcd'
- ^ab*cd\$: string starts with 'ab', ends with 'cd' (anything else in between)
- ^ab\dcd\$: string starts with 'ab', followed by a decimal digit, ends with 'cd'



Note:

It is possible to combine different criteria or expressions within the same filter in which case an implicit AND operation between the criteria within the same filter is done by the system.

Application definition example

The following example provides a basic configuration example with the application FTP made of two protocol signatures ftp_control and ftp_data; the application is mapped into the application group file transfer:

Create the application group.

```
configure application-assurance group 1:1 policy
  app-group "File Transfer"
  exit
```


Create the application.

```
configure application-assurance group 1:1 policy
  application "FTP"
  app-group "File Transfer"
  exit
```

Create the app-filters.

```
configure application-assurance group 1:1 policy
  app-filter
    entry <1..65535> create
      protocol eq "ftp_data"
      application "FTP"
      no shutdown
    exit
    entry <1..65535> create
      protocol eq "ftp_control"
      application "FTP"
      no shutdown
  exit
```



Note:

Once the application is created, the operator is expected to configure the collection of statistics at the subscriber level for this new application (usually only for business VPNs).

User-defined applications

General recommendations

In order to classify traffic properly, it is recommended to follow the guidelines and best practices defined in this section before creating a new application:

- Analyze the application traffic
 - Identify what type traffic is used (Wireshark®).
 - Use the application the same way the end user would use it, the same application can create various flows.
- Configure the appropriate App-Filters
 - Following the analysis of the application done above, create the application.
 - Follow the App-filter best practices chapter to understand in which range to add the filters.
 - More than one App-filter can be required to identify a single application.

AppDB/default AA policy

The default AA policy called AppDB (Application Database) is provided by Nokia and should be used on most deployments. Contact your regional support organization for more details on how to obtain it.

This configuration includes applications and application-groups most providers can use by default and is designed to allow the addition of any custom entries required by service providers to identify additional services/applications.

The following customization options exist:

- Generating a configuration for a particular region (such as APAC)
- Generating a new configuration (or updating a configuration) containing specific applications

Before adding new entries to the template and customizing the configuration, it is recommended to follow the next guidelines on app-filters and ranges. These guidelines are key to allow an easy upgrade path from the policy configuration provided by Nokia.

App-filters

App-filters are an ordered list of entries. It is important to keep the order of this list consistent with the classification objective.

For instance, a common configuration mistake is to configure a filter rule for the HTTP protocol signature before HTTP expression filters. If that were the case, then app-filters using HTTP expressions would not be used as the system would find an acceptable match with the protocol signature before walking the list of expressions configured. This mistake is described in the following example:

```
entry 100 create
  description "Default HTTP Protocol"
  protocol eq "http"
  application "HTTP"
  no shutdown
exit
entry 110 create
  description "Google"
  expression 1 http-host eq ".*.google.com$"
  application "Google"
  no shutdown
exit
```

This is an incorrect App-filter order. App-filter entry #100 will always match before the http expression entry #110.



Note:

It is not necessary to specify a protocol when defining an expression filter, the protocol is implicit based on the type of expression match criteria used (for instance, HTTP, SIP, H323).

App-filter ranges

The App-filter list is an ordered list, it is key to configure each app-filter in the right order and in the proper range.

The operator can customize the policy and create applications and app-filters by using the following ranges shown in [Table 1: Customer reserved App-filter ranges](#) (other ranges are used by the default policy):

Table 1: Customer reserved App-filter ranges

| Range Name | Description | Start | End |
|--------------------|---|-------|------|
| Extended top range | Top range, matches before any other filters | 1 | 1499 |
| High priority | Top range for high priority matches | 2000 | 4999 |

| Range Name | Description | Start | End |
|------------------------|---|-------|-------|
| Expression range A | HTTP Host, Host+URI; optionally with IP/Port match | 19000 | 22999 |
| Expression range B | Other Expression Match; optionally with IP/Port match | 33000 | 34999 |
| Extended protocols | Protocol-signature + Port IP Dir. match | 40000 | 58999 |
| Custom protocols | Custom protocol signature match | 61000 | 61999 |
| Trusted/validate ports | 1st packet validate, 1st packet trusted match | 62000 | 63999 |

Ordering basics:

- Layer 7 expression-based filters are located before their parent protocol signature (for example, expression matches on http are located before the HTTP protocol app-filter; the same applies to TLS, SIP, H323, RTSP, Citrix).
- HTTP host and URI are located before the HTTP referer for accounting accuracy (for example, YouTube® from within Facebook® is classified as YouTube®)
- App-filters combining protocol signatures with Layer 4 port, IP protocol, IP address, or flow direction are always located before the protocol signature only filter range.

HTTP

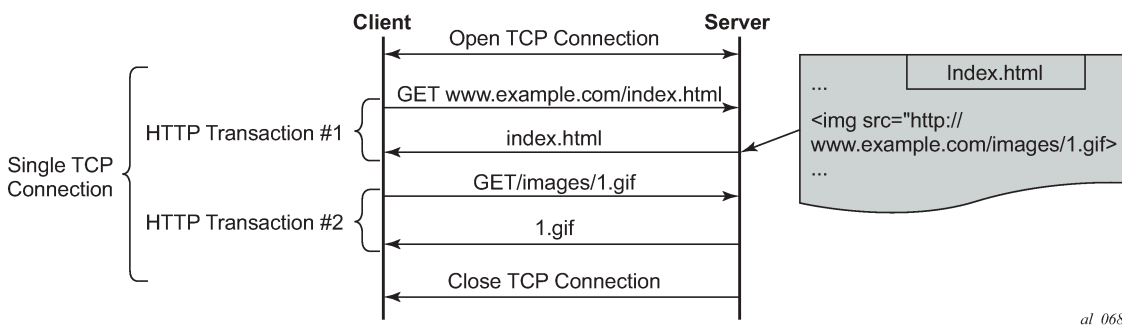
Protocol

HTTP is a client/server protocol using TCP/IP at the transport layer to deliver resources such as HTML files, images, videos and more.

HTTP 1.1 enables HTTP clients to use a persistent connection to a server allowing them to reuse the same TCP session for multiple HTTP transactions. Text, images, video, scripts and other objects can be downloaded individually in different transactions through the same TCP session.

Figure 2: HTTP persistent connection describes a typical persistent HTTP connection between a web client and a server with multiple HTTP transactions within the same TCP session:


Figure 2: HTTP persistent connection



User-defined expression-based HTTP applications will use the first HTTP transaction to classify the flow (optionally this behavior can be modified).

HTTP request

The following example shows the content of a typical HTTP request to wikipedia.org which includes the following header fields: HTTP Host, HTTP URI, HTTP User Agent and HTTP referer fields:



25452

- HTTP Host — Represents the domain name (does not include "http://").
- HTTP URI — The URL trailer after the host domain name (begins with slash "/").
- HTTP Referer — The address of the previous web page from which a link to the currently requested page was followed (in this example, the referer is `www.google.com` which means the user clicked on a link from a Google search pointing to `wikipedia.org`).
- HTTP User Agent — This identifies the web browser or application making the HTTP request.

Configuration examples

HTTP host (Wikipedia)

Classifying HTTP traffic from this web site can be done using a single expression tail anchored on the HTTP host:

```
configure application-assurance group 1:1 policy app-filter
  entry <1..65535> create
    description "Wikipedia Web Access" expression 1 http-host eq "*.wikipedia.org$"
    application "Wikipedia"
    no shutdown
  exit
```

This can be confirmed using Wireshark®.

Figure 3: Wireshark® www.wikipedia.org

| No. | Time | Source | Destination | Protocol | Info |
|-----|----------|----------------|----------------|----------|---------------------------------------|
| 149 | 4.474276 | 192.168.1.4 | 208.80.154.225 | TCP | 57881 > http [SYN] Seq=O Win=8192 Le |
| 172 | 4.508432 | 208.80.154.225 | 192.168.1.4 | TCP | http > 57881 [SYN, ACK] Seq=0 Ack=1 |
| 173 | 4.508543 | 192.168.1.4 | 208.80.154.225 | TCP | 57881 > http [ACK] Seq=1 Ack=1 Win=62 |
| 204 | 4.568615 | 192.168.1.4 | 208.80.154.225 | HTTP | GET / HTTP/1.1 |
| 207 | 4.615704 | 208.80.154.225 | 192.168.1.4 | TCP | http > 57881 [ACK] Seq=1 Ack=986 Win |
| 208 | 4.615807 | 208.80.154.225 | 192.168.1.4 | TCP | [TCP segment of a reassembled PDU] |
| 209 | 4.615635 | 208.80.154.225 | 192.168.1.4 | HTTP | HTTP/1.0 301 Moved Permanently |
| 210 | 4.617685 | 192.168.1.4 | 208.80.154.225 | TCP | 57881 > http [ACK] Seq=956 Ack=614 w |

| |
|--|
| <ul style="list-style-type: none"> ▪ Frame 204: 1039 bytes on wire (8312 bits), 1039 bytes captured (8312 bits) ▪ Ethernet II, Src: HonHaiPr_77:bf:c8 (4c:0f:6e:77:bf:c8), Dst: Netgear_d8:68:78 (c0:3f:0e:d8:68:78) ▪ Internet Protocol, Src: 192.168.1.4 (192.168.1.4), Dst: 208.80.154.225 (208.80.154.225) ▪ Transmission Control Protocol, Src Port: 57881 (57881), Dst Port: http (80), Seq: 1, Ack: 1, Len: 985 ▪ Hypertext Transfer Protocol <ul style="list-style-type: none"> ▪ GET / HTTP/1.1\r\n <ul style="list-style-type: none"> Host: en.wikipedia.org\r\n Connection: keep-alive\r\n User-Agent: Mozilla/5.0 (Windows NT 6,1; WOW64) AppleWebKit/537.4 (KHTML, like Gecko) Chrome/22.0.1229. |
|--|

al_0682W

Classification per URI within the same host

Operators may need to apply different charging rules to different content located on the same HTTP domain (different URI, same HOST).

Table 2: Classification rules for the ISP ON-NET content services displays an example of classification rules for the ISP ON-NET content services:

Table 2: Classification rules for the ISP ON-NET content services

| URL | Charging rule | AA application |
|--------------------------|--------------------|--------------------|
| www.ispdomain.com/video | Rule #1 – 0 Rating | ISP-Portal-Video |
| www.ispdomain.com/images | Rule #2 – Charge X | ISP-Portal-Images |
| www.ispdomain.com/* | Rule #3 – Charge Y | ISP-Portal-Default |

HTTP 1.1 can reuse the same TCP connection for many transactions to the same server. Classifying each HTTP transaction to www.ispdomain.com independently requires a specific AA configuration.

SR OS allows to selectively enable “http-match-all-requests” in app-filters to improve the system performance and limit the HTTP analysis per domain.

The following configuration example allows traffic classification of different URIs of the same domain (www.ispdomain.com) independently, therefore allowing differentiated charging and control:

- http-match-all-req is enabled on all host+uri app-filters to www.ispdomain.com
- default app-filter required to match any traffic to www.ispdomain.com

```
configure application-assurance group 1:1 policy
  app-filter
    entry <1..65535> create
```

```
description "Zero rated content"
expression 1 http-host eq "^www.ispdomain.com$"
expression 2 http-uri eq "^/video*"
http-match-all-req
application "ISP Portal Video"
no shutdown
exit
entry <1..65535> create
description "Image charging"
expression 1 http-host eq "^www.ispdomain.com$"
expression 2 http-uri eq "^/images*"
http-match-all-req
application "ISP Portal Images"
no shutdown
exit
entry <1..65535> create
description "Default charging"
expression 1 http-host eq "^www.ispdomain.com$"
http-match-all-req
application "ISP Portal Default"
no shutdown
exit
```

SSL/TLS (HTTPs)

Protocol

HTTPS uses SSL/TLS to encrypt traffic between the client and the server. Because this communication is encrypted, it is not possible to identify the HTTP Host or URI. However, AA can still identify the service requested by the subscriber by looking at the TLS certificate information or Server Name Indication exchanged in the clear before the TLS session is established.



Note:

SSL/TLS expression-based app-filters are not limited to HTTPS. HTTPS is not a protocol in itself, but it is HTTP traffic, tunneled encrypted into SSL/TLS on port 443.

SSL/TLS certificates

The following snapshot ([Figure 4: Wireshark@ HTTPS www.whatsapp.com](#)) from Wireshark shows the SSL/TLS certificate exchanged using the mobile application **whatsapp**®.

Figure 4: Wireshark® HTTPS www.whatsapp.com

| No. | Time | Source | Destination | Protocol | Info |
|-----|-----------|---------------|---------------|----------|--|
| 42 | 44.854067 | 192.11.231.83 | 50.23.142.168 | TCP | 33084 > https [SYN] Seq=0 Win=64240 L |
| 43 | 44.933347 | 50.23.142.168 | 192.11.231.83 | TCP | https > 33084 [SYN, ACK] Seq=0 Ack=1 |
| 44 | 45.213335 | 192.11.231.83 | 50.23.142.168 | TCP | 33084 > https [ACK] Seq=1 Ack=1 Win=12 |
| 45 | 45.342530 | 192.11.231.83 | 50.23.142.168 | SSLv3 | Client Hello |
| 46 | 45.448230 | 50.23.142.168 | 192.11.231.83 | TCP | https > 33084 [ACK] Seq=1 Ack=75 Win=6 |
| 47 | 45.851643 | 50.23.142.168 | 192.11.231.83 | SSLv3 | Server Hello |
| 48 | 45.853122 | 50.23.142.168 | 192.11.231.83 | TCP | [TCP segment of a reassembled PDU] |
| 49 | 45.853231 | 50.23.142.168 | 192.11.231.83 | TCP | [TCP segment of a reassembled PDU] |
| 50 | 46.042243 | 192.11.231.83 | 50.23.142.168 | TCP | 33084 > https [ACK] Seq=75 Ack=2777 w |
| 51 | 46.245518 | 192.11.231.83 | 50.23.142.168 | TCP | 33084 > https [ACK] Seq=75 Ack=4097 w |
| 52 | 46.334985 | 50.23.142.168 | 192.11.231.83 | SSLv3 | Certificate, Server Hello Done |

[Reassembled TCP Segments (4686 bytes): #47(1309), #48(1388), #49(1320), #52(669)]
 Secure Socket Layer
 SSLv3 Record Layer: Handshake Protocol: Certificate
 Content Type: Handshake (22)
 Version: SSL 3.0 (0x0300)
 Length: 4672
 Handshake Protocol: Certificate
 Handshake Type: Certificate (11)
 Length: 4668
 Certificates Length: 4665
 Certificates (4665 bytes)
 Certificate Length: 1377
 Certificate (id-at-commonname-*.whatsapp.net)-at-organizationalUnitName-Domain Control validated, id
 Certificate Length: 1250

The certificate information can be found in the Server Hello message sent by the server, capturing SSL/TLS (HTTPS) traffic from this application can be done using a single app-filter entry tail anchored on the TLS Common Name Certificate:

```
configure application-assurance group 1:1 policy
  app-filter
    entry <1..65535> create
      description "Whats App tls and image/voice/video traffic"
      expression 1 tls-cert-subj-common-name eq
        "*.whatsapp.net$"
      application "Whats App"
      no shutdown
    exit
```

Server name indication

SSL/TLS traffic can optionally be identified using the Server Name Indication (SNI) which is an extension to the TLS protocol.

The SNI is found in the TLS Client Hello, the http-host expression in the app-filter is reused to classify this traffic:

Figure 5: HTTPS SNI

| No. | Time | Source | Destination | Protocol | Info |
|-----|----------|---------------|---------------|----------|-------------------------------------|
| 4 | 0.088936 | 192.11.231.82 | 98.138.6.52 | TCP | iclpv-nlc > https [SYN] Seq=0 Win=1 |
| 5 | 0.165069 | 98.138.6.52 | 192.11.231.82 | TCP | https > iclpv-nlc [SYN, ACK] Seq=0 |
| 6 | 0.165136 | 192.11.231.82 | 98.138.6.52 | TCP | iclpv-nlc > https ACK] Seq=1 Ack=1 |
| 8 | 0.383867 | 192.11.231.82 | 98.138.6.52 | TLSv1 | Client Hello |

- Cipher Suites (36 suites)
- Compression Methods Length: 1
- Compression Methods (1 method)
- Extensions Length: 56
- Extension: server_name
Type: server_name (0x0000)

Data (30 bytes)

- Extension: elliptic_curves
- Extension: ec_point_formats
- Extension: SessionTicket TLS

```

0000 00 1e e5 7a 96 5f 00 0c 29 7e 53 cc 08 00 45 00 ...z... )~s...E.
0010 00 da 80 dS 40 00 80 06 69 2c c0 0b e7 52 62 8a ....O... i...Rb.
0020 06 34 05 72 01 bb 1f 6f 07 aS 3e de f1 43 50 18 .4.r...o ...>..CP.
0030 fc 00 6e 15 00 00 16 03 01 00 ad 01 00 00 a9 03 ...n.....
0040 01 4d 80 1d b4 c7 oc 86 06 8d 17 70 14 6c 85 ed .M..... ..p.1..
0050 ff a3 30 5c 56 87 c3 09 98 d3 e0 b3 9e a1 45 04 ..Olv... ..E.
0060 S1 00 00 48 00 ff c0 0a c0 14 00 88 00 87 00 38 Q..H... ..8
0070 c0 0f c0 05 00 84 00 35 00 39 c0 07 c0 09 c0 11 .....5 .9.....
0080 c0 13 00 45 00 44 00 33 00 32 c0 0c c0 0e c0 02 ...E.D.3 .2.....
0090 c0 04 00 96 00 41 00 04 00 05 00 2f c0 08 c0 12 ....A. ....
00a0 00 16 00 13 c0 0d c0 03 fe ff 00 0a 01 00 00 38 .....8
00b0 00 00 00 1e 00 1c 00 00 19 75 73 2e 64 61 74 61 .....us.data
00c0 2e 74 6f 6f 6c 62 61 72 2e 79 61 68 6f 6f 2e 63 toolbar .yahoo.c
00d0 6f 6d 00 0a 00 08 00 06 00 17 00 18 00 19 00 0b om.....
00e0 00 02 01 00 00 23 00 00 .....#...
    
```

```

configure application-assurance group 1:1 policy
  app-filter
    entry <1..65535> create
      description "Yahoo HTTP or TLS SNI"
      expression 1 http-host eq "*.yahoo.com$"
      application "Yahoo"
      no shutdown
    exit
    
```

SIP

Protocol

SIP is a signaling protocol used for controlling multimedia communication sessions such as voice and video over RTP. AA automatically monitors SIP control flows and associates RTP/RTCP media flows accordingly in the sip_rtp protocol signature.

The operator can use a SIP expression match criteria in app-filter to further refine traffic classification and identify any additional application on top of the default AA policy. This can be particularly useful in business VPNs to identify voice and telepresence applications.

AA supports SIP expression match criteria on SIP URI, SIP user agent and SIP media type. The following snapshot from Wireshark® shows a SIP control exchange using the voice-video application Vonage®

followed by the RTP media audio flow; the expression fields that can be matched using AA app-filters are highlighted:

Figure 6: SIP Wireshark® capture

```

Session Initiation Protocol
  Request-Line: INVITE sip:3102951568@k.voncp.com;transport=UDP SIP/2.0
  Method: INVITE
  Request-URI: sip:3102951568@k.voncp.com;transport=UDP
    Request-URI User Part: 3102951568
    Request-URI User Part: k.voncp.com
  [Resent Packet: False]
  Message Header
    From: "613-963-0148" <sip:16139630148@k.voncp.com>;tag=1019fb60-7196c445-2710-4e9485ff-7b9cb12-4e9485ff
    To: <sip:3102951568@k.voncp.com>
    Call-ID: 101a7de0-7196c445-2710-4e9485ff-229a8c45-4e9485ff@k.voncp.com
    CSeq: 1 INVITE
    Via: SIP/2.0/UDP 69.196.150.113:10000;branch=z9hG4bK-4e9485ff-f42b6c64-49ad5933
    P-Preferred-Identity: off
    Max-Forwards: 70
    Supported: replaces.timer.100rel
    User-Agent: VTA001346FE8BF111.4.1-r060815-1.00.09-20070402170142 1248967645135/1007551373 308
    Contact: <sip:16139630148@69.196.150.113:10000;transport=UDP>
    Min-SE: 0
    Content-Type: application/sdp
    Content-Length: 294
  Message Body
    Session Description Protocol
      Session Description Protocol version (v): 0
      Owner/Creator, Session Id (o): a0000 8644 6672 IN IP4 69.196.150.113
      Session Name (s): SIP Cal
      Connection Information (c): IN IP4 69.196.150.113
      Time Description, active time (t): 0 0
      Media Description, name and address (m): audio 10050 RTP/AVP 0 101 8 2 18
        Media Type: audio
        Media Port: 10050
  
```

al_0685

Configuration example

The following configuration example provides the configuration to classify Vonage® SIP/RTP desktop traffic using SIP URI expression:

```

configure application-assurance group 1:1 policy
  app-filter
    entry <1..65535> create
      description "Vonage"
      expression 1 sip-uri eq "*voncp.com*"
      application "Vonage"
      no shutdown
    exit
  
```

H323

Protocol

Similar to SIP, H323 is a signaling protocol used for controlling multimedia communication sessions such as voice and video over RTP. AA automatically monitors H323 control flows and associates the RTP media flow accordingly in the h323_rtp protocol signature.

The operator can use an H323 expression match criteria app-filter to further refine traffic classification and identify any additional application on top of the default AA policy. This can be particularly useful in business VPNs to identify voice and telepresence applications.

AA supports H323 expression match criteria on the H323 product ID. The following snapshot from Wireshark shows an H323 control exchange using the Telepresence application LifeSize® followed by the RTP media audio flow; the expression field that can be matched using AA app-filters is highlighted:

Figure 7: H323 Wireshark® capture

```

Transmission Control Protocol, Src Port: 61505 (61505), Dest Port: h323hostcall (1720), Seq: 1, Ack: 1, Len: 212
  TPKT, Version: 3, Length: 212
    Q.931
      Protocol discriminator: Q.931
      Call reference value length: 2
      Call reference flag: Message sent from originating side
      Call reference value: 461a
      Message type: SETUP (0x05)
      Bearer capability
      Display 'Conference Room'
      User-user
    H.225.0 CS
      H323-UserInformation
        h323-uu-pdu
          h323-message-body: setup (0)
            setup
              protocolIdentifier: 0.0.8.2250.0.5 (Version 5)
              sourceAddress: 3 items
              sourceInfo
                vendor
                  vendor
                    H.221 Manufacturer: Unknown (0xb500a11a)
                    productid: LifeSize Express 220
                    versionId: 4.7.10.14
              0 . . . . . mc: False
              .0 . . . . . undefinedNode: False
  
```

Configuration example

The following configuration example provides the configuration to classify LifeSize® H323/RTP traffic using the H323 product ID expression:

```

configure application-assurance group 1:1 policy
  app-filter
    entry <1.65535> create
      description "LifeSize H323 traffic"
      expression 1 h323-product-id eq "^LifeSize*"
      application "LifeSize"
      no shutdown
    exit
  
```

RTSP

Protocol

RTSP is a signaling protocol used for controlling media streaming content such as audio and video over RTP/RDT. AA automatically monitors the RTSP control flows and associates its RTP/RDT media flow with the `rtsp_rtsp` protocol signature.

The operator can use an RTSP expression match criteria app-filter to further refine traffic classification and identify any additional application on top of the default AA policy. This can be particularly useful to identify specific streaming applications.

AA supports RTSP expression match criteria on the RTSP Host, URI, User Agent. The following snapshot from Wireshark® shows an RTSP setup request to YouTube® followed by the RTP media audio flow; the expression fields that can be matched in RTSP SETUP request using AA app- filters are highlighted:



```
RTSP Header
SETUP rtsp://v3.cache7.c.youtube.com/ZTww=/0/0/0/video.3gp/trackID=13 RTSP/1.0
CSeq: 3
User-Agent: Mozilla/5.0 (BlackBerry; U; BlackBerry 9800; en) AppleWebKit/54.8+
x-wap-profile: "http://www.blackberry.net/go/mobile/profiles/uaprof/9800_unknown/6.0.0.rdf"
Transport: RTP/AVP;unicast;client_port=51132-51133;mode="PLAY"
```

25453

Configuration example

The following configuration example provides the configuration to classify YouTube® RTSP/RTP traffic using RTSP Host expression:

```
configure application-assurance group 1:1 policy
  app-filter
    entry <1..65535> create
      description "YouTube RTSP/RTP Video"
      expression 1 rtsp-host eq "*.youtube.com$"
      application "YouTube"
      no shutdown
    exit
  exit
```

Citrix

Protocol

Independent Computing Architecture (ICA) is a Citrix Systems® protocol used in Citrix's WinFrame, Citrix XenApp (formerly called MetaFrame/Presentation Server), and Citrix XenDesktop products.

Citrix makes it possible to run applications remotely on large servers, thus making better use of server resources while at the same time allowing people using other platforms to use the applications, for example, run Microsoft® Word on a UNIX workstation.

Citrix_ica protocol signature will detect any remote application using Citrix (the protocol needs to be unencrypted and configured to non-seamless). The Citrix ICA session is started from a client and can be anything from Remote Desktop, SAP to Microsoft® Word.

The Citrix expression match app-filter is used to classify traffic based on the Citrix-published application. This published application is configured on the server and in the preceding example, it can be for instance RDP, SAP, Word, XLS or Microsoft® Word depending how the server is configured.

Configuration example

```
configure application-assurance group 1:1 policy
  app-filter
    entry <1..65535> create
      description "Citrix SAP Application"
      expression 1 citrix-app eq "SAP"
      application "Citrix SAP"
      no shutdown
    exit
```

IP address and TCP/UDP port

Traffic from specific servers can be classified using IPv4/v6 server-address app-filter rules. It is used usually to identify traffic from an internal (on-net) server as opposed to an Internet (off-net) server.

The server-address app-filter automatically detects the client from the server by identifying which side opens the connection. It implicitly classifies traffic based on the server IP address or port number. For example, if A initiates a TCP connection to B, then flows from A to B and from B to A can be classified with a match on server-address B. Similarly, a flow initiated from B to A can be classified using a match on server-address A.

Server address

The following configuration example uses a server-address app-filter to classify traffic from server 10.1.1.1 in the application called Application-1:

```
configure application-assurance group 1:1 policy
  app-filter
    entry <1..65535> create
      description "Server #1 10.0.0.1"
      server-address eq 10.0.0.1/32
      application "Application-1"
      no shutdown
    exit
```

Server address and server port

The following configuration example uses server-address and server-port app-filters to classify traffic from server 10.0.0.2 on port 1234 in the application called Application-2. It is particularly useful when the same server is used to provide different services that need to be classified separately:

```
configure application-assurance group 1:1 policy
  entry <1..65535> create
    description "Server #2 10.0.0.2 port 1234 Only"
    server-address eq 10.0.0.2/32
    server-port eq 1234
    application "Application-2"
    no shutdown
  exit
```

Server port and protocol signature

It is possible to combine a protocol signature with a port number in the same app-filter, this is typically done in business VPNs for specific internal applications not detected using existing AA protocol signatures.

The following configuration example classifies a business VPN application running on TCP port 4000 and not detected by any other signatures. It combines the protocol signature unknown_tcp with the desired port number. This allows keeping the classification untouched for the rest of the protocols/applications and is the recommended approach:

```
configure application-assurance group 1:1 policy
  app-filter
    entry <1..65535> create
      description "Business VPN Application X Port 4000"
      server-port eq 4000
      protocol eq unknown_tcp
      application "Business VPN Application X"
      no shutdown
    exit
```

**Note:**

It is important to follow the app-filter range recommendations for a proper classification of traffic using IP address or port number.

Flow setup direction

Traffic can be classified based on flow-setup-direction app-filter. The flow setup direction can be either subscriber-to-network or network-to-subscriber.

Network side and subscriber side is AA terminology related to where AA is enabled:

- In broadband and mobile networks, AA is enabled per subscriber. This means the subscriber side represents the ESM/mobile/transit subscriber while the network side represents Internet or other subscribers.
- In business VPNs, AA is enabled on a VPN SAP/spoke SDP and the subscriber side represents the local VPN site (SAP/spoke/transit).

The following example shows the configuration to classify http traffic hosted by AA subscribers (for example, broadband subscribers running a web server):

```
configure application-assurance group 1:1 policy
  app-filter
    entry <1..65535> create
      description "HTTP Server on the subscriber side"
      flow-setup-direction network-to-subscriber
      protocol eq http
      application "HTTP Server"
      no shutdown
    exit
```

IP protocol

Traffic can be classified using an IP protocol number for non TCP/UDP traffic.

The following example provides the configuration to classify ICMP IPv4/v6 traffic:


```
configure application-assurance group 1:1 policy
  app-filter
    entry <1..65535> create
      description "ICMP v4"
      protocol eq "non_tcp_udp"
      ip-protocol-num eq icmp
      application "ICMP"
      no shutdown
    exit
    entry <1..65535> create
      description " ICMP v6"
      protocol eq "non_tcp_udp"
      ip-protocol-num eq ipv6-icmp
      application "ICMP"
      no shutdown
    exit
```

Custom protocol

Custom protocols can be used to classify TCP/UDP applications using hexadecimal string matching (up to 16 hex octets) at a configurable payload offset in the data payload. The expression string length and offset must not exceed 128 bytes.

To illustrate this feature the Solaris® application GoGlobal is used. It provides remote access to a server (similar to VNC®). The following snapshot ([Figure 6: SIP Wireshark® capture](#)) from Wireshark® shows a TCP SYN/ACK session establishment followed by the first data exchange:

Figure 8: Wireshark® GoGlobal



Wireshark® shows that each TCP session payload starts with 80DC0400 (no offset) after the three-way TCP handshake. As a result, the configuration required to classify this traffic is as follows:

```
configure application-assurance group 1:1 policy
  custom-protocol 1 ip-protocol-num tcp create
  description "goglobal tcp"
  expression 1 eq "\x80\xdc\x04\x00" offset 0 direction client-to-server
  no shutdown
exit
app-filter
  entry <1..65535> create
  description "GoGlobal "
  protocol eq "custom_01"
  application "GoGlobal"
  no shutdown
exit
```

Typical configuration mistakes

An operator creating new user-defined applications can make a few typical mistakes which are listed below:

- App-filters in shutdown state — The default app-filter state is shutdown. A **no shutdown** command must be executed in order for it to be enabled.
- App-filters with no match criteria — This is a more troublesome mistake as it will catch all the traffic entering the filter in a particular application.

Troubleshooting application identification

Show commands

Router/partition statistics

Partition level statistics are not updated in real time. Instead, statistics for a particular flow are updated either at flow closure or every five minutes. The five-minute sliding window interval is a common interval for all flows in an ISA MDA. Different ISA MDAs will have a different five-minute windows as this interval is set at the MDA boot time.

The following command can be used to view the statistics for all applications configured in the ISA Group 1, Partition 1:

```
show application-assurance group 1:1 application count
```

Alternatively, it is possible to sort the display by octets, packets, flows:

```
show application-assurance group 1:1 application count top [octets | packets | flows] [max-count <max-count>]
```

The operator can also identify which app-filters are being hit by the AA policy per partition (this command is not available per subscriber), it is particularly useful to identify which filters are used and optionally prune unnecessary app-filters from user-defined applications:

```
show application-assurance group 1:1 policy app-filter
```

**Note:**

The app-filter policy is usually relatively large, in which case additional 7x50 SR CLI functionality can be used to filter out the output and only show the relevant information.

The following example was created for the application FTP:

```
A:PE# show application-assurance group 1:1 policy app-filter | match "application \"FTP\""
pre-lines 3 post-lines 2
    exit
    entry 44300 create (2 flows, 1205 B)
        protocol eq "ftp_control"
        application "FTP"
        no shutdown
    exit
    entry 44301 create (2 flows, 1401 B)
        protocol eq "ftp_data"
        application "FTP"
        no shutdown
    exit
```

Because partition level statistics are not updated in real time, it is recommended for troubleshooting purposes to use subscriber statistics or sub-study statistics.

Subscriber statistics

Subscriber-level statistics can be updated in real time. AA is usually configured by the operator to collect subscriber-level statistics for all application groups in residential and Wifi, while business VPNs typically collect Application group and all applications for each site with AA enabled.

The following commands can be used to view per subscriber statistics for all app-groups or applications configured in ISA group 1, partition 1 for the ESM subscriber "Bob" or business VPN SAP 1/1/1:10:

```
show application-assurance group 1:1 aa-sub esm "bob" app-group count
show application-assurance group 1:1 aa-sub sap 1/1/1:10 application count
```

In case only app-group statistics are collected per subscriber, the aa-sub-study feature can be used to collect per application-level statistics for selected subscribers, as follows:

```
A:PE# configure application-assurance group 1:1 statistics aa-sub-study application
A:PE>config>app-assure>group>statistics>aa-sub-study# aa-sub esm "bob"
```

Once done, the system will show all application level statistics for this subscriber:

```
show application-assurance group 1:1 aa-sub-study esm "bob" application count
```

Similar to partition-level statistics, aa-sub and aa-sub-study statistics can be sorted by octets, packets, flows:

```
show application-assurance group 1:1 aa-sub-study esm "bob" application count top [octets |
packets | flows] [max-count <max-count>]
```

**Note:**

When the number of flows per ISA card reaches a threshold then per subscriber statistics are not available in real time anymore and only the snapshot command can be used to display the statistics recorded in the previous five-minute interval window:

```
show application-assurance group 1:1 aa-sub-study esm "bob" snapshot application
count
```

AppFilterMiss

The default policy configuration provides a failsafe application at the very end of the app-filter list to classify any remaining traffic in the AppFilterMiss application. There should never be any traffic in this application. This failsafe filter is used as a debug to make sure that there are no major issues in the configuration.

**Note:**

Traffic can typically be classified as AppFilterMiss when not all protocol signatures are mapped to a particular application. This could happen when upgrading to a new ISA software and enabling new protocol signature detection while not ensuring first that the correct application was provisioned. See the Release Note upgrade section for more details on AA signature upgrade.

Tools

Flow-record-search

Traditional show commands may not provide enough information when troubleshooting flow identification and the operator can use the ISA flow-record-search tool to dump the ISA flow table for more information. This feature comes with a large number of filtering options documented in the user guide.

Each flow gives visibility into: Flow ID, Sub-Type, Sub-Name, Initiator, Direction, Source IP, Dest. IP, IP Protocol, Source Port, Dest. Port, FC, DSCP, Classified, Protocol, Application, App- Group, Charging Group, Packets tx, Bytes Tx, Packets-discarded, Bytes-discarded etc.

See below for the most commonly used commands.

The following command shows all the flows in an ISA card per ISA group:partition (can be a very long output, up to 3M entries):

```
application-assurance group 1:1 flow-record-search isa 1/2
```

The following command shows all the flows per AA subscriber in a given group:partition:

```
tools dump application-assurance group 1:1 flow-record-search aa-sub esm "bob"
```

The following command shows all the active flows per AA subscriber in a given group:partition:

```
tools dump application-assurance group 1:1 flow-record-search aa-sub esm "bob" flow- status active
```

The flow-record-search command is also available with additional details by adding search-type detail at the end of the command line. Note that due to the length of the output it is recommended to paste the CLI output content in a notepad file.

HTTP host recorder

AA cflowd allows operators to export the HTTP domain extracted from HTTP flows to the NSP cflowd collector. This allows the operator to understand which HTTP hosts are visible in the network.

However, in case a cflowd collector is not deployed, AA provides the HTTP host recorder tool command to record HTTP hosts seen by AA. See the *Multiservice ISA and ESA Guide* for more details.

```
A:PE# show debug
debug
  application-assurance
    group 1:1
      http-host-recorder
        filter
          default-filter-action record
          record http-host-app-filter-candidates
        exit
        rate 100
        no shutdown
      exit
    exit
  exit
exit
A:PE# tools dump application-assurance group 1:1 http-host-recorder top bytes
```

Port recorder

This function is particularly useful in business VPN (it can also be used in residential networks). The port-recorder AA tool function is similar to the http-recorder. It allows the operator to record which ports are used on selected applications.

It is most commonly used with the applications Unidentified TCP and Unidentified UDP but it can be configured to record any other applications:

```
A:PE# show debug
debug
  application-assurance
    group 1:1
      port-recorder
        application "Unidentified TCP"
        application "Unidentified UDP"
        rate 100
        shutdown
      exit
    exit
  exit
exit
exit
exit

A:PE# tools dump application-assurance group 1:1 port-recorder top bytes
```

Conclusion

This example, which is intended for Application Assurance network architects and engineers, provides the information required to modify an existing AA policy following AA best practices and guidelines, and provides the necessary troubleshooting information to better understand application classification using Application Assurance.

Application Assurance — App-Profile, ASO and Control Policies

This chapter provides information about Application Assurance (AA) app-profile, Application Service Options (ASOs) and control policy configurations.

Topics in this chapter include:

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

Applicability

This example is applicable to all 7750, 7450 and 7750-SRc chassis supporting Application Assurance and was tested on SR OS Release 12.0.R4.

It is recommended to use the AppDB prior to configuring traffic control policies. The AppDB is a default configuration file to define all of the applications of interest, including all of the relevant application-groups, applications and app-filters to classify traffic, and can be obtained through Nokia's support organization.

Overview

In addition to providing valuable traffic analysis and statistics information using the 7750 Service Router (SR) or 7450 Ethernet Service Switch (ESS) and Application Assurance (AA), one of the key objectives of the AA solution is to provide the tools to manage subscriber traffic at the application level. Examples of traffic management actions include:

- Throttling low priority bandwidth hungry applications during peak hours.
- Prioritizing and remarking selected applications.
- Implementing a walled-garden environment providing open access to selected free web services only, redirecting all other requests from unregistered subscribers to a registration portal with payment services.
- Enrich HTTP Header with subscriber identification parameters to offer subscribers transparent access to premium content.
- In browser notification which triggers the display of administrative, informational or promotional messages in selected browser-sessions.
- Stateful session filtering with Application Level Gateway (ALG) support to protect subscribers against unsolicited flows.
- Parental control services interworking with an external Internet Content Adaptation Protocol (ICAP) server for rating the requested web sites.

Application traffic control policies can be applied as global policies for all subscribers, or they can be activated for individual subscribers or groups of subscribers.

This example describes the basics of activating Application Assurance on a given subscriber through the use of App-Profile and demonstrates the use of static or dynamic traffic control policies using Application

Service Options (ASOs) and Application QoS Policies (AQP). It also provides detailed information for configuring Bandwidth, Flow-Count and Flow-Rate Policing including Time of Day (ToD) policing. Other policy control actions can be found in the Advanced Configuration Guide or in the MS-ISA User Guide.

Configuration

Activation of AA Services

App-Profile

Application profiles (app-profile) enable application assurance services for a given Enhanced Subscriber Management (ESM), Distributed Subscriber Management (DSM), or transit subscriber, or for a SAP or spoke SDP which are commonly referred to as **AA-subscribers (aa-sub)**. Each app-profile is unique in the system and defines the services that the AA subscriber will receive.

Assigning an app-profile to an ESM subscriber affects every host of that subscriber. Similarly, applying an app-profile to a SAP/spoke SDP will affect all traffic within that SAP/spoke SDP.

App-profiles are defined at the AA group partition level (in case of a partitioned ISA-AA group), see the configuration example below:

```
A:BNG# configure
  application-assurance group 1:1 policy
    app-profile "1-1/15M" create
      description "App-Profile Description"
      divert
      characteristic "Parental Control" value "enabled"
      capacity-cost 15
    exit
```

The app-profile parameters are:

- **divert** — Diverts all traffic from and to this subscriber to an ISA-AA. Configuring **no divert** effectively disables all AA services for subscribers using this app-profile.
Default value: **no divert**.
- **characteristic** [*<characteristic-name>* **value** *<value-name>*] — one or more optional ASO service characteristics can be used to apply an AA control policy to the subscriber.
- **capacity-cost** *<cost>* — An application profile capacity cost is used to load balance AA subscribers across multiple ISA-AA cards. A common practice is to define a cost proportional to the expected peak BW for the subscribers using this profile (in Kbps or Mbps). The capacity cost is out of the scope of this example. The range is 1 to 65535, default 1.

This app-profile example uses the following naming convention:

<group-id>-<partition-id>/<BW>M where

- *<group-id>* — The ISA-AA group ID on which this profile is created.
- *<partition-id>* — The AA partition ID on which this profile is created.
- *<BW>* — Defines the maximum bandwidth used by the subscriber, which is used for aa-subscriber cost load balancing and subscriber rate limiting. The *M* stands for Mbps.

In general the operator can choose to use either ASO characteristics override or multiple app-profiles to apply different AA QoS policies to ESM Subscribers or Business VPN sites. For flexibility and scale it is recommended to use ASO overrides whenever possible. This is described in more details below.

**Note:**

Prior to using special characters in a policy object name the operator should verify the list of special characters supported by the 5620 SAM; for instance the 5620 SAM does not support the use of ":" in the app-profile name therefore it should be avoided.

Residential and Wi-Fi Services

The app-profile can be assigned or modified for ESM, DSM or Transit IP subscribers either at subscriber creation time or while the subscriber is in service:

- Subscriber creation — An app-profile can be assigned at subscriber creation time through RADIUS, DHCP Option 82, Local User Database, static configuration or through a default app-profile.
- In service app-profile modification — An app-profile can be dynamically modified in service through a RADIUS Change of Authorization (CoA). From software Release 12.0.R1 an app-profile can also be dynamically modified in service through Gx.

In case no app-profile is returned at subscriber creation by RADIUS, LUDB or DHCP, or when no static configuration is present, the system can apply a default app-profile if configured within the subscriber group-interface (or MSAP policy) sub-sla-mgmt:

```
sub-sla-mgmt
  def-app-profile "1-1/15M"
exit
```

Business VPN and other Service Interfaces

App-profiles are statically assigned to a given SAP, spoke SDP or transit prefix VPN site via the 5620 SAM or CLI.

The following configuration shows how to enable application assurance on a SAP or spoke SDP in a business VPRN service:

```
A:PE>config>service# vprn 100 customer 1 create
description "L3 Service Customer 1"
interface "to-site1" create
  address 192.168.1.1/24
  sap 1/1/10:11 create
  app-profile "1-1/15M"
  exit
interface "to-site2" create
  address 192.168.2.1/24
  spoke-sdp 12:100 create
  app-profile "1-1/15M"
  exit
no shutdown
```

Defining Application Service Options

ASOs for Traffic Control - Introduction

To determine which application control policies need to be applied to a AA-subscriber, an app-profile with a number of service characteristics (ASOs) is associated with each subscriber. These service characteristics are then used as match criteria in AQP policy rules to determine which rules to apply.

Therefore ASOs are service characteristics assigned to a subscriber and are used to identify the traffic control policy rule (AQP) applicable to a subscriber or a group of subscribers.

Most policy rules will be applicable to multiple subscriber profiles; nevertheless it is possible that a specific subscriber requires a dedicated policy.

ASO Characteristics and Values

For each service option that can be used by one or more subscribers, an ASO characteristic should be defined with a number of values that represent all available choices for that service characteristic. The names and values of the ASO characteristics are configurable string values; best practice is to use strings that provide a meaningful description of the service characteristic they represent.

Each ASO characteristic requires a default value and each app-profile inherits the default value of all the ASO characteristics created in a given partition unless a characteristic is referenced directly in the app-profile or overwritten as described below.

ASOs are defined at the AA group partition level (in case of a partitioned ISA-AA group). In the configuration example below two different ASO characteristics are defined: "Parental Control" and "P2P-Sub-DL":

```
BNG>config>app-assure# group 1:1 policy
app-service-options
  characteristic "Parental Control" create
    value "disabled"
    value "enabled"
    default-value "disabled"
  exit
  characteristic "P2P-Sub-DL" create
    value "500k"
    value "1M"
    value "unlimited"
    default-value "unlimited"
  exit
```

The ASO values and default value of a characteristic can be displayed using a show command:

```
A:BNG# show application-assurance group 1:1 policy app-service-option "P2P-Sub-DL"
=====
Application-Assurance Application Service Options
=====
Characteristic "P2P-Sub-DL"
Value                               Default
-----
1M                                   No
500k                                 No
unlimited                             Yes
=====
```

When configuring service characteristics for optional service options, it is recommended to configure a default value which will not trigger any AQP policy action (the default value does not match any AQP match criteria) such that the behavior of existing subscribers and app-profiles will not change until the operator specifically configures or signals a non-default characteristic value for the subscriber or the app-profile. In the example above "Parental Control" "disabled" and "P2P-Sub-DL" "unlimited" would have no corresponding AQP by design; therefore if these particular service options were applied to a subscriber they would not match a QoS policy entry.

How to Specify Service Options for AA Subscribers

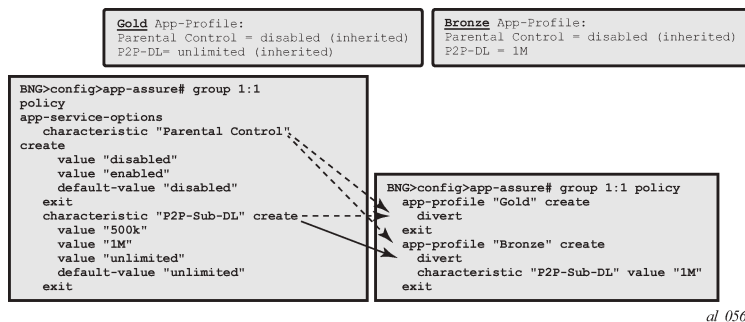
ASO Assignment in App-Profile

ASOs can be statically assigned in the app-profile; this type of ASO characteristic assignment is typically reserved to the default service options enabled on a large number of subscribers.

Figure 9: Service Tier Example using ASO, App-Profile and AQP shows an example of AA service definition (ASO and app-profile) for a Gold and Bronze service tier definition with the following characteristics:

- Two app-profiles *Gold* and *Bronze*
- *Gold* app-profile — No specific policy actions or ASO characteristics are configured statically in the app-profile.
- *Bronze* app-profile — A specific ASO characteristic and value is assigned to the profile to limit Peer to Peer download traffic to 1Mbps (this example does not show the app-qos-policy nor policer configuration, this will be described later).

Figure 9: Service Tier Example using ASO, App-Profile and AQP



al_0569

Each app-profile inherits the default values of all the ASO characteristics defined in a AA group-partition; in the example above this is reason why the app-profile Gold inherits "Parental Control" "disabled" and "P2P-Sub-DL" "unlimited". The app-profile Bronze inherits "Parental Control" "disabled" while "P2P-Sub-DL" "1M" is assigned to this profile statically.

The operator can identify per app-profile which characteristics values are inherited from their default value and which are statically assigned using the following show command:

```

*A:BNG# show application-assurance group 1:1 policy app-profile "Gold"
app-profile "Gold" create
divert
characteristic "P2P-Sub-DL" inherits default-value "unlimited"
    
```



```

        characteristic "Parental Control" inherits default-value "disabled"
    exit

A:BNG# show application-assurance group 1:1 policy app-profile "Bronze"
    app-profile "Bronze" create
        divert
        characteristic "P2P-Sub-DL" value "1M"
        characteristic "Parental Control" inherits default-value "disabled"
    exit
    
```



Note:

Using ASO overrides, described later, it is possible to implement the same choice of AA service options using a single app-profile.

ASO Overrides per Subscriber via RADIUS or Gx

Prior to SR OS 12.0.R1 the operator can assign (and modify: CoA) the app-profile per ESM or Transit-IP subscribers using the "Alc-App-Prof-Str" [26-6527-45] RADIUS attribute.

SR OS 12.0.R1 added support for ASO characteristic overrides for ESM and Transit-IP subscribers via RADIUS using the attribute "Alc-AA-App-Service-Options" [26-6527-193]. This attribute can be returned during the subscriber creation process or while the subscriber is in service through RADIUS CoA. Refer to SR OS 12.0 RADIUS Attributes Reference Guide for more details related to the use of the AA RADIUS attributes.

An example of a RADIUS CoA message returned to the system to modify both the app-profile and one ASO characteristic is provided below:

```

NAS-Port-Id = "1/1/5:4088"
Framed-IP-Address = 192.168.211.30
Alc-App-Prof-Str = "1-1/15M"
Alc-AA-App-Service-Options = "P2P-Sub-DL=1M"
    
```

The ASO characteristics and values assigned to a given subscriber (statically via app-profile or overridden) can be displayed using the following show command:

```

A:BNG# show application-assurance group 1:1 aa-sub esm "sub1" summary
=====
Application-Assurance Subscriber Summary (realtime)
=====
AA-Subscriber           : sub1 (esm)
ISA assigned            : 1/2
App-Profile             : 1-1/15M
App-Profile divert     : Yes
Capacity cost          : 1
Aarp Instance Id       : N/A
HTTP URL Parameters    : (Not Specified)
Last HTTP Notified Time : 2014/08/07 12:07:47
-----
Traffic                 Octets           Packets          Flows
-----
...
...
-----
Application Service Options (ASO)
-----
Characteristic          Value            Derived from
    
```

```
-----
P2P-Sub-DL          1M          dyn-override
Parental Control    disabled     default
=====
```

In the show command output above, the **derived from** field describes how the characteristics and values are assigned to the subscriber:

- app-profile — The characteristic's value statically configured in the app-profile.
- dyn-override — The characteristic's value received from RADIUS or Gx.
- default — The characteristic's default value inherited (not statically configured in the app-profile nor dynamically modified).

SR OS 12.0.R1 also introduced support for signaling the app-profile or ASO characteristics override via Gx, see [Application Assurance — App-Profile, ASO and Control Policies](#) for more details.

ASO Overrides for Business VPN and Other Services

Since SR OS 9.0.R1, ASO characteristic override values can be statically assigned to business VPN SAP, spoke SDP and transit prefix subscribers.

The operator can provision the AA policy override parameters, multiple characteristics overrides per AA-sub can be defined per override policy, see the configuration example below:

```
A:BNG>config>app-assure# group 1:1 policy-override
  policy aa-sub sap 1/1/5:210 create
    characteristic "P2P-Sub-DL" value "1M"
    characteristic "Parental Control" value "enabled"
  exit
```

Application Control Policies

App-QoS-Policy (AQP)

App-Profile / ASO / AQP Workflow Summary

App-profiles enable application assurance services for a given AA-subscriber. Each app-profile is unique in the system and defines the service that the AA subscriber will receive.

To determine which control policies need to be applied to an AA-subscriber, a number of service characteristics (ASO) are associated with each AA-subscriber.

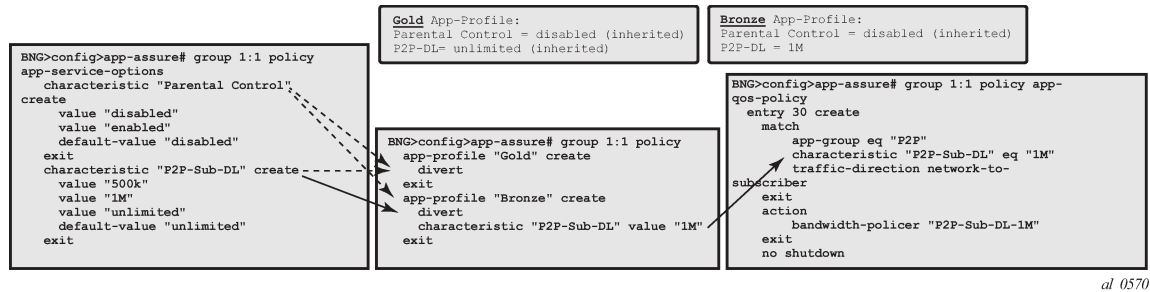
As described earlier, these service characteristics can either be configured directly within the app-profile or assigned using overrides and they are then used as match criteria in AQP policy rules to determine which application policy rules to apply.

The app-qos-policy (AQP) is an ordered list of entries defining policy actions for flows diverted to Application Assurance. Each AQP entry is composed of match criteria and action(s).

Flows are evaluated against all entries of the AA QoS policy defined in the AA group partition that the subscriber app-profile belongs to (in case of a partitioned AA group).

Figure 10: App-Profile, ASO, AQP Workflow Summary provides a configuration example summary with app-profile, ASO, AQP and policers:

Figure 10: App-Profile, ASO, AQP Workflow Summary



Match and Action Criteria

AQP Match Criteria

Multiple match criteria can be specified per AQP entry in which case the action will only apply to flows that match all criteria. The most common match criteria are: characteristic, application, app-group and charging-group.

The following AA match criteria can be used in an AQP:

- **app-group {eq | neq} <app-group name>**
- **application {eq | neq} <app name>**
- **charging-group {eq | neq} <charging-group-name>**
- **traffic-direction {subscriber-to-network | network-to-subscriber | both}**
- **characteristic <characteristic-name> eq <value-name>**: up to 4 characteristics and values per AQP
- **ip-protocol-num {eq | neq} <protocol-id>**
- **src-ip {eq | neq} <ip-address> or ip-prefix-list <ip-prefix-list-name>**
- **dst-ip {eq | neq} <ip-address> or ip-prefix-list <ip-prefix-list-name>**
- **src-port {eq | neq} <port-num> or range <start-port-num> <end-port-num>**
- **dst-port {eq | neq} <port-num> or range <start-port-num> <end-port-num>**
- **dscp {eq | neq} <dscp-name>**
- **aa-sub <aa-sub-name>**

AQP Actions

The following AA traffic control policies can be specified in an AQP:

- **drop**
- **bandwidth-policer <policer-name>**
- **flow-count-limit <policer-name>**
- **flow-rate-limit <policer-name>**

- **remark dscp in-profile** <dscp-name> **out-profile** <dscp-name>
- **remark fc** <fc-name>
- **remark priority** <priority-level>
- **http-error-redirect** <redirect-name>
- **http-redirect** <redirect-name> **flow-type** <flow-type> — Redirect traffic to a landing page
- **mirror-source [all-inclusive]** <mirror-service-id>
- **session-filter** <session-filter-name> — Session filter firewall
- **url-filter** <url-filter-name>: category based URL Filtering using ICAP
- **http-notification** <http-notification-name>
- Additional drop actions:
 - **error-drop**: configure a drop action for packets cut-through due to IP packet errors (bad IP checksums, tcp/udp port 0, etc.)
 - **overload-drop**: configure a drop action for packets cut-through due to overload
 - **fragment-drop**: configure a drop action for IP fragmented packets

Default Versus Application-Specific AQP Policies

Application QoS Policy

It usually requires the examination of a few packets to identify the protocol/application of a flow. When AQP entries are defined to match on IP header criteria (IP address, IP prefix list, TCP/UDP Port Number, IP Protocol, DSCP) or application criteria (application, App-Group or charging group), the AQP action will only be applied to matching application flows after a flow has been classified as a given application.

Default QoS Policy

If the AQP entry does not include match criteria against application (application, app-group and charging-group) or IP header information (IP address, IP prefix list, TCP/UDP port number, IP protocol, DSCP) then the AQP policy will be applied to all matching flows starting with the first packet of a flow before protocol and application identification is complete. Such AQPs are called default subscriber policies.

For an AQP to be qualified as a default subscriber policy, the match criteria must be limited to any combination of ASO characteristic values, traffic direction and optional AA subscriber name.

AQP match and actions for the default QoS policy and application QoS policy are summarized in [Table 3: Default QoS Policy, Application QoS Policy Table](#) :

Table 3: Default QoS Policy, Application QoS Policy Table

| Policy | AQP Match | AQP Action |
|-------------|--|--|
| Default QoS | ASO characteristic/values traffic direction | Remark FC, DSCP, Priority Bandwidth, flow-count, flow-rate policing Session-filter |

| Policy | AQP Match | AQP Action |
|-----------------|--|--|
| | aa-sub | Url-filter Mirror Error-drop, overload-drop, fragment-drop Drop |
| Application QoS | ASO characteristic/values traffic direction aa-sub application app-group charging-group IP address, IP Prefix List TCP/UDP Port Number DSCP IP Protocol Number | Remark FC, DSCP, Priority Bandwidth, flow-count, flow-rate policing HTTP Notification HTTP Redirect HTTP Enrichment Mirror Drop |

To ensure fair access to the ISA-AA bandwidth and flow resources, it is recommended to configure default AQP policy entries limiting bandwidth and flow resources per AA sub.

Figure 11: Default Downstream Bandwidth Policing shows a default subscriber policy limiting the downstream bandwidth (network-to-subscriber direction) to 25Mbps per subscriber:

Figure 11: Default Downstream Bandwidth Policing

```

7750>config>app-assure# group 1:1 policy
  app-service-options
    characteristic "access-rate" create
      value "100M"
      value "25M"
      default-value "100M"
    exit
  exit
  app-profile "1-1/25M" create
    description "25Mbps Site/Subscriber"
    divert
    characteristic "access-rate" value "25M"
    capacity-cost 25
  exit

7750>config>app-assure# group 1:1 policy
  app-qos-policy
    entry 500 create
      match
        traffic-direction network-to-subscriber
        characteristic "access-rate" eq "25M"
      exit
      action
        bandwidth-policer "DefltPol-Sub-BW-DS-25Mbps"
      exit
    no shutdown

7750>config>app-assure# group 1
  policer "DefltPol-Sub-BW-DS-25Mbps" type dual-bucket-bandwidth granularity subscriber create
    description "Default Policer for BW DL of Subscriber 25Mbps"
    rate 25000
    mbs 470
  exit
  
```

al_0571

Implicit Default Subscriber Policy

Session-filter, url-filter, overload-drop, fragment-drop and error-drop can only be used as part of a default subscriber policy; therefore these actions are not compatible with application or IP header match criteria within the same AQP.

AQP Entries Evaluation

Multiple AQP Match Entries Per Flow

A single flow can match multiple AQP entries, in which case multiple actions can be selected based on the AQP entry's order (the lowest number entry has the highest priority); the drop action takes precedence over any other AQP entry. The maximum numbers of actions that can be applied on a single flow are:

- 1 drop action
- Any combination of (applied only if no drop action is selected)
 - Up to 1 mirror action
 - Up to 1 FC, 1 priority and 1 DSCP remark action
 - Up to 4 BW policers (1 single rate AA-Sub, 1 dual rate AA-Sub, 2 single rate system level)
 - Up to 12 flow policers (3 subscriber flow-count, 3 subscriber flow-rate, 3 system flow-count, 3 system flow-rate)
 - Up to 1 HTTP Redirect
 - Up to 1 HTTP Error Redirect
 - Up to 1 HTTP Enrichment
 - Up to 1 URL-Filter
 - Up to 1 HTTP-Notification
 - Up to 1 Session-Filter Firewall
- 1 error-drop
- 1 overload-drop
- 1 fragment-drop

An AQP entry match that would cause the above limits to be exceeded is ignored (no actions from that rule are selected) and the conflict counter for this AQP is incremented.

The operator can display hits and potential conflicts per AQP entry using the following show command:

```
A:BNG# show application-assurance group 1:1 policy app-qos-policy
=====
Application QoS Policy Table
=====
Entry          Admin State          Flow Hits          Flow Conflicts
-----
30             in-service           0                  0
-----
No. of AQP entries: 1
=====
```

AQP Evaluation

Flows are evaluated against all entries of the AA QoS Policy at different steps during the lifetime of the flow:

- **Flow creation** — The default subscriber policy AQP entries for matching flows are applied starting with the first packet of a flow so before application identification completes.
- **Application identification completion**— The application QoS policies are applied once flow identification has been completed.



Note:

The default QoS policy entries are applied to the subscriber's flows for packets received before and after application identification is completed.

- **Policy change** — When a configuration change is applied to the AA policy by executing the commit command on the AA group:partition policy, all diverted flows for subscribers using this policy partition will be evaluated again against all AQP entries. This re-evaluation happens as a paced background task; hence AQP control changes may not be applied immediately to all existing flows.

Policing

Policers

AA policer templates are configured as part of the AA Group configuration by specifying the policer name, type and granularity. Policers are unidirectional by definition so that separate policers must be defined per flow direction if the traffic needs to be policed in both directions (a separate AQP for each flow direction is therefore required as well).

The operator can configure the following types of policers:

- Bandwidth Policers
 - Single bucket system level
 - Single bucket AA subscriber level
 - Dual bucket AA subscriber level
- Flow Count Policer: system or AA subscriber level
- Flow Setup-Rate Policer: system or AA subscriber level

Subscriber level policers are instantiated per AA sub, meaning:

- The system automatically uses a dedicated policer for every single subscriber, even when multiple subscribers match the same AQP entry.
- The same policer can be referenced in different AQP entries; in this case all subscribers' flows matching any of these AQP entries are policed by the same subscriber policer. Example: if the same subscriber level policer '1Mbps' is referenced in AQP entry 100 matching application BitTorrent and in AQP entry 110 matching application EDonkey, then the sum of both the BitTorrent and EDonkey traffic cannot exceed 1Mbps.

System level policers on the other hand are shared by all AA subscribers matching a given AQP entry. These policers are typically used in residential and Wi-Fi service deployments to limit the total bandwidth for an application or application group, for all subscribers or for a group of subscribers on the system or partition. An example would be a system level 500Mbps policer to limit the aggregated downstream bandwidth of "Peer to Peer" applications for all subscribers with a "Bronze" app-profile to 500Mbps.

**Note:**

In case multiple ISA-AA cards are used per system, the overall maximum throughput using a system level policer is equal to the policer rate limit times the number of ISA cards in the system.

Bandwidth Policing

Single Bucket Subscriber/System Bandwidth Policer

Single bucket policers police the matching traffic against a configured peak-information-rate (PIR). Traffic above the PIR can be marked as out of profile or dropped.

The configuration template for a single rate bandwidth policer is as follows:

```
BNG>config>app-assure# group 1
  policer <policer-name> type single-bucket-bandwidth
                                granularity {subscriber|system} create
    description <string>
    rate <pir-rate-in-Kbps>
    mbs <max-burst-size-in-Kbytes>
    adaptation-rule pir {max|min|closest}
    tod-override <tod-override-id>
    action permit-deny|priority-mark
```

where:

- **action** — Defines the action that must be taken by the policer for non-conforming traffic.
- **permit-deny** — Non-conforming packets will be dropped.
- **priority-mark** — Non-conforming traffic will be marked as out of profile (increasing the chances that non-conforming packets will be discarded in case of congestion on the egress queues).
- **rate** — Peak information rate in Kbps.
- **mbs** — Maximum burst size in Kbytes.
- **adaptation-rule pir <max|min|closest>** — The policers work at discrete operational rates supported by the hardware. The adaptation rule specifies how the actual operational policer rate (supported by the hardware) must be selected as compared to the configured PIR. During operation, both the operational and configured rate can be displayed using the operational **show application-assurance group <n> policer <policer-name> detail** command.
- **tod-override** — Defines a time of day override policy applicable to a policer, this is described in more detail at the end of the policing section.

A single bucket subscriber level policer configuration example is shown below:

```
BNG>config>app-assure# group 1
  policer "P2P-Sub-DL-1M" type single-bucket-bandwidth granularity subscriber create
    rate 1000
    mbs 19
  exit
```

A single bucket system level policer configuration example is shown below:

```
BNG>config>app-assure# group 1
  policer "P2P-Sys-DL-100M" type single-bucket-bandwidth granularity system create
    rate 100000
```



```
mbs 1875
exit
```

Dual Bucket Subscriber Bandwidth Policer

Dual-bucket policers police the matching traffic against a configured peak information rate (PIR) and committed information rate (CIR). Traffic below CIR is marked in profile, traffic between CIR and PIR is marked as out of profile, and traffic above the PIR is dropped.

Dual-bucket policers can only be used as subscriber policers; system policers cannot be defined as dual-bucket policers.

The configuration is similar to the single-bucket policer, but adds the configuration of a CIR and a Committed Burst Size (CBS), and the action cannot be configured:

```
BNG>config>app-assure# group 1
  policer <policer-name> type dual-bucket-bandwidth
                                granularity {subscriber|system} create
    description <string>
    rate <pir-rate-in-Kbps> cir <cir-rate-in-Kbps>
    mbs <max-burst-size-in-Kbytes>
    cbs <committed-burst-size-in-Kbytes>
    adaptation-rule pir {max|min|closest} cir {max|min|closest}
```

A dual-bucket subscriber level policer configuration example is shown below:

```
BNG>config>app-assure# group 1
  policer "P2P-Sub-DL-2M-DB" type dual-bucket-bandwidth granularity subscriber create
    rate 2000 cir 1000
    cbs 19
    mbs 38
  exit
```

MBS/CBS Calculation for Bandwidth Policers

The default MBS/CBS value of a bandwidth policer is set to 0. This value can and should be modified by the operator to allow proper interworking with TCP based applications.

The formula to calculate the MBS or CBS buffer size, as documented in RFC 6349, *Framework for TCP Throughput Testing*, is:

$$\text{Buffer (B)} = \text{Rate (bps)} / 8 * \text{RTT (s)}$$

For Internet applications it is recommended to use a common Round Trip Time (RTT) of 150 msec.

An example using a single bucket subscriber level policer rate of 1000 Kbps:

$$\text{MBS (B)} = 1,000,000 / 8 * 0.150 = 18750 \text{ Bytes or } 190 \text{ KB.}$$

Note that these policer values may need to be further adjustment depending on the application.

Flow Rate Limit Policer

Flow rate limit policers police the maximum number of new flows that are accepted per second for matching traffic. The configuration is similar to the single-bucket bandwidth policer, with the rate and MBS now expressed in flows/sec and flows, respectively.

```
BNG>config>app-assure# group 1
  policer <policer-name> type flow-rate-limit granularity {subscriber|system} create
    description <string>
    rate <flow-rate-in-flows/sec>
    mbs <max-burst-size-in-flows>
    adaptation-rule pir {max|min|closest}
    action permit-deny|priority-mark
```

This type of policer is primarily used for the default subscriber AQP policy in order to limit the maximum number of flow/seconds allocated per AA subscriber.

Note that in case the policer is used as part of the default AA subscriber policy then the **priority-mark** action has the effect to cut-through non conformant traffic in the ISA instead of drop using **permit-deny**.

Flow Count Limit Policer

Flow count limit policers police the maximum number of concurrent flows for matching traffic:

```
BNG>config>app-assure# group 1
  policer <policer-name> type flow-count-limit granularity {subscriber|system} create
    description <string>
    action permit-deny|priority-mark
    flow-count <max-number-of-flows>
```

This type of policer is primarily used for the default subscriber AQP policy in order to limit the maximum number of concurrent flows allocated per AA subscriber.

Note that the "priority-mark" has the effect to cut-through non conformant traffic in the ISA instead of drop using "permit-deny".

Time of Day Policing

Software Release 11.0.R1 introduced support for time-of-day (ToD) policer override. Up to 8 override rates with time of day specifications can be defined per policer, this time of day override using the system local time.

ToD overrides are supported for all policer types described in the previous section (bandwidth, flow-count, flow-rate) and can be configured using either daily or weekly patterns.

The configuration of ToD override on daily or weekly basis is shown in the following template:

```
BNG>config>app-assure# group 1
  policer "P2P-Sub-DL-1M-TOD" type single-bucket-bandwidth
    granularity subscriber create
    action permit-deny
    rate 1000
    mbs 19
    adaptation-rule pir closest
    tod-override <override-id>
    description <string>
```

```
time-range daily start <start-time> end <end-time>
                        [on <day> [<day>...(upto 7 max)]]
time-range weekly start <day,start-time> end <day,end-time>
rate 2000
mbs 38
```

where:

- **tod-override** <override-id> — Up to 8 override-ids (with value 1-255) can be configured per policer.
- **time-range** — Can be configured to be triggered.
 - On a daily basis at the indicated start/end-time on the specified days.
 - On a weekly basis at the indicated start day+time and end-day+time.
 - Times can be indicated as <hh>:<mm> with a 15-minute granularity for the minutes (mm = 0 | 15 | 30 | 45).

A configuration example for a single bucket system level bandwidth policer with the following ToD-override patterns follows:

- Default Rate Limit: 300Mbps
- Rate Limit override to 100Mbps between 5PM and 10PM
- Rate Limit override to 200Mbps between 10PM and 12PM

```
BNG>config>app-assure# group 1
  policer "P2P-Sys-DL-300M-TOD" type single-bucket-bandwidth
                                granularity system create
  description "Peer to Peer Policer System level Policer"
  rate 300000
  mbs 5625
  tod-override 1 create
    description "Override busy hour #1"
    time-range daily start 17:00 end 22:00
    rate 100000
    mbs 1875
    no shutdown
  exit
  tod-override 2 create
    description "Override busy hour #1"
    time-range daily start 22:00 end 24:00
    rate 200000
    mbs 3750
    no shutdown
  exit
```

The operator can display which policing rate is applied at any moment in time together with all configured override rates using the following command:

```
show application-assurance group <n> policer <policer-name> detail
```

Design and Configuration Examples

Default AA QoS Policy

To ensure fair access for all subscribers to the ISA-AA resources, and avoid that a disproportionate amount of ISA-AA resources are used by one or more subscribers which are misbehaving or receiving large traffic bursts from the Internet, it is recommended to configure the following three types of subscriber-level default AA QoS policies:

- **A default bandwidth policer** to limit the downstream bandwidth per subscriber (upstream bandwidth is already limited by ESM/SAP access ingress IOM QoS).
- **A default flow count policer** to limit the maximum number of active flows per traffic direction per subscriber. The operator can choose to drop or cut-through non conforming traffic.
- **A default flow rate policer** to limit the maximum flow setup rate per traffic direction per subscriber. The operator can choose to drop or cut-through non conforming traffic.

The minimum set of app-profiles used in a network is typically determined by the different access bandwidth rates; services characteristics are then used for each profile to apply a default QoS policy to limit bandwidth and flow resources accordingly.

In theory, it is possible to configure a set of default policers for every individual access bandwidth rate that is offered to a subscriber. This would however result in a large number of policers and corresponding ASO values plus app-profiles that need to be configured. Therefore, a best practice guideline is to define a small number of bandwidth ranges (not more than five to ten) that cover the full offered access bandwidth spectrum, and define for each bandwidth range a default bandwidth policer plus flow policers with appropriate limits.

As an example, assuming a residential deployment with 2 bandwidth ranges of up to 25Mbps and 100Mbps, the configuration below provides:

- Complete ASO and app-profile configuration.
- Default QoS policy for subscribers in the 25Mbps range including bandwidth.
- Flow count and flow rate policers are configured by default as permit-deny. Non conforming traffic is dropped which is common for residential deployments; alternatively the operator can decide to configure these policers as priority-mark to cut-through traffic in the ISA-AA.

In this example the resources are limited per subscriber based on their access rate maximum speed from which flow count and flow rate are derived.

App-Profile and ASO

The configuration below provides the app-profile and ASO characteristics used for the default subscriber AQP policy for the 25Mbps and 100Mbps access bandwidth range:

```
BNG>config>app-assure# group 1:1 policy
  app-service-options
    characteristic "access-rate" create
      value "100M"
      value "25M"
      default-value "100M"
    exit
  exit
  app-profile "1-1/25M" create
    description "25Mbps Site/Subscriber"
    divert
    characteristic "access-rate" value "25M"
    capacity-cost 25
```

```

exit
app-profile "1-1/100M" create
  description "100Mbps Site/Subscriber"
  divert
  characteristic "access-rate" value "100M"
  capacity-cost 100
exit

```

Default Bandwidth Policing – 25Mbps AA-Sub

```

BNG>config>app-assure# group 1
  policer "DefltPol-Sub-BW-DS-25Mbps" type dual-bucket-bandwidth
                                     granularity subscriber create
  description "Deflt downstream BW policer for 25Mbps Subs"
  rate 25000
mbs

```

The AQP entry below will act as a default AQP policy since it does not include application or IP Header match criteria:

```

BNG>config>app-assure# group 1:1 policy
  app-qos-policy
  entry 500 create
    description "Deflt downstream BW policer for 25Mbps Subs"
    match
      traffic-direction network-to-subscriber
      characteristic "access-rate" eq "25M"
    exit
  action
    bandwidth-policer "DefltPol-Sub-BW-DS-25Mbps"
  exit
  no shutdown
exit

```



Note:

A similar configuration can be implemented for the 100Mbps access rate service option.

Default Flow-Count-Limit Policing – 25Mbps AA-Sub

```

BNG>config>app-assure# group 1
  policer "DefltPol-Sub-FlowCount-US-25Mbps" type flow-count-limit
                                     granularity subscriber create
  description "Deflt policer to limit active upstream flows for 25Mbps Subs"
  flow-count 10000
  action permit-deny
exit
  policer "DefltPol-Sub-FlowCount-DS-25Mbps" type flow-count-limit
                                     granularity subscriber create
  description "Deflt policer to limit active downstream flows for 25Mbps Subs"
  flow-count 10000
  action permit-deny
exit

```

The AQP entry below will act as a default AQP policy since it does not include application or IP Header match criteria:

```
BNG>config>app-assure# group 1:1 policy app-qos-policy
  entry 510 create
    description " Deflt policer to limit active upstream flows for 25Mbps Subs"
    match
      traffic-direction subscriber-to-network
      characteristic "access-rate" eq "25M"
    exit
    action
      flow-count-limit "DefltPol-Sub-FlowCount-US-25Mbps"
    exit
    no shutdown
  exit
  entry 515 create
    description " Deflt policer to limit active downstream flows for 25Mbps Subs"
    match
      traffic-direction network-to-subscriber
      characteristic "access-rate" eq "25M"
    exit
    action
      flow-count-limit "DefltPol-Sub-FlowCount-DS-25Mbps"
    exit
    no shutdown
  exit
```



Note:

A similar configuration can be implemented for the 100Mbps access rate service option.

Default Flow-Rate-Limit Policing – 25Mbps AA-Sub

```
BNG>config>app-assure# group 1
  policer "DefltPol-Sub-FlowRate-US-25Mbps" type flow-rate-limit
    granularity subscriber create
    description "Deflt policer to limit upstream flow setup rate for 25Mbps Subs"
    rate 200
    action permit-deny
  exit
  policer "DefltPol-Sub-FlowRate-DS-25Mbps" type flow-rate-limit
    granularity subscriber create
    description "Deflt policer to limit downstr flow setup rate for 25Mbps Subs"
    rate 200
    action permit-deny
  exit
```

The AQP entry below will act as a default AQP policy since it does not include application or IP Header match criteria:

```
BNG>config>app-assure# group 1:1 policy app-qos-policy
  entry 520 create
    description "Deflt policer to limit upstream flow setup rate for 25Mbps Subs"
    match
      traffic-direction subscriber-to-network
      characteristic "access-rate" eq "25M"
    exit
    action
      flow-rate-limit "DefltPol-Sub-FlowRate-US-25Mbps"
    exit
```

```

    no shutdown
  exit
  entry 525 create
    description "Deflt policer to limit downstr flow setup rate for 25Mbps Subs"
    match
      traffic-direction network-to-subscriber
      characteristic "access-rate" eq "25M"
    exit
    action
      flow-rate-limit "DefltPol-Sub-FlowRate-DS-25Mbps"
    exit
    no shutdown
  exit
exit

```

**Note:**

A similar configuration can be implemented for the 100Mbps access rate service option.

Application BW Policing (Per Subscriber)

The configuration example below provides a per AA subscriber peer-to-peer rate limit of 1Mbps. It does not include the app-profile configuration since the ASO characteristic and values can be either statically configured within the app-profile or dynamically signaled through RADIUS or Gx using ASO overrides.

AA subscribers with service characteristic "P2P-Sub-DL" value of "1M" will have a bandwidth policer of 1Mbps applied to peer to peer traffic in the network to subscriber direction:

```

BNG>config>app-assure# group 1
  policer "P2P-Sub-DL-1M" type single-bucket-bandwidth granularity subscriber create
    description "Per-subscr BW policer to limit P2P downstream traffic to 1Mbps"
    rate 1000
    mbs 19
    action permit-deny
  exit

BNG>config>app-assure# group 1:1 policy
  app-service-options
    characteristic "P2P-Sub-DL" create
      value "10M"
      value "1M"
      value "unlimited"
      default-value "unlimited"
  exit

BNG>config>app-assure# group 1:1 policy app-qos-policy
  entry 30 create
    description "Per-subscr BW policer to limit P2P downstream traffic to 1Mbps"
    match
      app-group eq "Peer to Peer"
      traffic-direction network-to-subscriber
      characteristic "P2P-Sub-DL" eq "1M"
    exit
    action
      bandwidth-policer "P2P-Sub-DL-1M"
    exit
    no shutdown
  exit
exit

```

Conclusion

This example provides detailed information to properly configure and use app-profiles, ASOs and AQPs to successfully configure application policy control rules using Application Assurance.

Application Assurance — Asymmetry Removal

This chapter describes Application Assurance asymmetry removal configurations.

Topics in this chapter include:

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

Applicability

This chapter was originally written for and configured on SR OS Release 11.0.R1. The CLI in the current edition corresponds to SR OS Release 14.0.R4.

The prerequisites for this chapter are a base understanding of AA configuration and operation for single homed deployments. This chapter applies to dual-homed SAPs and spoke SDPs configurations, in a business or residential AA context. AARP is not used for ESM AA subscribers.

Overview

This chapter is intended for Application Assurance (AA) network architects and engineers. It provides best practices recommendations to configure AA asymmetry removal.

Asymmetry means that the two directions of a traffic flow (to-sub and from-sub) take different paths through the network. Asymmetry removal is a means of eliminating traffic asymmetry between a set of dual-homed SAP or spoke SDP endpoints. This can be across endpoints within a single node or across a pair of inter-chassis link connected routers, which is the topology explained in this chapter. Asymmetry removal ensures all packets of a dual-homed AA subscriber are diverted to an AA ISA in order to achieve accurate per subscriber traffic identification and policy enforcement.

Traffic asymmetry is created when there are dual-homed links for a service, and the links are simultaneously carrying traffic. Asymmetry removal for transit subscribers must be implemented in the first routed hop on the network side of the subscriber management point, so there will be a deterministic and fixed SAP/spoke SDP representing the downstream subscriber management node. This ensures there are no more than two paths that the flows can take, both covered by the asymmetry removal solution.

Configuration

Application Assurance Redundancy Protocol (AARP) provides the data plane connectivity for dynamically keeping a dual-homed AA subscriber's traffic on the same ISA-AA for AA processing. An AARP instance is configured between the dual-homed routers to establish connectivity with the same AARP instance number on each node.

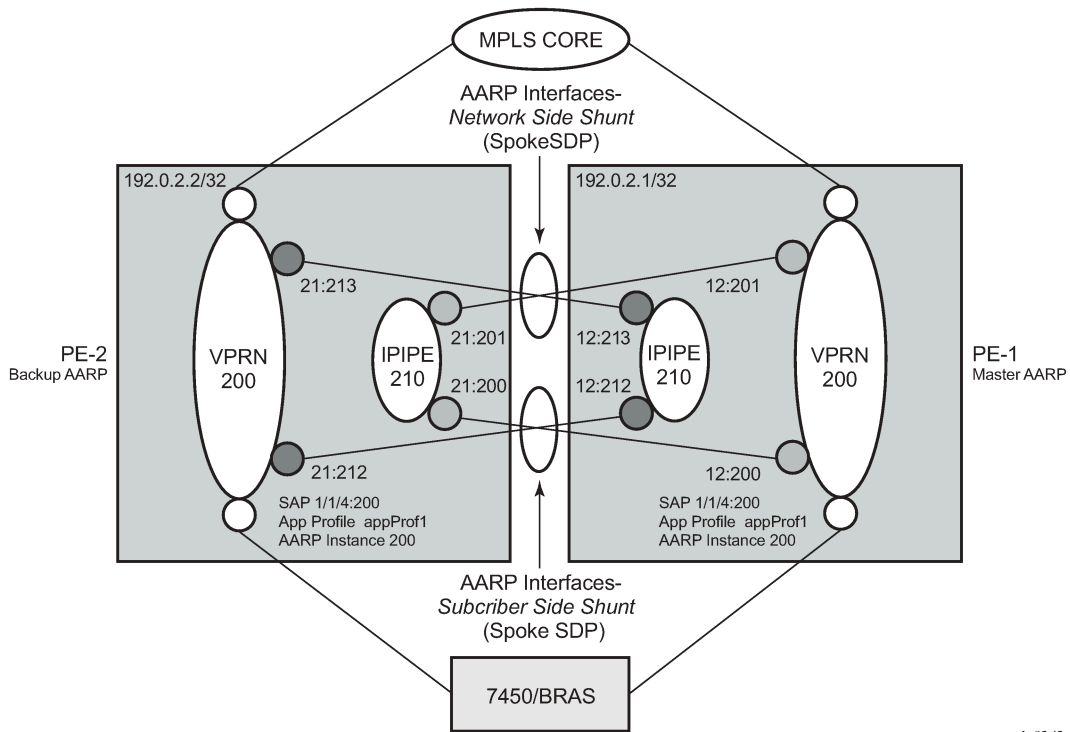
When asymmetry exists between dual-chassis redundant systems, lpipe spoke SDPs are used to interconnect these services between peer nodes over an Inter-Chassis Link (ICL). The following sections explain the configuration and operation of the services for use with the Application Assurance Redundancy Protocol.

AARP service configuration

The following services must be configured to establish communications between the AARP instances in each of the paired nodes.

- Network topology is a VPRN (or IES) service configured in each node, with a dual-homed SAP from each node to a downstream access element.
- Assumes starting point with AA ISAs installed with identical AA policy and divert enabled in each node.
- Also, the system needs basic routing and LDP configuration for the SDP and the spoke SDPs to be established.

Figure 12: AA asymmetry removal topology



al_0242

Table 4: AA asymmetry removal topology

| On PE-2 | On PE-1 |
|----------------------------|----------------------------|
| system ip: 192.0.2.2 | system ip: 192.0.2.1 |
| dual-homed service: 200 | dual-homed service: 200 |
| dual-homed sap: 1/1/4:200 | dual-homed sap: 1/1/4:200 |
| app-profile diverting: yes | app-profile diverting: yes |

Configuration commands for AARP

To enable AARP, AARP instances and AARP interfaces on both nodes must be configured. AARP operation has the following dependencies between the nodes:

- Shunt links configured and operationally up, both subscriber side shunt and network side shunt.
- Peer communications established between nodes, AARP instance operational status will be up when peers are communicating.
- Dual-homed sap/spoke SDP configured with a unique AARP instance (matched by dual-homed interface).
- App-profile configured against sap/spoke SDP with divert enabled (making the sub an aa-sub). The app-profile is the trigger to divert the traffic in the node with the active AARP instance to one of the ISAs in that node, per normal AA divert behavior.

Begin with PE-2:

```
configure
  application-assurance
    aarp 200 create
      description "aarp protecting a dual-homed sap"
      priority 100
      peer 192.0.2.1
      no shutdown
    exit
  exit
exit
```

lpipe shunt configuration

```
configure
  service
    sdp 21 mpls create
      far-end 192.0.2.1
      ldp
      keep-alive
      shutdown
    exit
    no shutdown
  exit
  ipipe 210 customer 1 vc-switching create
    service-mtu 1556
    spoke-sdp 21:200 create
      aarp 200 type subscriber-side-shunt
      no shutdown
    exit
    spoke-sdp 21:201 create
      aarp 200 type network-side-shunt
      no shutdown
    exit
    no shutdown
  exit
exit
exit
```

Dual-homed and interface shunt configuration

```
configure
  service
    vprn 200 customer 1 create
```

```

description "VPRN 200 Dual Homed Routed Service"
aarp-interface "subside_1" create
    spoke-sdp 21:212 create
        aarp 200 type subscriber-side-shunt
        no shutdown
    exit
exit
aarp-interface "netside_1" create
    spoke-sdp 21:213 create
        aarp 200 type network-side-shunt
        no shutdown
    exit
exit
interface "int-BRAS-1" create
    sap 1/1/4:200 create
        aarp 200 type dual-homed
        app-profile "app-prof-1"
    exit
exit
no shutdown
exit
exit
exit

```

Then similarly configure the associated AARP configuration on PE-1:

```

configure
    application-assurance
        aarp 200 create
            description "aarp protecting a dual-homed sap"
            priority 200
            peer 192.0.2.2
            no shutdown
        exit
    exit
exit

```

lpipe shunt configuration

```

configure
    service
        sdp 12 mpls create
            far-end 192.0.2.2
            ldp
            keep-alive
            shutdown
        exit
        no shutdown
    exit
    lpipe 210 customer 1 vc-switching create
        service-mtu 1556
        spoke-sdp 12:212 create
            aarp 200 type subscriber-side-shunt
            no shutdown
        exit
        spoke-sdp 12:213 create
            aarp 200 type network-side-shunt
            no shutdown
        exit
        no shutdown
    exit
exit

```

```
exit
```

Dual-homed and interface shunt configuration

```
configure
service
  vprn 200 customer 1 create
  aarp-interface "subside_1" create
  spoke-sdp 12:200 create
  aarp 200 type subscriber-side-shunt
  no shutdown
  exit
exit
aarp-interface "netside_1" create
spoke-sdp 12:201 create
aarp 200 type network-side-shunt
no shutdown
exit
exit
interface "int-BRAS-1" create
sap 1/1/4:200 create
description "AA enabled SAP"
aarp 200 type dual-homed
app-profile "app-prof-1"
exit
exit
no shutdown
exit
exit
exit
```

Show commands for AARP

Verify correct configuration on each node. The following output displays the example configuration for PE-1.

Starting with the AARP instance in each node, verify that the AARP instance operational state is up (if everything is properly configured as intended):

```
*A:PE-1# show application-assurance aarp 200
=====
AARP Instance 200
=====
Description      : aarp protecting a dual-homed sap
Admin State      : Up                               Oper State      : Up
Local IP         : 192.0.2.1                         Peer IP         : 192.0.2.2
Local State      : master                           Peer State      : backup
Local Priority    : 200                             Peer Priority    : 100

Local Flags      : none
Peer Flags       : none
Peer End-Point   : none

Master Selection Mode      : minimizeSwitchovers

-----
Service References
-----
Service      Reference      Reference Type
```

```

-----
VPRN 200          1/1/4:200          Dual-Homed
Ipipe 210         12:212          Subscriber-Side Pipe Shunt
Ipipe 210         12:213          Network-Side Pipe Shunt
VPRN 200          12:200          Subscriber-Side AARP-Interface Shunt
VPRN 200          12:201          Network-Side AARP-Interface Shunt
-----
No. of service references: 5
-----
=====
*A:PE-1#

```

Verifying that the AARP instance is up is an indication that the dual-node communications for AARP is working (instance, shunts, etc.). In addition, in the preceding output, verify on both PE nodes that the intended SAPs are dual-homed for that instance.

Now a detailed review of the configured AARP shunt infrastructure services can be shown to make sure they are all properly configured with intended AARP parameters (such as AARP ID and Type on the network and subscriber side shunts) as displayed in the following output:

```

*A:PE-1# show service id 210 all

=====
Service Detailed Information
=====
Service Id       : 210                Vpn Id          : 0
Service Type    : Ipipe
Name            : (Not Specified)
Description     : (Not Specified)
Customer Id     : 1                  Creation Origin  : manual
Last Status Change: 10/03/2016 11:45:51
Last Mgmt Change : 10/03/2016 11:45:51
Admin State     : Up                 Oper State      : Up
MTU             : 1556
Vc Switching    : True
SAP Count       : 0                  SDP Bind Count  : 2
CE IPv4 Discovery : n/a             Keep address    : No
CE IPv6 Discovery : n/a             Stack Cap Sig   : n/a

Eth Legacy Fault Notification
-----
Recovery Timer   : 10.0 secs           Admin State     : outOfService

-----
ETH-CFM service specifics
-----
Tunnel Faults    : ignore

-----
Service Destination Points(SDPs)
-----
Sdp Id 12:212   -(192.0.2.2)
-----
Description     : (Not Specified)
SDP Id          : 12:212                Type           : Spoke
Spoke Descr    : (Not Specified)
Split Horiz Grp : (Not Specified)
VC Type        : Ipipe                 VC Tag         : 0
Admin Path MTU : 0                   Oper Path MTU  : 1556
Delivery       : MPLS
Far End        : 192.0.2.2

```

```

Tunnel Far End      : 192.0.2.2          LSP Types           : LDP
Hash Label          : Disabled           Hash Lbl Sig Cap    : Disabled
Oper Hash Label     : Disabled
Entropy Label       : Disabled

Admin State         : Up                 Oper State          : Up
MinReqd SdpOperMTU : 1556
Acct. Pol          : None                Collect Stats       : Disabled
Ingress Label       : 262141             Egress Label        : 262139

---snip---

Application Profile: None
Transit Policy      : None
AARP Id             : 200
AARP Type           : subscriber-side-shunt

---snip---

-----
IPIPE Service Destination Point specifics
-----
Configured CE IPv4 Addr: n/a                Peer CE IPv4 Addr : 0.0.0.0

-----
Sdp Id 12:213 - (192.0.2.2)
-----
Description        : (Not Specified)
SDP Id             : 12:213                Type                : Spoke
Spoke Descr        : (Not Specified)
Split Horiz Grp    : (Not Specified)
VC Type            : Ipipe                 VC Tag              : 0
Admin Path MTU     : 0                     Oper Path MTU       : 1556
Delivery           : MPLS
Far End            : 192.0.2.2
Tunnel Far End     : 192.0.2.2          LSP Types           : LDP
Hash Label         : Disabled           Hash Lbl Sig Cap    : Disabled
Oper Hash Label    : Disabled
Entropy Label      : Disabled

Admin State         : Up                 Oper State          : Up
MinReqd SdpOperMTU : 1556
Acct. Pol          : None                Collect Stats       : Disabled
Ingress Label       : 262140             Egress Label        : 262138

---snip---

Application Profile: None
Transit Policy      : None
AARP Id             : 200
AARP Type           : network-side-shunt

---snip---
=====

```

Next, the configuration of the VPRN service of the dual-homed SAP can be reviewed to ensure it reflects the attached endpoints for the shunt lpipe spoke SDPs:

```

*A:PE-1# show service id 200 all
=====
Service Detailed Information
=====

```

```

Service Id      : 200                Vpn Id         : 0
Service Type   : VPRN
Name           : (Not Specified)
Description    : (Not Specified)
Customer Id    : 1                  Creation Origin : manual
Last Status Change: 10/03/2016 11:45:51
Last Mgmt Change : 10/03/2016 11:45:51
Admin State    : Up                Oper State     : Up

Route Dist.    : 64496:200         VPRN Type     : regular
Oper Route Dist : 64496:200
Oper RD Type   : configured
AS Number      : None              Router Id      : 192.0.2.1
ECMP           : Enabled           ECMP Max Routes : 1
Max IPv4 Routes : No Limit
Auto Bind Tunnel
Resolution     : disabled
Max IPv6 Routes : No Limit
Ignore NH Metric : Disabled
Hash Label    : Disabled
Entropy Label  : Disabled
Vrf Target    : target:64496:200
Vrf Import    : None
Vrf Export    : None
MVPN Vrf Target : None
MVPN Vrf Import : None
MVPN Vrf Export : None
Car. Sup C-VPN : Disabled
Label mode    : vrf
BGP VPN Backup : Disabled
BGP Export Inactv : Disabled

SAP Count     : 1                  SDP Bind Count : 2
VSD Domain    : <none>

```

---snip---

Service Destination Points(SDPs)

Sdp Id 12:200 - (192.0.2.2)

```

Description    : (Not Specified)
SDP Id         : 12:200                Type           : Spoke
Spoke Descr    : (Not Specified)
VC Type       : n/a                   VC Tag         : n/a
Admin Path MTU : 0                     Oper Path MTU  : 1556
Delivery      : MPLS
Far End       : 192.0.2.2
Tunnel Far End : 192.0.2.2                LSP Types     : LDP
Hash Label    : Disabled               Hash Lbl Sig Cap : Disabled
Oper Hash Label : Disabled
Entropy Label  : Disabled

Admin State    : Up                Oper State     : Up

```

---snip---

```

Application Profile: None
Transit Policy    : None
AARP Id           : 200
AARP Type         : subscriber-side-shunt

```



```

---snip---
-----
IPIPE Service Destination Point specifics
-----
Configured CE IPv4 Addr: n/a                Peer CE IPv4 Addr : 0.0.0.0
-----
Sdp Id 12:201  -(192.0.2.2)
-----
Description      : (Not Specified)
SDP Id           : 12:201                Type                : Spoke
Spoke Descr     : (Not Specified)
VC Type         : n/a                    VC Tag              : n/a
Admin Path MTU  : 0                      Oper Path MTU       : 1556
Delivery        : MPLS
Far End         : 192.0.2.2
Tunnel Far End  : 192.0.2.2             LSP Types           : LDP
Hash Label      : Disabled               Hash Lbl Sig Cap    : Disabled
Oper Hash Label : Disabled
Entropy Label   : Disabled

Admin State     : Up                      Oper State          : Up
-----
---snip---

Application Profile: None
Transit Policy   : None
AARP Id          : 200
AARP Type        : network-side-shunt
-----
---snip---

```

Continuing deeper into the same VPRN service show output, or using the following show command, it can be verified that the dual-homed SAP itself is properly configured and associated with that service and AARP instance:

```

*A:PE-1# show service id 200 sap 1/1/4:200 detail
=====
Service Access Points(SAP)
=====
Service Id       : 200
SAP              : 1/1/4:200             Encap              : q-tag
Description      : AA enabled SAP
Admin State     : Up                      Oper State         : Up
Flags           : None
Multi Svc Site  : None
Last Status Change : 10/03/2016 11:45:51
Last Mgmt Change  : 10/03/2016 11:45:51
Sub Type        : regular
Dot1Q Ethertype : 0x8100                    QinQ Ethertype     : 0x8100
Split Horizon Group: (Not Specified)

Admin MTU        : 1518                    Oper MTU           : 1518
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      : None                Dynamic Hosts      : Enabled
Avl Static Hosts   : 0                   Tot Static Hosts   : 0
Calling-Station-Id : n/a

Application Profile: app-prof-1
Transit Policy     : None
AARP Id           : 200
AARP Type         : dual-homed

Oper Group        : (none)                Monitor Oper Grp   : (none)
Host Lockout Plcy : n/a
Lag Link Map Prof : (none)
Bandwidth         : Not-Applicable
```

---snip---

Network to subscriber traffic flow


When the AARP is operationally up, AARP tracks which ISA is the master ISA for each dual-homed AARP instance and uses the inter-chassis services (spoke SDP AARP shunts) to move all traffic for each instance traffic to the node with the Master ISA.

Looking at traffic in the network to subscriber direction ([Figure 13: Network to subscriber traffic flow](#)):

- Traffic arriving on PE-1 is diverted to the local master ISA, processed, then proceeds to the egress SAP.
- Traffic arriving on PE-2 with the backup AARP interface is sent to the master node for AA processing. The ingress FP forwards packets to network-side-shunt AARP interface for remote AA divert.
- Arriving on PE-1, the packets on the AARP Ipipe are diverted to the master ISA where the packets are processed as if this traffic was traveling in the to-sub direction towards the dual-homed endpoint on PE-1, then returned to PE-2.
- Entering PE-2, the traffic from the subscriber side shunt interface is not diverted to ISAs in that node and egresses on the AARP instance SAP.

With this behavior, traffic always returns to the original ingress node before egressing toward the subscriber (network path for the flows are not modified).

Figure 13: Network to subscriber traffic flow




al_0243

Subscriber to network traffic flow

Looking at traffic in the subscriber to network direction ([Figure 14: Subscriber to network traffic flow](#)):

- Traffic arriving on PE-1 is diverted to the local master ISA, processed, then proceeds to the egress SAP.
- Traffic arriving on PE-2 with the backup AARP ISA is sent to the master node for AA processing (not diverted to an ISA in PE-2). The ingress FP forwards packets to subscriber-side-shunt AARP interface for remote AA divert.
- Arriving on PE-1, the packets on the AARP Ipipe are diverted to the master ISA where the packets are processed as if the traffic was flowing in the from-sub direction on the dual-homed endpoint, then returned to PE-2 over the Ipipe's AARP subscriber-side-shunt.
- Entering PE-2, the traffic from the network side shunt interface is forwarded by the IES/VPRN service to its destination.

Figure 14: Subscriber to network traffic flow



al_0243

Typical configuration mistakes

Operators configuring AARP can make some typical mistakes listed below that will keep the AARP instance in operational state down:

- The spoke SDP AARP shunt instances' IDs must be aligned with the respective spoke SDP on the peer node: if not, it will result in a flag indicating **shunt mismatch** in the show output.
- Ipipe service MTU alignment — The Ipipe service MTU values must be the same in both nodes, otherwise it will result in the services be in operational status UP, but the AARP instance will remain down.

Conclusion

This chapter is intended for Application Assurance (AA) network architects and engineers to provide the information required to understand and configure dual-node asymmetry removal following the intended service configuration as used by the AARP implementation.

Application Assurance — Best Practices for ISA and Host IOM Overload Protection

This chapter provides information about Application Assurance best practices for ISA and host IOM overload protection.

Topics in this chapter include:

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

Applicability

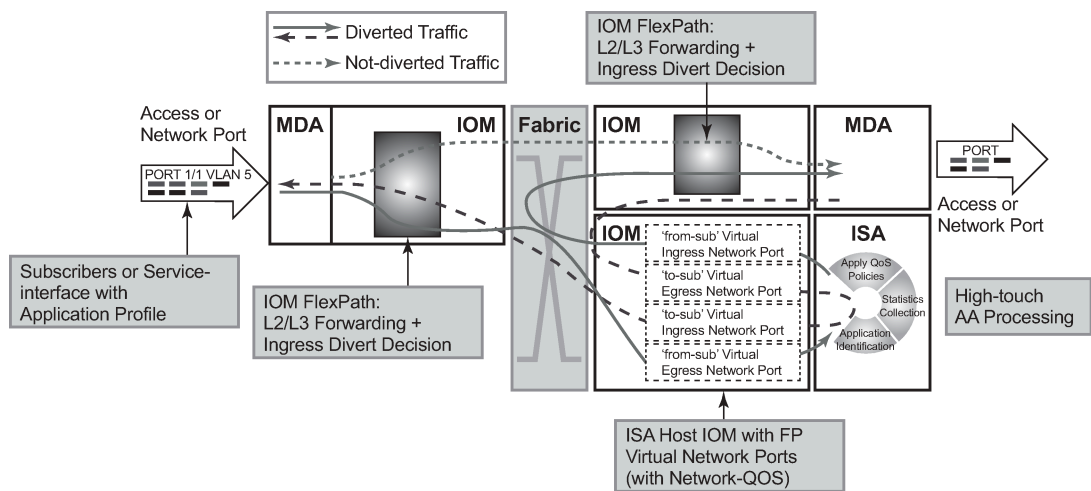
The information and configuration in this chapter is based on SR OS Release 12.0.R4.

Overview

The multiservice integrated services adapter (MS-ISA) is a processing resource module installed on an ISA host IOM. This example describes the best practices for configuration and monitoring of the system to ensure proper engineering of the system resources involved in AA ISA capacity planning.

As shown in [Figure 15: System packet datapath to AA ISA](#), traffic is diverted to an AA ISA by provisioning of an application profile (app-profile) for a subscriber or SAP service context. SR OS then automatically handles traffic diversion for both directions of traffic for that AA subscriber context, through one of the AA ISAs in the AA group where that app-profile is defined.

Figure 15: System packet datapath to AA ISA



a1_0535

The following elements in the SR OS node must be properly engineered for any given AA deployment. Each element is described in this section:

1. ISA capacity cost and load balancing across ISAs.
2. ISA host IOM network egress QoS. Host IOM egress network ports weighted-average shared buffer pool thresholds (within the egress QoS configuration for each group) are used for overload cut-through processing.
3. ISA resources and statistics collection.
 - Flows
 - Traffic volume (bandwidth)
 - Subscribers
 - Flow setup rate
 - ISA overload cut-through
 - ISA default subscriber policies

ISA capacity planning approach

This example illustrates an approach to the configuration of the AA system to address these considerations:

- IOM/ISA-AA network egress QoS configuration should be designed to treat the ISA as a network port with normal network port maximum delay (by MBS).
- Within the ISA, fair access to the ISA-AA bandwidth and flow resources must be ensured: it is recommended that default application QoS policy (AQP) policy entries be configured limiting bandwidth and flow resources per AA subscriber.
- Thresholds for SNMP alerts that indicate a high load on ISA processing should be configured: capacity cost, flow, bandwidth.
- Capacity tracking in live deployments should be performed for parameters that can affect overload: flow setup rate, bandwidth, and subscriber-count per ISA.
- Use of other scale related consumable AA resources against system maximum limits. This includes parameters such as statistics records, transit-ip table entries, and transit-prefix TCAM entries, which should be planned and periodically tracked. These limits will not affect overload of the ISA but may affect intended service operation.
- For recommendations of the specific parameters to watch in a given deployment as well as the values of the system limits for a given release, contact your regional support organization.

AA overload and resource monitoring

Overload is a condition where the total packet processing requirements for traffic arriving on a given ISA exceeds the available resources, resulting in the host IOM egress buffers reaching a configured "overload" threshold. Above this threshold, the ISA can be configured to forward excess traffic (called overload cut-through). If cut-through is not enabled and the overload condition continues, the egress queue MBS threshold will eventually be reached, after which packets will be discarded. Even if overload cut-through is enabled, any egress traffic that exceeds the maximum bus capacity of the ISA queue discard threshold will be discarded.

ISA capacity overload events are supported within the system resource monitoring and logging capabilities if the traffic and resource load crosses any of the following high and low load thresholds on a per-ISA basis. Exceeding one of these thresholds does not in itself indicate an overload state.

1. Host IOM egress network ports weighted-average shared buffer pool thresholds (within the egress QoS configuration for each AA group) are used for triggering and removing overload cut-through processing. Care should be taken in the configuration of these buffers, as the IOM flexpath has significant buffer capacity that can result in latency larger than the network SLA acceptable guidelines. A properly engineered configuration will have large enough buffering to not trigger ISA overload unnecessarily (due to normal bursts with a reasonable traffic load) but will not incur excessive latency prior to triggering the overload state.
2. ISA capacity cost: if the capacity cost of all subscribers on the ISA exceeds a threshold, an event is raised but the overload condition is not set (unless other resources are exhausted). ISA overload or traffic cut-through does not occur simply if capacity cost is exceeded. It is used to capacity plan an intended load for the ISA, proportional to resource use per subscriber, in order to generate events prior to overload to allow appropriate action to mitigate the resource consumption (such as provisioning more ISAs).
3. Flow table consumption (number of allocated flow resources in use): the flow table high-watermark threshold warnings are for proactive notification of a high load. The ISA will cut-through new flows when the "flow resources in use" is at the maximum flow limit. Reaching the flow limit does not generate backpressure to the IOM, nor is the ISA considered in an overload state. Flow usage thresholds are different from bit-rate/packet-rate/flow-setup-rate thresholds in that when the flow table high-watermark threshold is exceeded, the ISA will no longer be operating as application-aware for the flows with no context. The default subscriber policy is applied to traffic that required a flow record but was unable to allocate one, which is a similar behavior to overload cut-through.

The following terms are used to describe flow resources:

- Maximum flows: the maximum AA flow table size for a given release.
 - Flows: on the show screens, the "flows" field is an indication of the number of unique 5-tuple entries in the flow table. This includes active and inactive flows; inactive will age out of the table after a period of inactivity that is dependent on the protocol used.
 - Active flows: the number of flows with traffic in the current reporting interval.
 - Flow resources in use: the number of allocated flows in the flow table. This number is greater than the number of active flows, reflecting inactive flows and flows pre-allocated for some dynamic protocols (control + data channels) and for some UDP traffic.
4. Traffic volume: traffic rate in bytes/sec and packets/sec is the dominant cause of ISA overload in most network scenarios, when the ISA is presented with more traffic than it can process. This results in the ISA internal ingress buffers reaching a threshold that causes backpressure to the IOM egress queues (toward the ISA), allowing the ISA to process the packets it already has. This internal backpressure mechanism is normal behavior, allowing burst tolerance at the IOM-to-ISA interface; thus backpressure is not in itself an indication of overload. Overload occurs when the bursts or the load of traffic is sustained long enough to reach the ISA host IOM network port egress weighted-average shared buffer threshold. The actual amount of traffic that can be passed through an ISA is dependent on the application traffic mix, flow density, and AA policy configurations and will vary by network type and by region. The bit-rate and packet-rate watermarks can be used to provide event notification when the traffic rates exceed planning expectations.
 5. Flow setup rate: this is generally proportional to total traffic volume, and as such can be a factor in ISA overload. The flow setup rate is the rate at which new flows are presented to the ISA, each resulting in additional tasks that are specific to flow state creation; thus the ISA has a sensitivity to flow setup rates

as fewer cycles are available for other datapath tasks when the flow setup rate is high. In residential networks, flow setup rates of 3 k to 5 k flows/sec per Gbps of traffic are common. The flow setup watermarks can be used to provide event notification when the rate exceeds planning expectations.

ISA Overload Models

For an ISA overload strategy, there are two design options for configuring the overload behavior of the system:

- **Host IOM egress discards:** in this model, the philosophy is to treat AA packet processing resources in the same way as a network interface (of somewhat variable capacity depending on the traffic characteristics). When too much traffic is presented to the ISA, it backpressures the host IOM egress, which will buffer packets. If the egress buffer thresholds are exceeded, the ISA will discard according to the egress QoS slope policy. This is configured by not enabling **isa-overload-cut-through** and use of appropriate egress QoS policies. Firewall or session filter deployments may use this model.
- **Overload cut-through:** the ISA group can be enabled to cut-through some traffic if an overload event occurs, triggered when the IOM network port weighted-average queues depth exceeds the weighted-average shared high-watermark threshold. In this ISA state, some packets are cut-through from application analysis but retain subscriber context with the default subscriber policy applied. This mode of deployment is intended for situations where it is preferable to forward packets even if not identified by AA than to drop/discard the packet. For example, if AA is providing value-added services (VAS) such as In-Browser Notification (IBN), analytics, or traffic rate limiting, this would usually be the preferred model as the underlying service should be preserved even if capacity to provide the VAS is not available.

Note that even with overload cut-through enabled, there is a hardware-based maximum ISA throughput of approximately 11 Gbps for MS-ISA and 40 Gbps for MS-ISA2. If this is exceeded on a sustained basis, IOM egress discards may still occur.

Understanding Packet and Protocol Cut-Through

Traffic can be cut-through the ISA-AA card on a packet-by-packet basis, in which case packets do not go through AA identification and subscriber application policy. The conditions that trigger cut-through include:

- **Overload (IOM egress network port weighted-average shared buffer threshold):** excess traffic bypasses all AA processing except for the default subscriber policy
- **Non-conformant IP packet:** traffic bypasses all AA processing except IP protocol checks and the default subscriber policy. Optionally, these packets can be discarded in AA.
- **Flow table full:** for new 5-tuples sent to the ISA, if the flow table is full, the packets are cut-through the ISA and only the default subscriber policy is applied.



Note:

The default subscriber policy is a set of AQP rules that apply AQP match criteria limited to Application Service Options (ASO), aa-sub, and traffic-direction starting with the first packet of a flow, with no match conditions based on AA identification (application, app-group, charging-group, IP header). Packets will be either `denied_by_default_policy` or `cut_through_by_default_policy`, depending on the policer action configuration in the AQP rules.

For cut-through traffic, no flow records exist but it is counted under per-subscriber protocol statistics as one of the following counters, depending on the case:

- `cut_through` — Statistics for any packet that could not map to a flow, but that has a valid subscriber ID. This can be an error packet, fragmented out-of-order, no flow resource, invalid TCP flags, etc. This is the most important count for indicating overload cut-through, as it counts all traffic in overload cut-through mode (when the weighted-average threshold has been crossed).
- `denied_by_default_policy` — Packets that are dropped due to a default policy with a flow-based policer (flow rate or flow count) with action discard.
- `cut_through_by_default_policy` — Packets that failed to pass flow-based policers with an action of priority-mark.

An example of overload cut-through statistics in the CLI is shown below:

```
A:BNG# show application-assurance group 1 protocol count
=====
Application-Assurance Protocol Statistics
=====
Protocol                Disc          Octets        Packets        Flows
-----
advanced_direct_connect 0%             0             0             0
aim                     0%             0             0             0
amazon_video            0%             0             0             0
ares                    0%             0             0             0
bbm                     0%             0             0             0
betamax_voip            0%             0             0             0
bgp                     0%             0             0             0
bittorrent              0%          678428534     5322929       1036129
cccam                   0%             0             0             0
citrix_ica              0%             0             0             0
citrix_ima              0%             0             0             0
cnnlive                 0%             0             0             0
cups                    0%             0             0             0
cut_through             0%          5299435739    10603771      0
cut_through_by_default_policy 0%             0             0             0
cvs                     0%             0             0             0
daap                    0%             0             0             0
dcerpc                  0%             0             0             0
denied_by_default_policy 0%             0             0             0
```

Configuration

This example illustrates a typical configuration of an SR OS node for AA for each of the configuration topics.

AA traffic load test environment

Application assurance identifies every byte and every packet of hundreds of real-world applications using per-flow stateful analysis techniques. It is a challenge to find test equipment that can accurately emulate full scale (10 Gbps to 40 Gbps) with traffic mixes and flow behaviors representing hundreds of thousands of end users with application clients across a range of devices. Some specialized stateful test equipment can emulate large traffic rates, but even the best will have equipment-specific patterns and behaviors not representative of live traffic. Therefore, the best scenario to engineer the AA overload configuration is by iteration in live deployments: setting an initial target and modifying the configuration based on ISA performance under load.

For a lab test of ISA throughput and loading, Nokia uses stateful test equipment which supports emulation of various service provider traffic mix profiles suitable for generating overload conditions; however, it is outside the scope of this document to configure AA throughput tests.

The operator should be aware that use of unrealistic, non-stateful traffic generators can result in a high level of unknown traffic, with the ISA performance impacted by continually trying to identify large numbers of packets of no real application type. This, combined with cut-through for invalid IP packets, can result in ISA overload and traffic cut-through (due to overload or invalid IP packets) at traffic levels not representative of actual ISA performance on real traffic.

ISA capacity cost and load balancing across ISAs

These AA group-level commands define the load balancing parameters within an ISA group.

```
*A:BNG# configure isa
  application-assurance-group 1 aa-sub-scale residential create
  no description
  no fail-to-open
  isa-capacity-cost-high-threshold 304000
  isa-capacity-cost-low-threshold 272000
  partitions
  divert-fc be
  no shutdown
exit
```

The following should be noted related to this configuration:

- Up to 7 primary and 1 backup ISAs are allowed. If the AA services are considered "value added" and not part of a paid service, backups are usually not used because the "fail to fabric" capability keeps the underlying service running.
- The default behavior in case of ISA failure is "no fail-to-open", which means "fail-to-wire"; if an ISA fails, traffic is forwarded as if **no divert** was configured
- Threshold for sending capacity-cost SNMP traps: the unit used for capacity cost is a variable defined in the network design; in this example, it is expressed in Mbps of the subscriber total BW UP+DOWN with a high watermark set to $7600 \text{ Mbps} \times 40 = 304000$ (where 40 is an oversubscription ratio). The low watermark is equal to $6800 \text{ Mbps} \times 40 = 272000$.
- Partitions should always be enabled to configure additional policies in the future (for example, wifi/business)
- **divert-fc** configuration applies to the AA group: in this example, FC BE Internet is the only diverted FC; this is typical for AA residential and Wlan-GW deployments. For VPN services, typically all datapath FCs are diverted to AA.

ISA-AA host IOM - network egress shared memory and QoS

The amount of shared memory allocated per port, along with the network port egress QoS policy, determine the maximum delay for traffic diverted to Application Assurance.

This maximum network port delay is typically determined by the operator and must be used to define the proper QoS configuration to apply to the ISA-AA ports; this QoS configuration may be the same (typically) as what is applied to regular network ports on the SR OS node.

On the line cards there is shared network egress memory per ISA-AA port, with the ISA-AA is represented by two network ports on the host IOM:

- "from-sub": for traffic sent from the subscriber to the network
- "to-sub": for traffic sent from the network to the subscriber

```
configure isa application-assurance-group 1
    qos
        egress
            from-subscriber
                pool
                    slope-policy "default"
                    resv-cbs default
                exit
                queue-policy "network-facing-egress"
                port-scheduler-policy "network-facing"
            exit
            to-subscriber
                pool
                    slope-policy "default"
                    resv-cbs default
                exit
                queue-policy "network-facing-egress"
                port-scheduler-policy "network-facing"
            exit
        exit
    no shutdown
```

The amount of shared memory reserved for each egress network port is determined by the speed of the port (10 Gbps for MS-ISA and 40 Gbps for MS-ISA2) and the **egr-percentage-of-rate** ratio configuration.

MS-ISA uses by default 1000% and 500% of the rate respectively for to-sub and from-sub ports, while MS-ISA2 uses by default 100% for both to-sub and from-sub ports.

It is typically recommended that these values be adjusted when MS-ISA and a high-speed Ethernet MDA are mixed on the same IOM, because in this context the amount of shared memory allocated to the Ethernet MDA should be increased by reducing the MS-ISA network ports memory allocation ratio. If two MS-ISAs are installed on the same IOM, the system will by default allocate 50% of the network egress shared memory to each ISA. In addition, an operator may adjust these values in case the actual network-to-subscriber versus subscriber-to-network ratio is significantly different in the production network, in order to achieve the expected maximum tolerated network delay.

The operator can modify the **egr-percentage-of-rate** per port using the following command:

```
A:BNG# configure port 1/2/fm-sub
A:BNG>config>port# info detail
-----
    modify-buffer-allocation-rate
        egr-percentage-of-rate 500
    exit
-----
A:BNG# configure port 1/2/to-sub
A:BNG>config>port# info detail
-----
    modify-buffer-allocation-rate
        egr-percentage-of-rate 1000
    exit
```

Network egress scheduling/queuing priority is for all ISAs within a group defined at the AA ISA group level

An example below with ISA-AA and 2 x 10G Eth MDA:

```
7750# configure port <slot>/<isa-aa-mdm>/fm-sub
      modify-buffer-allocation-rate
      egr-percentage-of-rate 65

7750# configure port <slot>/<isa-aa-mdm>/to-sub
      modify-buffer-allocation-rate
      egr-percentage-of-rate 130
```

In this example, the configuration defines:

- from-sub — Approximately 190 msec worth of buffer at 2500 Mbps.
- to-sub— Approximately 190 msec worth of buffer at 5000 Mbps.
- The buffer can be further refined from the network QoS policy.

For MS-ISA2, each MS-ISM flexpath will default the buffer allocation rate to 100%, which is a suitable value assuming that both modules in a slot are MS-ISA2 (which is the MS-ISM configuration), or that the I/O module has a similar traffic rate as the MS-ISA2 (which is also the case in the 10x10GE and 1x100GE versions of the MS-ISA2 line cards).

Configuring ISA resources and stats collection

The following are the key consumable resources in an AA ISA:

- Flows
- Bandwidth
- Subscribers
- Flow setup rate

The AA group should be configured with watermark thresholds where each ISA will generate SNMP events when resources reach this level.

- Per-ISA-card resource usage watermarks trigger SNMP traps to the management system (5620 SAM)
- The values defined below can be refined based on the network characteristics in term of flows and bandwidth per ISA after the initial deployment

```
7750# configure application-assurance
-----
      flow-table-low-wmark 90
      flow-table-high-wmark 95
      flow-setup-high-wmark 66500
      flow-setup-low-wmark 63000
      bit-rate-high-wmark 7600
      bit-rate-low-wmark 6800
```

In this example, the usage SNMP watermarks are configured for:

- Flow table: 95%/90% (maximum 4M flows on MS-ISA)
- Flow setup rate: configured to 95%/90% (of maximum 70k fps on MS-ISA)
- Bit rate/total diverted throughput

The **show app-assure group status detail** command is used to display basic ISA health status:

- # aa-sub, active aa-sub, bitrate, flows in use, flow setup rate
- statistics for all ISAs combined or per ISA

```
A:BNG# show application-assurance group 1 status detail
=====
Application-Assurance Status
=====
Last time change affecting status : 05/30/2014 17:18:34
Number of Active ISAs             : 4
Flows                             : 214007945881
Flow Resources In Use             : 2955164
AA Subs Created                   : 70567
AA Subs Deleted                   : 10544
AA Subs Modified                  : 0
Seen IP Requests Sent             : 0
Seen IP Requests Dropped         : 0
-----
Current      Average      Peak
-----
Active Flows      : 2911508      2769454      4582522
Flow Setup Rate (per second) : 33923        29400        67865
Traffic Rate (Mbps) : 7620          7238         22628
Packet Rate (per second) : 1254138      1182571      3044376
AA-Subs Downloaded : 69887        66129        70567
Active Subs       : 23131        19737        38114
-----
Packets      Octets
-----
Diverted traffic      : 7437950197613      5530634242355947
Diverted discards    : 0                  0
  Congestion         : 0                  0
  Errors              : 0                  N/A
Entered ISA-AAs      : 7437950180191      5530634229794634
Buffered in ISA-AAs  : 22                 29849
Discarded in ISA-AAs : 97790              47801217
  Policy              : 0                  0
  Congestion         : 0                  0
  Errors              : 97790             47801217
Modified in ISA-AAs
  Packet size increased : 0                  0
  Packet size decreased : 0                  0
Errors (policy bypass) : 28283549           21160338635
Exited ISA-AAs       : 7437950082379      5530634181963568
Returned discards    : 0                  0
  Congestion         : 0                  0
  Errors              : 0                  N/A
Returned traffic     : 7437950054070      5530634162337570
=====
```

This can also be run on a per-ISA basis:

```
show application-assurance group 1 status isa <slot/port> detail
```

Note that for MS-ISA2, there is a maximum AA packet rate of 7 M pps; under most known traffic mix scenarios, the ISA should be safely below this packet rate when at maximum bandwidth throughput. However, it is worth periodically checking this value, because if the maximum packet rate is exceeded, and overload cut-through will result. (For MS-ISA, the maximum packet rate supported is high enough to not be feasible with realistic application-based traffic mixes).

The ISA aa-performance record should always be enabled in a network for capacity planning purposes in order to properly plan when to add new ISA cards if required and to monitor the network health:

```
*A:BNG>config>isa# info
-----
  application-assurance-group 1 aa-sub-scale residential create
    no description
    primary <slot/port>
    backup <slot/port>
    no fail-to-open
    isa-capacity-cost-high-threshold 304000
    isa-capacity-cost-low-threshold 272000
    partitions
    statistics
      performance
        accounting-policy 7
        collect-stats
      exit
    exit
    divert-fc be
    no shutdown
  exit
```

The commands highlighted in bold above will export information on the total traffic load and resource utilization of the ISA card:

- Flows — active flows, setup rates, resource allocation
- Traffic rates — bandwidth, packets
- Subscribers — active, configured, statistics resource allocation in use

The AA statistics collection configuration refers to accounting policies that are also defined in the SR OS node:

```
*A:BNG>config# log
  file-id 7
  description "ISA Performance Stats"
  location cf2:
  rollover 15 retention 12
  exit
  accounting-policy 7
  description "ISA Performance Stats"
  collection-interval 15
  record aa-performance
  to file 7
  no shutdown
  exit
```

From the AA performance record the following fields in [Table 5: Tracking ISA load in the reporting interval](#) can be used as to tracking ISA load in the reporting interval (typically a 15 to 60 minute period):

Table 5: Tracking ISA load in the reporting interval

| Record name | Type | Description | Load planning use |
|-------------|------------|---|--------------------------------------|
| dco | cumulative | octets discarded due to congestion in MDA | Should be 0; ISA internal congestion |

| Record name | Type | Description | Load planning use |
|-------------|--------------|--|---|
| dcp | cumulative | packets discarded due to congestion in MDA | Should be 0; ISA internal congestion |
| dpo | cumulative | octets discarded due to policy in MDA | Not related to load planning |
| dpp | cumulative | packets discarded due to policy in MDA | Not related to load planning |
| pbo | cumulative | octets policy bypass | Not used. Traffic was for an invalid subscriber and the group was "no fail-to-open" |
| pbp | cumulative | packets policy bypass | Not used. Traffic was for an invalid subscriber and the group was "no fail-to-open" |
| nfl | cumulative | number of flows | informative |
| caf | intervalized | current active flows | informative |
| aaf | intervalized | average active flows | informative |
| paf | intervalized | peak active flows | Check vs max |
| cfr | intervalized | current flow setup rate | informative |
| afr | intervalized | average flow setup rate | Check meets expected norms; increasing over time increases load |
| pfr | intervalized | peak flow setup rate | informative |
| ctr | intervalized | current traffic rate | informative |
| atr | intervalized | average traffic rate | Check meets expected norms; increasing over time increases load |
| ptr | intervalized | peak traffic rate | Check vs max |
| cpr | intervalized | current packet rate | informative |
| apr | intervalized | average packet rate | informative |
| ppr | intervalized | peak packet rate | informative |
| cds | intervalized | current diverted subscribers | informative |
| ads | intervalized | average diverted subscribers | informative |
| pds | intervalized | peak diverted subscribers | Check vs max and expected norms; increasing over time increases load |

| Record name | Type | Description | Load planning use |
|-------------|--------------|-------------------|--|
| rfi | intervalized | flows in use | Check vs max and expected norms; increasing over time increases load |
| rcc | cumulative | ISA capacity cost | Check meets expected norms; increasing over time increases load |

The intended deployment model is for this statistic record to be collected by a centralized service-aware management system along with all other AA records and to be stored in a reporting and analysis management database for subsequent analytics purposes, such as trending charts or setting thresholds of key values. It is recommended that a CRON script be used to export the AA performance record to a storage server for post processing if no reporting and analysis management tool is deployed:

- If no reporting and analysis management tool is deployed in the network, it is possible to automatically collect the XML accounting files and provide high-level reporting through an XML-to-CSV conversion.
- The simplest approach is to configure a CRON script on the SR OS node to automatically retrieve the CF accounting file (alternatively, any other scripting mechanism with an interval smaller than the retention period can be used)
- It is recommended that the rollover interval of the file-id policy be modified to 6H or above in order to collect fewer files while keeping the same collection interval.

```
*A:BNG# file type cf2:/script
file copy cf2:/act/*.gz ftp://login:password@IP-ADDRESS/acct/router1/

*A:BNG>config>cron# info
-----
    script "test-ftp-act"
        location "cf2:/script"
        no shutdown
    exit
    action "cron1"
        results "ftp://login:password@IP-ADDRESS/results/router1-result.log"
        script "test-ftp-act"
        no shutdown
    exit
    schedule "schedule1"
        interval 36000
        action "cron1"
        no shutdown
    exit
```

The schedule : interval 36000 is in seconds (10 hours).

With this XML to CSV export mechanism, a spreadsheet can be used by the network engineer to periodically track the ISA resource utilization.

ISA overload cut-through

The system can be configured to react to overload based on the weighted-average (WA) queue depth of the shared network port buffer pool from-sub and to-sub. Overload cut-through is typically recommended for use of AA for value-added services where, in the event of overload, the preference is for the ISA to

continue to pass packets without AA processing. However, firewall use cases will prefer to drop excess traffic in the event of overload, in which case overload cut-through may not be desired.

In addition to triggering an alarm, further packets sent to the ISA after the WA high-watermark threshold is reached are cut-through immediately by the ISA card without application identification or subscriber policy processing, if the **isa-overload-cut-through** command is enabled.

The WA queue depth is typically configured based on the maximum tolerated delay for the service diverted and the amount of shared buffer space allocated from the IOM.

AA deployment recommended settings:

- high watermark — 33% of the maximum MBS for all diverted network queues
- low watermark — 5% of the maximum MBS for all diverted network queues

The recommended high and low watermarks assume that the sum of the network port egress queues MBS size is 100% of the shared buffer. If this network queue maximum size is further reduced in the network QoS policy, the watermark values must be adapted proportionally; for example, if the total MBS size cannot exceed 50% of the shared buffer, then the watermark values would be divided by 2: the high watermark = $33\% / 2 = 16\%$, the low watermark = $5\% / 2 = 2\%$. Adjusting the MBS and the **wa-shared-high-wmark** and **wa-shared-low-wmark** values proportionately ensures that the MBS point (after which discards occur) is above the WA shared high-watermark threshold; otherwise, the ISA will not ever overload if MBS discards are occurring first.

```
A:BNG# configure isa application-assurance-group 1
      isa-overload-cut-through
      qos
        egress
          from-subscriber
            wa-shared-high-wmark 16
            wa-shared-low-wmark 2
          exit
        to-subscriber
          wa-shared-high-wmark 16
          wa-shared-low-wmark 2
        exit
      exit
    exit
```

The **show isa group** commands can be used to verify that overload cut-through is enabled.

```
*A:BNG>show isa application-assurance-group 1
=====
ISA Application-assurance-groups
=====
ISA-AA Group Index      : 1
Description             : (Not Specified)
Subscriber Scale        : residential
WLAN GW Group Index    : N/A
Primary ISA-AA         : 1/2 up/active
Backup ISA-AA          : 2/1 down
Last Active change     : 07/02/2014 12:17:45
Admin State            : Up
Oper State             : Up
Diverted FCs          : be
Fail to mode           : fail-to-wire
Partitions             : enabled
QoS
  Egress from subscriber
    Pool                : default
    Reserved Cbs       : default
```

```

Slope Policy      : default
Queue Policy     : default
Scheduler Policy  :
Egress to subscriber
Pool             : default
Reserved Cbs    : default
Slope Policy     : default
Queue Policy     : default
Scheduler Policy  :
Capacity Cost
High Threshold   : 4294967295
Low Threshold    : 0
Overload Cut Through : enabled
Transit Prefix
Max IPv4 entries : 0
Max IPv6 entries : 0
Max IPv6 remote entries : 0
HTTP Enrichment
Max Packet Size  : 1500 octets
=====

```

To monitor the load status of an ISA, enter the following CLI command.

```

*A:BNB>show application-assurance group 1 status isa 5/1 cpu
=====
Application-Assurance ISA CPU Utilization
(Test time 993791 uSec)
=====
Management CPU Usage
-----
Name                CPU Time      CPU Usage
                   (uSec)
-----
System              14277         1.43%
Management          61101         6.15%
Statistics          69850         7.02%
Idle                848563        85.39%
=====
Datapath CPU Usage
-----
Name                CPU Time      CPU Usage
                   (uSec)
-----
System              14277         1.43%
Packet Processing   61101         6.15%
Application ID      69850         7.02%
Idle                848563        85.39%

```

Additionally, the system log files can be used to examine the AA overload history to determine when the overload state was entered and exited. It can be helpful to send AA events to a separate log using the following configuration:

```

log
  filter 45
    default-action drop
    entry 10
      action forward
      match
        application eq "application_assurance"
      exit
    exit
  exit

```

```
log-id 45
  description "application-assurance log"
  filter 45
  from main
  to memory 500
exit
```

The log files can then be examined to see if overload has occurred, and how frequently. If overload occurs with any regularity, it is a situation that should be addressed. Below is an example of a log file showing AA overload:

```
A:BNG# show log log-id 45
=====
Event Log 45
=====
Description : application-assurance log
warning: 13 events dropped from log
Memory Log contents [size=500  next event=16  (not wrapped)]

15 2014/08/14 17:00:32.66 EST WARNING: APPLICATION_ASSURANCE #4433 Base
"ISA AA Group 1 MDA 5/1 exiting overload cut through processing."

14 2014/08/14 17:00:32.55 EST WARNING: APPLICATION_ASSURANCE #4431 Base
"ISA-AA group 1 MDA 5/1 wa-shared buffer use is less than or equal to 1% in the to-subscriber
direction or corresponding tmnxBsxIsaAaGrpToSbWaSBuf0vld notification has been disabled."

13 2014/08/14 17:00:32.06 EST WARNING: APPLICATION_ASSURANCE #4432 Base
"ISA AA Group 1 MDA 5/1 entering overload cut through processing."

12 2014/08/14 17:00:32.05 EST WARNING: APPLICATION_ASSURANCE #4430 Base
"ISA-AA group 1 MDA 5/1 wa-shared buffer use is greater than or equal to 35% in the to-
subscriber direction."
```

The primary indicator to look at in CLI statistics for ISA load indication is datapath CPU usage. Regardless of the configuration and traffic profiles in use, datapath CPU usage gives a consistent indication of whether the ISA is under heavy load (the cause of overload is the inability of the ISA to perform more tasks). The idle datapath time is not proportionate to bandwidth throughput, but if idle datapath CPU usage is under 5%, this indicates an approaching maximum processing load.

At an average datapath use of 95-100% (less than 5% idle) the ISA is creating latency and backpressuring the host IOM egress. It is the best way to know how close to overload the ISA has been. Attempting to examine data throughput statistics such as bit rate, flow setup rate and packet rate to predict overload is not recommended, as these are quite variable under normal circumstances and are not directly correlated to overload. Once in overload, the data statistics (volume, setup rate, and so on) are useful for determining what threshold traps to put in place for the future, but the needed thresholds will always be specific to the live deployment traffic mix and policy configuration.

Below is an example of the status for an ISA that is fully loaded but not yet in overload:

```
*A:BNG>show application-assurance group 1 status isa 5/1 cpu
=====
Application-Assurance ISA CPU Utilization
=====

-----
Management CPU Usage (Test time 999636 uSec)
-----
Name                CPU Time      CPU Usage
                   (uSec)
-----
```

| | | |
|--|--------------------|-----------|
| System | 1540 | 0.15% |
| Management | 14 | ~0.00% |
| Statistics | 643955 | 64.42% |
| ICAP Client | 603 | 0.06% |
| Idle | 353524 | 35.37% |
| ----- | | |
| Datapath CPU Usage (Test time 999735 uSec) | | |
| ----- | | |
| Name | CPU Time (uSec) | CPU Usage |
| ----- | | |
| System | 188374 | 18.84% |
| Packet Processing | 534203 | 53.43% |
| Application ID | 277158 | 27.72% |
| Idle | 0 | 0.00% |
| ----- | | |

In this example, 0% idle datapath CPU means the ISA is fully used. When the Datapath CPU Usage Idle average is in the 5-10% range consistently, the ISA should be considered "full"; to add new subscribers, more ISAs are required.

If the excessive traffic condition persists, backpressure from the ISA to the IOM will buffer packets in the egress buffers, and when the egress MBS is exceeded, the ISA host IOM will indicate diverted discards due to congestion if cut-through is not enabled:

```
*A:BNG>show application-assurance group 1 status detail
=====
Application-Assurance Status
=====
Last time change affecting status : 08/12/2014 13:16:15
Number of Active ISAs             : 1
Flows                             : 235754165
Flow Resources In Use             : 12000000
AA Subs Created                   : 14224
AA Subs Deleted                   : 0
AA Subs Modified                  : 1
Seen IP Requests Sent             : 0
Seen IP Requests Dropped         : 0
-----
Current      Average      Peak
-----
Active Flows      : 8452434      3786948      10632607
Flow Setup Rate (per second) : 246578      65104      298677
Traffic Rate (Mbps) : 33702      13229      35813
Packet Rate (per second) : 6847697      2466118      6945936
AA-Subs Downloaded : 14224      13710      14224
Active Subs       : 14224      9934      14224
-----
Packets      Octets
-----
Diverted traffic      : 8924242848      5983284952320
Diverted discards : 752486      729147667
Congestion      : 752486      729147667
Errors               : 0      N/A
Entered ISA-AAs     : 8923417360      5982508976617
Buffered in ISA-AAs : 57      19277
Discarded in ISA-AAs : 0      0
Policy              : 0      0
Congestion          : 0      0
Errors              : 0      0
Modified in ISA-AAs
```

| | | |
|------------------------|--------------|---------------|
| Packet size increased | : 0 | 0 |
| Packet size decreased | : 0 | 0 |
| Errors (policy bypass) | : 0 | 0 |
| Exited ISA-AAs | : 8923417303 | 5982508957340 |
| Returned discards | : 0 | 0 |
| Congestion | : 0 | 0 |
| Errors | : 0 | N/A |
| Returned traffic | : 8923285123 | 5982432640249 |
| ===== | | |

ISA default subscriber policy

Default subscriber policy — AQP with match criteria not using App-ID or 5-tuple. Match **only** includes traffic direction and/or ASO characteristic and/or subscriber name.

It is recommended that each ISA be configured with some default subscriber policies that get applied to all subscribers at all times, independent of application flow ID, and even when an ISA is in overload cut-through. These policies protect the ISA resources and provide fairness of resource allocation between subscribers by limiting the ISA resources that can be consumed by a single subscriber. A starting point for the recommended policers is (in all cases, network-specific tuning is recommended):

- Per-subscriber flow rate policer: value more than expected maximum peak per-subscriber rate for active subscribers. The policer protects one subscriber from attacking the network with an excessive flow rate and affecting ISA flow rate resources used by other customers. A typical rate for residential networks could be 100 fps per subscriber.
- Per-subscriber flow count policer: value more than expected maximum per-subscriber flow count for active subscribers. The policer protects one subscriber from consuming excessive flow counts and affecting ISA flow resources used by other customers.
- Downstream bandwidth per subscriber: to a value more than the maximum rate supported by the service, or to less than the maximum per-subscriber capability of the ISA, whichever is lower. For fixed networks, several default policer rates are recommended using a per-sub ASO value for low, medium and large rate ranges set at a rate related to the subscriber access speed. For example, for an FTTH service the per-sub policers could be set at 3 value ranges: below 25Mbps, with another at 100Mbps sub policer for services between 25Mbps and 100Mbps, and another sub-policer for rates between 100Mbps and 1Gbps. The settings for a mobile 3G network rate may be 1Mbps and in an LTE network the rate may be vhc10Mbps.

For a CLI example of a default subscriber policy, see [Application Assurance — App-Profile, ASO and Control Policies](#).

Conclusion

Any deployment of Application Assurance should include careful capacity planning of the ISA resources, with an appropriate ISA overload strategy, whether for overload cut-through to keep excess traffic flowing, or with a discard policy engineered in the host IOM egress QoS policies.

ISA resource use should be monitored via appropriately configured resource thresholds, events, log files, XML records and show screens to ensure that sufficient ISA resources are available as required.

Application Assurance — DNS IP Cache

This chapter describes the Application Assurance - DNS IP Cache.

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 R5.

Overview

This chapter is intended for network architects and engineers working with Application Assurance (AA). It provides an example to configure the DNS IP Cache, which can be used by operators to ensure that traffic will be correctly classified to prevent charging fraud.

This chapter assumes that the reader has a basic understanding of AA policy (application-filters, applications, application-qos-policies). Additionally, the reader is expected to know the basics of HTTP and DNS.

Nokia recommends using the AA policy configuration templates provided by the Nokia Application Database (AppDB). The AppDB contains an optimized configuration to classify all supported applications and includes all relevant application groups, applications, and app-filters. It can be obtained through Technical Support.


AA uses various methods to classify traffic. Criteria include L3/4 information (IP addresses and ports), or L7 information (HTTP hostnames or other HTTP headers). Classification of Internet services using L3/4 information alone is not very common because IP addresses may change (service providers may add or remove IP addresses), or IP addresses may be region-specific. A more reliable method is to use L7 information, usually the HTTP hostname.

To access a website, after receiving the URL, a browser will perform two basic tasks:

1. Generate a DNS request to resolve the hostname to an IP address
2. Send an HTTP request to the web server, using the IP address from the preceding step, and display the web server response to the user browser

[Figure 16: Basic Message Flow When Accessing Website](#) shows the basic message flow.

Figure 16: Basic Message Flow When Accessing Website



27196

HTTP is not a strict protocol. Several web server implementations have not implemented the standard completely and may allow requests that appear invalid. The most common case is an HTTP request in which the destination IP address of the packet does not correspond to the hostname.

For example, consider that:

- The IP address of www.domain1.com is a.b.c.d (IP1).
- The IP address of www.domain2.com is e.f.g.h (IP2).

If the web server running at IP1 receives an HTTP request containing the hostname www.domain2.com , it will reply and provide the content of www.domain1.com . This may happen because several web server implementations (for example, HTTP1.0) ignore the hostname requested.

This may lead to misclassification and incorrect handling of the flow. More specifically, some users may take advantage of this, as follows:

- Assume that www.domain2.com is a free-of-charge (zero rated) site and www.domain1.com is a chargeable site. Any classification and charging function (including AA) that uses only the hostname will classify the traffic as traffic toward www.domain2.com .
- As a result, the user will receive the content of www.domain1.com for free.

There are Internet applications that use this technique to bypass the operator charging policies. For malicious purposes, some applications manipulate (forge) the HTTP headers. Subscribers appear to be browsing a zero-rated site, while traffic is directed to a chargeable site. The applications tunnel traffic to a proxy server. The HTTP host header is a zero-rated site and the host that the subscriber wants to visit is inserted in another header (usually the x-online-host header).

To protect against such host header attacks, AA has implemented the DNS IP Cache feature.

AA will populate a cache with the DNS responses generated when the user requests to access a website, by inspecting DNS traffic. During traffic classification, apart from checking the hostname included in the

HTTP request, AA will ensure that the destination IP of the packet matches an entry contained in the cache table.

For example, consider that the cache currently contains the following:

| Domain | IP address |
|-----------------|--------------|
| www.domain1.com | 172.20.1.135 |
| | 172.20.1.136 |

The operator can configure an app-filter with match criteria, as follows:

- Hostname equal to www.domain1.com
- Server address contained in the cache

The system will then verify that the packet contains the hostname configured and the packet destination is an IP address contained in the cache. Only if both criteria are true, the app-filter will match.

Regardless of how access to www.domain1.com is charged, the operator will be certain that this traffic will be correctly classified and charged and no other traffic will be charged at this rate (which may be zero-rated).

For additional protection, the operator can also configure a list of trusted DNS servers. The cache will be populated only with the responses from those DNS servers. This ensures that the cache will not be corrupted with responses from malicious DNS servers.

Operators are advised to use the DNS IP Cache feature for any web traffic that the operator charges at a lower rate or for web traffic that is free. This includes operator websites or partner sites.

Some users will attempt to exploit the operator configuration, trying to obtain content at reduced rates or free of charge. Enabling the feature in the configuration will validate the hostname included in the HTTP request, preventing charging fraud.

Configuration

The example configuration comprises a 7750 SR configured with ISA-AA. The configuration requires a web client and Internet access through SR. The same general DNS IP cache configuration also applies to AA deployed in VSR, MG, and CMG systems.

To pass traffic between two endpoints and verify the configuration, an Epipe service can be created.

The following example assumes:

- The user has the IP address 192.168.2.11.
- The zero-rated website is "www.nokia.com".
- The IP address of the trusted DNS server is 192.0.2.1.

The following sections show how to configure a DNS IP cache to contain the IP addresses of the website www.nokia.com. Only the IP addresses provided by the trusted DNS server will be added in the cache. Finally, an app-filter will be configured whose match criteria will be both the domain and the IP addresses contained in the cache.

To verify the configuration, the operator needs to generate DNS traffic and HTTP traffic to www.nokia.com (as shown in [Figure 16: Basic Message Flow When Accessing Website](#)).

Creating a DNS IP Cache

Create a **dns-ip-cache** entry and populate it with DNS responses for "*.nokia.com\$", only from the DNS server 192.0.2.1:

```
*A:Dut-C>config>app-assure>group# info
-----
dns-ip-cache "nokia-cache" create
  description "dns cache for nokia"
  dns-match
    domain "domain1" expression "*.nokia.com$"
    server-address 192.0.2.1
  exit
no shutdown
-----
```

Create an app-filter

Nokia recommends using the AA AppDB for a complete library of application and app-filter definitions. The AppDB contains verified configurations for all common applications. Operators are advised to consult the AppDB, then edit one of the predefined app-filters for the application to be verified.

It is common for operators to define app-filters for domains that are not included in the AppDB, such as on-net websites:

```
app-filter
  entry 100 create
    expression 1 http-host eq "*.sponsor1-operator.com$"
    server-address eq dns-ip-cache "dns-ip-cache1"
    application "Sponsor Content #1"
  no shutdown
```

Create the Application for the Nokia Website

The configuration is as follows:

```
*A:Dut-C>config>app-assure>group# info
-----
policy
  begin
    application "Nokia NET" create
      description "Nokia website"
      app-group "Web"
    exit
  commit
exit
-----
```

Create an App-filter for Nokia Traffic, Which Uses the DNS IP Cache

In the example, it is assumed that Nokia content is fetched from the host www.nokia.com.

```
*A:Dut-C>config>app-assure>group# info
-----
```

```

policy
  begin
    app-filter
      entry 61000 create
        description "Nokia Web Access"
        expression 1 http-host eq "*.nokia.com$"
        server-address eq dns-ip-cache "nokia-cache"
        application "Nokia NET"
        no shutdown
      exit
    exit
  commit
exit
-----

```

The preceding app-filter will match traffic if the hostname matches *.nokia.com\$ and if the destination IP address of the packet matches an entry in the cache.

Create application-qos-policy entries

Create a default subscriber policy to inspect DNS responses. This application QoS policy (AQP) is used to populate the cache. No match criteria are needed. AA will only analyze DNS traffic to populate the cache.

```

*A:Dut-C>config>app-assure>group# info
-----
policy
  begin
    app-qos-policy
      entry 50 create
        action
          dns-ip-cache "nokia-cache"
        exit
        no shutdown
      exit
    exit
  commit
exit
-----

```

Operational Considerations

1. When configuring the DNS IP cache (chapter 3.1), configuration of the following parameters was omitted:

- size (10)
- high-wmark (90)
- low-wmark (90)

The default values—shown in parentheses—were used.

- **size** refers to the maximum number of entries in the cache (it can be between 10 and 32000). The cache size is set to values above 1000 in live deployments, depending on the operator use cases. If the cache becomes full, new IP addresses will be ignored.

- **high-wmark** defines the high watermark value for the DNS IP cache (percentage). When the number of IP addresses stored in the cache exceeds the threshold defined, the system will generate a trap.
 - **low-wmark** defines the low watermark value for the DNS IP cache (percentage). Assuming the high watermark value was reached and a trap was generated, the system will clear the trap if the number of IP addresses stored in the cache drops below the low watermark value.
2. Apart from the number of IP addresses that can be stored in the cache (10 in our example), there are also a maximum number of domains that can be stored. These maximum values are indicated in the **show** command outputs described later.
 3. The AQP configured to populate the cache uses a default subscriber policy. A default subscriber policy does not contain any match criteria against application (application, app-group, charging-group) or L3/4 information (IP address or prefix, TCP/UDP port, IP protocol, DSCP). The AQP will be applied to all matching flows starting with the first packet of a flow. Match criteria can only be by application service options, traffic direction, and optional AA subscriber name.
 4. The app-filter was created by consulting app-filters defined in the AppDB. The entry ranges have been defined in the AppDB. The ranges defined ensure that after Nokia upgrades the AppDB, the custom protocols (defined by the operator) will not be affected. Therefore, the operator should use an entry in the correct range, so that the configuration is not removed after a policy upgrade done with the policy sync tool. User-defined application groups must prefix their description with the text "Custom", thus facilitating future policy upgrades.
 5. If the cache becomes full, new DNS responses will still increase the "Cache full count" and will still get a hit on "DNS Total responses", "Domain name matched", and "Domain & server matched", assuming that they match (see the next section for the **show** command output). An IP address will expire and be removed from the cache if it is not included in any DNS response for a specified period of time.

Troubleshooting and Debugging

Using the **tools** and **show** commands, the operator can verify whether DNS responses are present in the cache and that the configuration is valid.

Show Commands

The following show commands can be used to evaluate the configuration created and verify whether traffic has matched.

The following command shows the DNS IP cache:

```
*A:Dut-C# show application-assurance group 1 dns-ip-cache "nokia-cache"

=====
Application Assurance Group 1 dns-ip-cache "nokia-cache"
=====
Admin Status           : Up
AQP Ref                : Yes
Domain expressions     : 1 (out of 32)
Server addresses       : 1 (out of 64)
High watermark         : 90%
Low watermark          : 80%
Cache size             : 10
```

| ISA | Usage | (%) | Alarm State | Hit Count |
|-----|-------|------|-------------|-----------|
| 1/2 | 0 | 0.00 | clear | 8 |
| 3/1 | 0 | 0.00 | clear | 8 |

The preceding output provides the following information:

- The administrative status of the DNS IP cache "nokia-cache" is "Up". An administrative status of "Down" indicates that the DNS IP cache is not in use.
- The DNS IP cache "nokia-cache" is referenced by an AQP.
- The number of configured domain expressions (in our example, "*.nokia.com\$") equals 1 and the maximum value (shown in brackets) is 32.
- The number of server addresses is 1 and the maximum value (shown in brackets) is 64.
- The number of times a cache lookup was successful (hit count) is 8. A hit count of 0 indicates a possible configuration error.

The following command displays the content of each ISA card:

```
*A:Dut-C# show application-assurance group 1 dns-ip-cache "nokia-cache" isa 1/2
=====
Application Assurance Group 1 dns-ip-cache "nokia-cache" ISA 1/2
=====
Admin Status           : Up
AQP Ref                : Yes
Domain expressions     : 1 (out of 32)
Server addresses       : 1 (out of 64)
High watermark         : 90%
Low watermark          : 80%
Cache size             : 10
-----
ISA 1/2 DNS Stats
-----
DNS
  Total responses      : 5
  Domain name matched : 5
  Domain & server matched : 5
Cache
  Total entries added  : 2
  Total entries removed : 1
Usage
  1 (10.00%) threshold alarm clear
  Full count           : 0
  Hit count            : 8
  Miss count           : 0
=====
```

The preceding output provides additional information about the:

- total number of DNS responses that were analyzed (Total responses)
- number of times the domain name defined in the DNS match criteria matched a DNS response (Domain name matched)
- number of times both the domain name and server address defined in the DNS match criteria matched a DNS response (Domain & server matched)

- total number of entries added in the cache (2 in the preceding example)
- total number of entries removed from the cache (1 in the preceding example). An entry will be removed from the cache if an internal timer expires. This ensures that the cache does not maintain very old IP addresses, which may no longer be valid. If the IP address is re-learned, the timer will be reset.
- successful/unsuccessful IP address lookups in the cache (Hit/Miss count)

If one or more of the preceding values indicates an error, the configuration should be checked. The domain/server configured may not match the hostname requested (for example, the operator configured the domain "*.nokia.com" while the user requested the hostname "*.nokia.co.uk").

The following command shows the created app-filter:

```
*A:Dut-C# show application-assurance group 1 policy app-filter 61000
  app-filter
    entry 61000 create (0 flows, 0 B)
      description "Nokia Web Access "
      expression 1 http-host eq "*.nokia.com$"
      server-address eq dns-ip-cache "nokia-cache"
      application "Nokia NET"
      no shutdown
    exit
  exit
```

The following command shows the AQP used to populate the cache and verifies that traffic is matched.

```
*A:Dut-C# show application-assurance group 1 policy app-qos-policy 50

=====
Application QoS Policy Entry 50 (Default Subscriber Policy)
=====
Description : (Not Specified)
Admin State : in-service
Hits       : 2 flows
Conflicts  : 0 flows

Match :

Action :
  DNS IP Cache           : nokia-cache
=====
```

The Admin State "in-service" indicates that the AQP is in use. If no traffic has matched and the number of hits remains zero, the configuration should be checked.

Tools Commands

The following **tools** command dumps the initial status of the DNS IP cache (assuming that the server with IP address 192.168.1.100 can be accessed via FTP):

```
*A:Dut-C# tools dump application-assurance group 1 dns-ip-cache "nokia-cache"
url ftp://username:password@192.168.1.100/tmp/mylog.log
```

This command dumps the contents of the cache in a file. Adding a URL is optional; however, for cache tables that contain a large number of entries, it is a better option. The command will generate a text file that contains the following:

```

=====
Application-Assurance dns-ip-cache "nokia-cache"
Current Time:          "11/08/2017 12:45:32" (UTC)
  group:                1
  isa:                  3/2
  admin state:         no shutdown
  max-entries:         10
=====
ip-address  creationTime(UTC)  lastUpdated(sec)  numDNSResponses
lastMatchTime(UTC)  numTimesMatched

Total entries in-use:  0
=====

```

DNS traffic is generated to populate the cache and HTTP messages are sent to verify that the filter matched correctly. Afterward, the following command is launched to dump the cache.

```

*A:Dut-C# tools dump application-assurance group 1 dns-ip-cache "nokia-cache"
url ftp://username:password@192.168.1.100/tmp/mylog.log

=====
Application-Assurance dns-ip-cache "nokia-cache"
Current Time:          "11/13/2017 09:02:25" (UTC)
  group:                1
  isa:                  3/1
  admin state:         no shutdown
  max-entries:         10
=====
ip-address  creationTime(UTC)  lastUpdated(sec)  numDNSResponses  lastMatchTime(UTC)
numTimesMatched
8.1.17.21  "11/13/2017 08:31:35"  136                4                "11/13/2017 09:00:09"
  7

Total entries in-use:  1
=====

```

The preceding log shows that the cache entry matched four times (numDNSResponses). Traffic was sent to this address seven times (numTimesMatched).

Conclusion

This chapter describes how an operator can ensure that traffic will match the correct filter and ensure that fraudulent traffic will not be misclassified. By using the DNS IP Cache feature, operators can ensure that traffic will be correctly charged.

Application Assurance — GTP Roaming Firewall

This chapter describes Application Assurance GTP roaming firewall.

Topics in this chapter include:

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

Applicability

The configuration and information in this chapter are based on SR OS Release 20.2.R2.


Overview

Wireless network operators rely on GPRS Tunneling Protocol (GTP) for the delivery of mobile data services across the access network. However, GTP is not designed to be secure, exposing the mobile access network to attacks from both its own subscribers and its partner networks.

The Application Assurance (AA) SR OS 20.2.R2 firewall feature extends AA-Integrated Service Adapter (AA-ISA) application-level analysis to provide an in-line stateful service that integrates into a 7750 Service Router. The feature provides protection for mobile operator infrastructure against attacks from compromised mobile gateways: Serving Gateways (SGWs) or Packet data network Gateways (PGWs).

AA stateful packet filtering, combined with AA L7 classification and control, provides operators with advanced, next-generation firewall (FW) functionality. This AA stateful FW feature runs on AA-ISA and, using stateful inspection, not only inspects packets at Layers 3 to 7, but also monitors and keeps track of the connection state. [Figure 17: AA GTP roaming FW deployment](#) shows an example AA GTP roaming FW deployment.

Figure 17: AA GTP roaming FW deployment



36100

S8/Gp AA FW deployment

AA FW is deployed as a GTP FW on S8/Gp (or S5/Gn) interfaces, either as part of a 7750 SR router in the form of an AA-ISA hardware module or as a separate Virtual SR (VSR) appliance. AA FW provides operators with network security, such as:

- GTP protocol validation, which checks for anomaly attacks that involve malformed, corrupt, or spoofed traffic:
 - header length checks
 - Information Element (IE) length validation
 - invalid reserved field validation
 - reserved IE validation
 - missing mandatory IE validation
 - sequence number validation
 - Tunnel Identification (TEID) validation - blocks GTP tunnel creations that have not been signaled correctly
- PGW and SGW redirection protection
- GTP-in-GTP check

- Handover control to prevent session hijacking
- Source address (User Equipment (UE)) anti-spoofing protection
- Protection against unauthorized Public Land Mobile Network (PLMN) and/or Access Point Name (APN) access:
 - filter message-based APN, International Mobile Subscriber Identity (IMSI) prefix
- Protection against unsupported GTP message types:
 - filter messages, based on message types and/or message length
- Protection against flooding attack:
 - GTP traffic bandwidth policing, which limits the GTP bandwidth from a roaming partner SGW/PGW
 - GTP tunnel limiting, which limits the number of concurrent GTP tunnels and/or the setup rate of these tunnels from a roaming partner SGW/PGW
- Protection against IP fragmentation-based attacks:
 - drop rules for IP fragmentation of GTP messages

AA FW supports both GTPv1 and GTPv2. It is typically deployed as an L3/VPDN service. SAPs/spokes are diverted to AA for a GTP FW. L2/VPLS connectivity is supported by AA. AA transit subscribers (identified by SGW IPs) are auto-created under the parent-diverted SAPs/spokes.

UE IP address anti-spoofing

Source address spoofing is initiated by a malicious UE that hijacks (spoofs) an IP address of another UE and invokes a download from a malicious server on the Internet. After the download begins, the malicious UE exits the session. The UE under attack (receiving the download traffic) gets charged for traffic it did not solicit.

AA FW associates the GTP-c messages of the UE IP address IE with the GTP-u packets to ensure that the packets carried in the upstream have the correct source IP address (inner IP within the GTP-u tunnel). Because the UE address is negotiated within the PDP context creation handshake, any packets originating from the UE that contain a different source address are detected by AA FW and dropped.

To enable UE IP address anti-spoofing protection, the operator needs to enable "validate-source-ip-addr", as follows:

```
*A:Dut-C>config>app-assure>group>
+---gtp-filter <gtp-filter-name> [create]
|   +---gtp-tunnel-database
|   +---validate-source-ip-addr
```

GTP TEID validation

Compromised mobile gateways (GSNs) can send storms of GTP traffic with invalid GTP TEIDs to cause a denial of service (DoS) attack. By inspecting GTP-c messages, AA FW supports stateful correlation of upstream and downstream GTP flows (DstIP + TEID) of the same PDN session. AA drops packets with TEIDs that have not been negotiated correctly.

The operator can enable AA to drop GTP traffic with an invalid TEID using:

```
*A:Dut-C>config>app-assure>group>
```

```
+---gtp-filter <gtp-filter-name> [create]
|   +---gtp-tunnel-database
|   +---validate-gtp-tunnels
```


GTP anomaly prevention/sequence number checks

Protocol anomaly attacks involve malformed or corrupt packets that typically fall outside of the protocol specifications. Packets are denied by AA FW if they fail the sanity check. The following are some examples of GTP sanity checks:

- invalid GTP header length
- invalid IE length
- invalid reserved fields
- invalid sequence number
- missing mandatory IEs

Also, AA FW performs sequence number validation whereby it ensures no out-of-sequence GTP packets. By default, sequence number validation is disabled. To enable it, the following CLI command can be used:

Figure 18: CLI command



36101

GTP packets with wrong sequence numbers are dropped when validate-sequence-number is enabled.

GTP message-type filtering

AA FW performs GTP message validation, in which packets with invalid message types (that are not applicable to the roaming interfaces) are denied by the AA GTP-c inspection command:

Table 6: Denied GTP message types for roaming interface

| | GTP-u port | GTP-c port |
|----------------------------|------------|--|
| Denied GTPv1 message types | None | GTPU_PDU GTPV1_END_MARKER GTPV1_MSG_ERR_IND GTPV1-ALL-MBMS message-types GTPV1-ALL-Location management message-types |

| | GTP-u port | GTP-c port |
|----------------------------|------------|---|
| Denied GTPv2 message types | N/A | GTP_PKT_ERROR_INDICATION GTP_PKT_DNLK_DATA_FAIL_INDICATION GTP_PKT_STOP_PAGING_INDICATION GTP_PKT_CRE_INDR_TNL_REQ GTP_PKT_CRE_INDR_TNL_RSP GTP_PKT_DEL_INDR_TNL_REQ GTP_PKT_DEL_INDR_TNL_RSP GTP_PKT_RELEASE_BEARERS_REQ GTP_PKT_RELEASE_BEARERS_RSP GTP_PKT_DNLK_DATA GTP_PKT_DNLK_DATA_ACK GTP_PKT_MOD_ACCESS_BEARERS_REQ GTP_PKT_MOD_ACCESS_BEARERS_RSP |

Also, AA FW allows the operator to further restrict allowed message types (shown in the following table) by configuring GTP message type filter entries to deny (or allow) the following message types:

Table 7: Allowed GTP message types (Cat-1)

| | GTP-u port | GTP-c port |
|-----------------------------|--|--|
| Allowed GTPv1 message types | GTPV1_MSG_ECHO_REQ GTPV1_MSG_ECHO_RESP GTPV1_SUPP_EXT_HDR_NOTIF GTPV1_MSG_ERR_IND GTPV1_END_MARKER GTPU_PDU | GTPV1_MSG_ECHO_REQ GTPV1_MSG_ECHO_RESP GTPV1_SUPP_EXT_HDR_NOTIF GTPV1_MSG_VER_NOT_SUPP_IND GTPV1_MSG_PDP_CREATE_REQ GTPV1_MSG_PDP_CREATE_RESP GTPV1_MSG_PDP_UPD_REQ GTPV1_MSG_PDP_UPD_RESP GTPV1_MSG_PDP_DEL_REQ GTPV1_MSG_PDP_DEL_RESP GTPV1_MSG_NET_INIT_REQ GTPV1_MSG_NET_INIT_RESP GTPV1_MSG_MSINFO_REQ GTPV1_MSG_MSINFO_RESP |
| Allowed GTPv2 | N/A | GTP_PKT_ECHO_REQ GTP_PKT_ECHO_RSP |

| | GTP-u port | GTP-c port |
|---------------|------------|---|
| message types | | GTP_PKT_VERSION_NOT_SUPPORTED GTP_PKT_CRE_SES_REQ GTP_PKT_CRE_SES_RSP GTP_PKT_MOD_BEARER_REQ GTP_PKT_MOD_BEARER_RSP GTP_PKT_DEL_SES_REQ GTP_PKT_DEL_SES_RSP GTP_PKT_CHG_NOT_REQ GTP_PKT_CHG_NOT_RSP GTP_PKT_MOD_BEARER_CMD GTP_PKT_MOD_BEARER_FAIL_INDICATION GTP_PKT_DEL_BEARER_CMD GTP_PKT_DEL_BEARER_FAIL_INDICATION GTP_PKT_BEARER_RESOURCE_CMD GTP_PKT_BEARER_RESOURCE_FAIL_INDICATION GTP_PKT_SUSPEND_NOTIFICATION GTP_PKT_SUSPEND_ACK GTP_PKT_RESUME_NOTIFICATION GTP_PKT_RESUME_ACK GTP_PKT_CRE_BEARER_REQ GTP_PKT_CRE_BEARER_RSP GTP_PKT_UPD_BEARER_REQ GTP_PKT_UPD_BEARER_RSP GTP_PKT_DEL_BEARER_REQ GTP_PKT_DEL_BEARER_RSP GTP_PKT_TRACE_SESSION_ACTIVATION GTP_PKT_TRACE_SESSION_DEACTIVATION GTP_PKT_UPDATE_PDN_CONNECTION_SET_REQ GTP_PKT_UPDATE_PDN_CONNECTION_SET_RSP GTP_PKT_DELETE_PDN_CONNECTION_SET_REQ GTP_PKT_DELETE_PDN_CONNECTION_SET_RSP |

By default, the GTP message filter allows all GTP messages.

To configure GTPv2 message filtering, the following command is used:

```
*A:Dut-C>config>app-assure>group>
+---gtp-filter <gtp-filter-name> [create]
```

```

+---gtpc-inspection
| +---message-type-v2
| | |
| | | +---default-action {permit|deny}
| | | |
| | | +---entry <entry-id> value <gtpv2-message-value> action {permit|deny}
| | | | no entry <entry-id>

```

To configure GTPv1 message filtering, the following command is used:

```

*A:Dut-C>config>app-assure>group>
| +---message-type
| | |
| | | +---default-action {permit|deny}
| | | |
| | | +---entry <entry-id:1..255> value <gtpv1-message-value> action {permit|deny}
| | | | no entry <entry-id>

```



Note:

If the operator configures a message type that is invalid for the roaming interface to be denied, it will be dropped and counted under that filter entry (and not tagged as dropped due to "wrong-interface" in the event log). However, configuring the message-type filter to "permit" a message type that is invalid for the roaming interface will not take effect, because the packet with the specified message type will be dropped by the GTP-c protocol inspection process.

Unauthorized APN attack – APN filtering

APN filtering checks GTP-c messages to determine if a roaming subscriber is allowed to access a specified external network (/aka APN). The "create-session-request" and "create pdp request" GTP message types contain an APN IE in the header of a GTP packet. An APN IE consists of an external network ID (for example, nokia.com) and, optionally, a unique ID that identifies the operator PLMN.

APN filtering prevents malicious UEs from initiating a "create PDP/session request" flood attack toward the PGW/GGSN for invalid or disallowed APNs. The operator can configure an AA GTP filter to perform APN filtering to restrict roaming subscribers access to specific external networks.

An APN filter, an IMSI prefix, and an SGSN address pool can be used together to filter GTP packets, as follows:

```

*A:Dut-C>config>app-assure>group>
+---gtp-filter <gtp-filter-name> [create]
| +---imsi-apn-filter //NEW and all its children attributes
| | |
| | | +---default-action {permit|deny} //default permit
| | | |
| | | +---entry <entry-id: 1031..2030> create
| | | | + apn <string 0|1..32 characters>
| | | | |---no apn
| | | | + src-gsn ip-prefix-list <ip-prefix-list>
| | | | + src-gsn <ip address prefix>
| | | | |---no src-gsn
| | | | + action {permit|deny} //default permit
| | | |---no entry <entry-id>

```

By default, AA FW permits all APNs.

Unauthorized PLMN access – IMSI prefix filtering

The PLMN of a subscriber home network is identified by combining the Mobile Country Code (MCC) and Mobile Network Code (MNC). MCC-MNC is also known as the International Mobile Station Identity (IMSI) prefix. The IMSI prefix acts as a PLMN identifier.

GTP IMSI prefix filters can be configured to deny GTP incoming traffic from invalid roaming partners. Conversely, GTP IMSI prefix filters can allow only incoming traffic from those network operators that have signed roaming agreements. Any GTP packets with IMSI prefixes not matching the configured prefixes are dropped.

As shown in the [Unauthorized APN attack – APN filtering](#) section, IMSI filter entry can also be optionally combined with an SGSN/SGW IP address (or IP address prefix list) to further restrict allowed IMSI prefix traffic to specific SGSN/SGW nodes.

Unauthorized network access

An attacker, using an unauthorized GSN, can cause a DoS attack using spoofed PDP context delete messages (DoS attack) or spoofed update PDP context requests to hijack existing sessions. Such attacks can also spoof a create PDP context request to gain unlawful Internet access. Session hijacking can come from either the SGW/SGSN or the PGW/GGSN. An unauthorized GSN can hijack GTP tunnels or cause a DoS attack by intercepting another GSN and redirecting traffic to it.

Operators can use AA FW to configure pools of trusted GSN IP addresses (using AA IP-prefix-list) to stop spoofed requests from untrusted GSNs. AA IP prefix lists can be configured to model GSN groups, as follows:

```
*A:Dut-C>config>app-assure>group#
  ip-prefix-list ip-prefix-list-name [create]
    prefix ip-prefix/ip-prefix-length [name prefix-name]
```

These lists are then referenced in session filters, such that only sessions that match the lists can be "permitted", as follows:


```
*A:Dut-C>config>app-assure>group# session-filter
  default-action deny
  entry # Configure an entry in the session filter
    match
      src-ip # Configure IPs that correspond to authorized SGW/SSGN
    action
      permit
```

Configuration

AA GTP filtering functionality is enhanced in SR OS Release 20.2.R2 to include support for the AA FW feature related to the GTP roaming interface. The GTP filters are optional Application QoS Policy actions (AQP actions). AQPs have partition-level scope, which allows different FW policies to be implemented by using AA partition concepts within the same AA-ISA.

The configuration topology in [Figure 19: Configuration topology](#) shows how the VSR equipped with AA FW functionality provides protection for the S8 interfaces.

Figure 19: Configuration topology



36101

Pre-setup requirements

Configuration of a VSR router is required if a 7750 SR is not already used in the access network on the S8 interfaces. If a 7750 SR is already deployed, AA-ISA must be configured.

See the [Application Assurance — Stateful Firewall](#) chapter for basic knowledge about AA-FW functionality.

Platform-dependent configuration

VSR

For GTP FW deployment in VSR, the following configuration supports load balancing of traffic across multiple CPU cores:

```
*A:7750-1>config> config>isa>aa-grp# vm-traffic-distribution-by-teid
```

7750 hardware

GTP FW deployment in 7750 hardware is only supported by ISA2:

```
*A:7750-1>config> config>isa>aa-grp# minimum-isa-generation 2
```

Allocation of memory for stateful GTP processing

To support stateful GTP processing (for example, TEID, sequence number, and UE IP validation FW operations), the operator must configure the system to allocate sufficient memory resources, as follows:

```
*A:7750-1>config>isa>aa-grp# shared-resources
                        gtp-tunnel-database 100
```

Configuration to divert SAPs/VPRN traffic into AA-ISA

In this configuration example, one VPRN is used per wireless roaming partner network. In the example, two roaming partner networks are used for illustration. In real networks, this number is much bigger.

However, before configuring SAPs for diversion, the operator can optionally define some Application Service Option (ASO) characteristics to provide different FW policies for different roaming partners, as follows:

```
*A:7750-1>config>app-assure>group 1:1 policy
begin
app-service-options
characteristic "FW-Protection" persist-id 1 create
value "OFF" persist-id 1
value "ON" persist-id 2
default-value "OFF"
exit
characteristic "strict-FW-Protection" persist-id 2 create
value "OFF" persist-id 1
value "ON" persist-id 2
default-value "OFF"
exit
exit
commit
```

For more information about ASO configuration, see the [Application Assurance — App-Profile, ASO and Control Policies](#) chapter.

After configuring any ASO characteristics, define an application profile and transit IP policy; for example:

```
*A:7750-1>config>app-assure>group$ info
-----
policy
begin
app-profile "default" create
description "App profile that applies to the whole SAP"
divert
characteristic "FW-Protection" value "ON"
exit
app-profile "strict-FW" create
description "App profile that applies strict FW rules to the SAP"
divert
characteristic "FW-Protection" value "ON"
characteristic "strict-FW-Protection" value "ON"
exit
commit
exit
transit-ip-policy 1 create
def-app-profile "strict-FW"
detect-seen-ip
transit-auto-create
```



```
no shutdown
exit
exit
```

Traffic of these two VPRNs needs to be diverted into AA-ISA to provide firewall protection. Apply the following policies to the SAPs of the two VPRNs:

```
*A:7750-1>config#
service
customer 1 name "1" create
description "Default customer"
exit
customer 2 name "2" create
exit
vprn 100 name "100" customer 2 create
interface "to-NetA" create
exit
exit
vprn 200 name "200" customer 1 create
interface "to-site1" create
exit
exit
vprn 100 name "100" customer 2 create
description "L3 Service roaming partner 2"
route-distinguisher 100:2
interface "to-NetA" create
address 192.168.1.1/24
sap 1/2/3 create
app-profile "default"
exit
exit
exit
vprn 200 name "200" customer 1 create
description "L3 Service roaming partner 1"
route-distinguisher 200:1
interface "to-site1" create
address 192.168.2.1/24
static-arp 1.1.1.2 00:ff:02:00:00:01
sap 1/2/1 create
transit-policy ip 1
app-profile "strict-FW"
exit
exit
exit
exit
```

This configuration achieves the following.

1. Roaming traffic is diverted to AA-ISA for FW protection.
2. Customer 1 traffic will have a "strict" FW rule attribute, while customer 2 traffic will be subject to basic FW rules.
3. Within AA-ISA, the customer 1 diverted SAP is treated as a parent SAP.
 - Instead of treating the whole SAP as a single subscriber, subscribers are auto-created within this SAP, based on the IP address of the SGWs/SSGNs.

**Note:**

If the operator does not require per SGW/SSGN control (such as limiting the total bandwidth of a SGW to prevent DoS attack), the "transit IP policy" from the SAP configuration can be removed.

This will cause AA to treat the whole SAP as a single subscriber, as in the case of the customer 2 SAP.

Configuration FW events log

To configure a log that captures events related to various AA firewall actions:

```
*A:7750-1# configure application-assurance group 1:1
*A:7750-1>config>app-assure>group# event-log "FW_events_log" create
*A:7750-1>config>app-assure>group>evt-log$ buffer-type circular
*A:7750-1>config>app-assure>group>evt-log$ max-entries 100000
*A:7750-1>config>app-assure>group>evt-log$ no shutdown
*A:7750-1>config>app-assure>group>evt-log$ exit
*A:7750-1>config>app-assure>group# info
-----
---snip---
event-log "FW_events_log" create
buffer-type circular
max-entries 100000
no shutdown
exit
```



Note:

Alternatively, due to the limited size of the log and the large amount of traffic AA can handle, it is recommended that the operator use the syslog mechanism instead of local logging, as follows:

```
config>app-assure>group>event-log>syslog
```

The event log can be referenced in various FW actions that are configured later in this chapter.

```
*A:7750-1# tools dump application-assurance group 1:1 event-log "FW_event_log"
                                                    isa 1/1
=====
Application-Assurance event-log "FW_events_log"
Current Time:      "05/19/2020 19:03:58" (UTC)
  group[:partition]:  1:1
  isa:                1/1
  admin state:       no shutdown
  buffer-type:       circular
  max-entries:       100000
=====
Event-source
Action      SubType      SubName      Direction Src-ip
Dst-ip      Ip-protocol Src-port Dst-port Timestamp
"gtp filter gtp-filter-partner1 reason: filtered-gtp-message-type, teid: 0x0001d100,
MT: 36, version: 2"
  transit    "1_10.10.68.1/32"      from-sub 10.10.68.1      deny
10.10.68.3      udp          2123      2123      "05/19/2020 18:54:50"
"gtp filter gtp-filter-partner1 reason: filtered-gtp-message-type, teid: 0x00019100,
MT: 37, version: 2"
  transit    "1_10.10.68.1/32"      to-sub   10.10.68.3      deny
10.10.68.1      udp          2123      2123      "05/19/2020 18:54:55"

Total Records:    2
=====
```

The following command clears all the entries within the specified log:

```
*A:7750-1# clear application-assurance group 1:1 event-log "FW_events_log"
```

Configuration to limit total traffic from SGWs

Nokia recommends that a total limit be placed on how much bandwidth and how many flows an SGW/SGSN can generate toward the network. The exact limit values are a function of the number of end devices that are served by the roaming partner SGW/SGSN and capacity limits of the HPLMN PGW/GGSN, plus some additional margin.

In the following example, it is assumed that traffic from each roaming SGW will not exceed 1200 concurrent flows/second (serving about 200 roaming UEs) and 50 Mb/s. These need to be replaced in actual deployments with appropriate values that reflect the specific network deployment.

```
*A:7750-1>config>app-assure>group# info
-----
    policer "limit_roamingSGW_Flows" type flow-count-limit
                                     granularity subscriber create
        flow-count 1200
        gtp-traffic
    exit
    policer "limit_roamingSGW_bw" type single-bucket-bandwidth
                                     granularity subscriber create
        rate 50000
        mbs 500
    exit
-----
```

Apply the configured policers as actions from within the default subs-policy AQP entry:

```
*A:7750-1>config>app-assure>group>policy>aqp#
    entry 500 create
        description "limit per SGW flow and b/w- partner 1"
        match
            traffic-direction subscriber-to-network
            characteristic "strict-FW-Protection" eq "ON"
        exit
        action
            bandwidth-policer "limit_roamingSGW_bw"
            flow-count-limit "limit_roamingSGW_Flows"
        exit
        no shutdown
    exit
```

For GTP traffic flow count policing, it is important that "aqp-initial-lockup" is enabled:

```
*A:7750-1# configure application-assurance group 1:1 aqp-initial-lockup
```



Note:

All the preceding actions apply to the traffic direction "subscriber-to-network". These actions do not apply to traffic in the other direction (downlink), because the purpose of the AA FW is to protect the network resources from upstream traffic from compromised roaming partner SGWs.



Note:

No policers are placed for the traffic of customer 2, because its profile does not have "strict policing" enabled. A policer can be configured to limit the total bandwidth and flows from all SGWs served to the customer 1 SAP, as follows:

```
*A:7750-1>config>app-assure>group# info
-----
    policer "limit_roamingSGWs_total_Flows" type flow-count-limit
                                         granularity subscriber create
        flow-count 12000
        gtp-traffic
    exit
    policer "limit_roamingSGWs_total_bw" type single-bucket-bandwidth
                                         granularity subscriber create
        rate 500000
        mbs 5000
    exit
-----
```

Apply the configured policers as actions from within the default subs-policy AQP entry, as follows:

```
*A:7750-1>config>app-assure>group>policy>aqp#
    entry 501 create
        no shutdown
        description "limit total SGW flow and b/w- partner 2 "
        match
            traffic-direction subscriber-to-network
            characteristic "strict-FW-Protection" eq "OFF"
        exit
        action
            bandwidth-policer "limit_roamingSGWs_total_bw"
            flow-count-limit "limit_roamingSGWs_total_Flows"
        exit
    exit
```

GTP filtering – disallow traffic from unauthorized SGWs

To use GTP filtering to disallow traffic from unauthorized SGWs, perform the following steps:

1. Create AA IP lists
 2. Use AA IP lists in session filters and AQPs
 3. Reference session filters within AQPs
1. Create AA IP lists, by creating an AA IP prefix list that contains SGW IP addresses or range of addresses for each customer, as follows:

Roaming partner 1

```
*A:7750-1>config>app-assure# group 1:1
    ip-prefix-list "Roaming1_ALL_SGWs" create
        description "SGWs subnet-partner 1"
        prefix 172.16.100.0/24
    exit
    ip-prefix-list "Roaming2_ALL_SGWs" create
        description "SGWs subnet for roaming partner2"
        prefix 172.16.110.100/30
    exit
```

Roaming partner 2

```
*A:7750-1>config>app-assure>group#
  ip-prefix-list "Roaming2_ALL_SGWs" create
    description "SGWs subnet for roaming partner2"
    prefix 172.16.110.100/30
  exit
```

2. The AA IP prefix lists can be referenced and used in AA FW rules using session filters and AQPs, as follows:

```
*A:7750-1>config>app-assure>group#
  session-filter "restricted_access_partner1" create
    description "SGWs_allowed_partner1"
    default-action deny
    entry 10 create
      description "allow GTP-u from authorized subnets"
      match
        ip-protocol-num udp
        src-ip ip-prefix-list "Roaming1_ALL_SGWs"
        dst-port eq 2152
      exit
      action permit
    exit
  entry 11 create
    description "allow GTP-c from authorized subnets"
    match
      ip-protocol-num udp
      src-ip ip-prefix-list "Roaming1_ALL_SGWs"
      dst-port eq 2123
    exit
    action permit
  exit
  entry 20 create
    description "allow DNS"
    match
      ip-protocol-num *
      src-ip ip-prefix-list "Roaming1_ALL_SGWs"
      dst-port eq 53
    exit
    action permit
  exit
  session-filter "restricted_access_partner2" create
    description "SGWs_allowed_partner2"
    default-action deny event-log "FW_events_log"
    entry 10 create
      description "allow GTP-u from authorized subnets"
      match
        ip-protocol-num udp
        src-ip ip-prefix-list "Roaming2_ALL_SGWs"
        dst-port eq 2152
      exit
      action permit
    exit
  entry 11 create
    description "allow GTP-c from authorized subnets"
    match
      ip-protocol-num udp
      src-ip ip-prefix-list "Roaming2_ALL_SGWs"
      dst-port eq 2123
    exit
```

```

        action permit
    exit
    entry 20 create
        description "allow DNS"
        match
            ip-protocol-num *
            src-ip ip-prefix-list "Roaming2_ALL_SGWs"
            dst-port eq 53
        exit
        action permit
    exit
exit

```

**Note:**

Optionally, you can combine the session filter entries for the two roaming partners into a single session filter (for scale reasons). AA supports a total of 300 session filters. If there are less than 300 roaming partners, you can use a session filter per partner for customization purposes (related to, for example, IP subnets). If the number of partners is greater than the maximum number of session filters, you need to aggregate entries into a fewer number of session filters. Be aware of overlapping IP addresses from different roaming partner networks/VPRNs.

3. The configured session filters need to be referenced within AQPs, as follows:

```

*A:7750-1>config>app-assure>group>policy>aqp#
    entry 510 create
        description "apply FW rules for roaming partner 1"
        match
            traffic-direction subscriber-to-network
            characteristic "strict-FW-Protection" eq "ON"
        exit
        action
            session-filter "restricted_access_partner1"
        exit
        no shutdown
    exit
    entry 511 create
        description "apply FW rules for roaming partner 2"
        match
            traffic-direction subscriber-to-network
            characteristic "strict-FW-Protection" eq "OFF"
        exit
        action
            session-filter "restricted_access_partner2"
        exit
        no shutdown
    exit

```

Restrict downstream traffic (optional)

Operators can optionally restrict downstream traffic to specific destinations and protocols, as follows:

```

*A:7750-1>config>app-assure>group#
    session-filter "restricted_downstream_traffic_1" create
        description "allow only traffic to only signed up partners"
        default-action deny
    entry 10 create
        description "allow GTP-u from authorized subnets"
        match

```

```

        ip-protocol-num udp
        dst-ip ip-prefix-list "Roaming1_ALL_SGWs"
        dst-port eq 2152
    exit
    action permit
exit
entry 11 create
description "allow GTP-c to authorized subnets"
match
    ip-protocol-num udp
    dst-ip ip-prefix-list "Roaming1_ALL_SGWs"
    dst-port eq 2123
    exit
    action permit
exit
entry 20 create
description "allow DNS"
match
    ip-protocol-num *
    dst-ip ip-prefix-list "Roaming1_ALL_SGWs"
    exit
    action permit
exit
exit
session-filter "restricted_downstream_partner2" create
description "SGWs_allowed_partner2"
default-action deny event-log "FW_events_log"
entry 10 create
description "allow GTP-u to authorized subnets"
match
    ip-protocol-num udp
    dst-ip ip-prefix-list "Roaming2_ALL_SGWs"
    dst-port eq 2152
    exit
    action permit
exit
entry 11 create
description "allow GTP-c to authorized subnets"
match
    ip-protocol-num udp
    dst-ip ip-prefix-list "Roaming2_ALL_SGWs"
    dst-port eq 2123
    exit
    action permit
exit
exit

```

**Note:**

The preceding configuration provides the most flexibility and allows IP addresses to overlap between different partner networks. However, it comes at the cost of creating separate session filters for each partner. If IP addresses do not overlap, a single session filter is sufficient.

The session filters need to be referenced from AQP, as follows:

```

*A:7750-1>config>app-assure>group>policy>aqp#
    entry 514 create
        description "apply FW rules for roaming partner 1"
        match
            traffic-direction network-to-subscriber
            characteristic "strict-FW-Protection" eq "ON"
        exit
        action
    
```

```
        session-filter "restricted_downstream_partner1"
        exit
        no shutdown
    exit
    entry 513 create
        description "apply FW rules for roaming partner 2"
        match
            traffic-direction network-to-subscriber
            characteristic "strict-FW-Protection" eq "OFF"
        exit
        action
            session-filter "restricted_downstream_partner2"
        exit
        no shutdown
    exit
```

Configuration to protect against malformed packets

It is always recommended in FW deployments that overload-drop, error-drop, and fragment-drop are enabled within the default sub-policy, as follows:

```
*A:7750-1>config>app-assure>group>policy>aqp#
    entry 50 create
        description "drop error and fragmented packets"
        action
            overload-drop event-log "FW_events_log"
            error-drop event-log "FW_events_log"
            fragment-drop all event-log "FW_events_log"
        exit
        no shutdown
    exit
```



Note:

- The overload-drop action ensures that AA-ISA, if it gets overloaded, drops the excess traffic instead of cutting it through without applying FW rules.
- The error-drop action ensures that AA-ISA drops malformed IP packets.
- The fragment-drop all action allows the operator to drop all fragmented traffic, drop out-of-order fragments only, or allow fragments through. Because many network DoS attacks use IP fragmentation to initiate attacks, allowing fragments through is not recommended for firewall deployments. As a minimum, if fragmentation is used, the operator is recommended to configure AA to drop out-of-order fragmented packets.
- The preceding actions are applied to all traffic. Therefore, there are no AQP match conditions configured.

Plausibility of GTP messages and GTP message validation

To protect the network from malformed GTP packets and associated attacks as described in the overview section, a GTP filter needs to be created and referenced from an AQP entry.

1. Configure the GTP filter object to:
 - a. Enable GTP-c inspection so that the FW:
 - i. Ensures the correct IE states and order

- ii. Discards any GTP packet that contains an invalid message type for S5/S8 (Gn/Gp) interface
- b. Enable sequence number checking for GTP-c traffic (for partner 1 traffic)
- c. Enable the GTP filter to check and drop errored GTP packets (anomalies)
- d. Enable GTP message length checking (to minimize exposure to code injection attacks). The maximum is set here (for example, 1250 bytes). The value is operator dependent, and should be replaced with the figure used by the operator.
- e. Drop GTP-in-GTP encapsulated packets

```
*A:7750-1>config>app-assure>group>gtp#
  event-log "FW_events_log"
  gtpc-inspection
  gtp-filter "gtp-filter-partner1" create
    description "gtp-filter for partner 1"
    max-payload-length 1250
    event-log "FW_events_log" action deny
    gtp-in-gtp deny
    gtp-tunnel-database
    validate-sequence-number
  exit
exit
gtp-filter "gtp-filter-partner2" create
  description "gtp-filter for partner 2"
  max-payload-length 1250
  event-log "FW_events_log" action deny
  gtp-in-gtp deny
no shutdown
```

2. The configured GTP filters need to be referenced from AQP entries, as follows:

```
*A:7750-1>config>app-assure>group>policy>aqp#
  entry 512 create
    description "apply SGW GTP filter rules"
    match
      characteristic "strict-FW-Protection" eq "ON"
    exit
    action
      gtp-filter "gtp-filter-partner1"
    exit
    no shutdown
  exit
  entry 513 create
    description "apply SGW GTP filter rules"
    match
      characteristic "strict-FW-Protection" eq "OFF"
    exit
    action
      gtp-filter "gtp-filter-partner2"
    exit
    no shutdown
  exit
```

Filtering of GTP message types

In this configuration, traffic from partner 2 is considered "safe/trusted". Therefore, unlike traffic from partner 1, no additional GTP message-type filtering is applied to it, beyond the GTP Cat-1 (see [Table 7: Allowed GTP message types \(Cat-1\)](#)) message filtering applied as a result of enabling GTP-c inspection.

For roaming partner 1 traffic, the GTP filter is configured to block some Cat-1 optional message types (GTPv1 and GTPv2):

- GTPv2: Trance session activation/deactivation (this is optional for S8)
- GTPv1: Allows only the message types used by GTP-u and blocks GTPv1 message types used by GTP-c



Note:

By configuring GTP-c inspection, only Cat-1 message types (see [Table 7: Allowed GTP message types \(Cat-1\)](#)) are allowed and all others are denied. Therefore, there is little to no need for additional GTP message filtering configuration.

The GTP filter is configured as follows:

```
*A:7750-1>config>app-assure>group>gtp#
    event-log "FW_events_log" action deny
    gtpc-inspection
    gtp-filter "gtp-filter-partner1" create
      description "gtp-filter for partner 1"
      max-payload-length 1250
      event-log "FW_events_log" action deny
      message-type
        default-action deny
        entry 1 value "echo-request" action permit
        entry 2 value "echo-response" action permit
        entry 3 value "error-indication" action permit
        entry 4 value "supported-extension-headers-notification"
                                                    action permit
        entry 5 value "end-marker" action permit
        entry 6 value "g-pdu" action permit
      exit
    message-type-gtpv2
      default-action permit
      entry 524 value "trace-session-activation" action deny
      entry 525 value "trace-session-deactivation" action deny
    exit
    gtp-in-gtp deny
    imsi-apn-filter
      default-action deny
      entry 1031 create
        apn ANY_APN
        mcc-mnc-prefix 161379
        action permit
      exit
    exit
    gtp-tunnel-database
      validate-sequence-number
    exit
  exit
```

TEID validation

For roaming partner 1, to protect the network resources from spoofed TEIDs, the FW is recommended to verify that the TEIDs used in the GTP-u traffic are valid (that is, correctly negotiated via GTP-c), as follows:

```
*A:7750-1>config>app-assure>group>gtp#
    gtp-filter "gtp-filter-partner1"
      gtp-tunnel-database
```

```
        validate-gtp-tunnels
    exit
exit
```

Since roaming partner 2 network is trusted, no TEID validation is needed.

UE IP address anti-spoofing

It is a good practice to protect the network against UEs spoofing a different IP address, as follows:

```
*A:7750-1>config>app-assure>group>gtp#
    gtp-filter "gtp-filter-partner1"
        gtp-tunnel-database
            validate-source-ip-addr
        exit
    exit
```

This example applied to partner 1 traffic. Validation of source IP requires the use of a GTP tunnel database.

APN and IMSI filtering

In this example, only Home-Routed (HR) traffic from partner 1 is allowed, regardless of the APN. The rest is denied. This is achieved by configuring an IMSI prefix ("1613797") that corresponds to the Home network.

For roaming partner 2, MVNO traffic is allowed as well as HR traffic. This MVNO traffic (specific IMSI prefix = "1613400" in this example) is only allowed to attach to the mvnoguest.com APN, as follows:

```
*A:7750-1>config>app-assure>group>gtp#
    gtp-filter "gtp-filter-partner1"
        imsi-apn-filter
            default-action deny
            entry 1031 create
                apn ANY_APN
                mcc-mnc-prefix 161379
            action permit
        exit
    exit
exit
gtp-filter "gtp-filter-partner2"
    imsi-apn-filter
        default-action deny
        entry 1040 create
            apn mvnoguest.com$
            mcc-mnc-prefix 161340
        action permit
    exit
    entry 1041 create
        apn ANY_APN
        mcc-mnc-prefix 161379
    action permit
    exit
exit
exit
```

Limiting concurrent session creations

To further lower the risk of DoS attacks using massive amounts of session/PDN create messages, it is recommended that the operator configure the maximum concurrent number of endpoints (TIEDs) that an SGW can create.

In this example, the limit that is configured in the GTP filter corresponds to the maximum concurrent TEIDs that can be created by any SGW IP address, as follows:

```
*A:7750-1>config>app-assure>group>gtp#
    gtp-filter "gtp-filter-partner1"
    gtp-tunnel-database
        default-tunnel-endpoint-limit 400
    exit
```

Configuring FW statistics

To gain visibility into the traffic passing through the FW and the FW actions taken, it is highly recommended to enable "deny-admit" statistics, as follows:

```
A:7750-1>config>#
  log
    file-id 5
    location cf3:
  exit
  accounting-policy 5
    description "LogFileforAAFirewallAccounting"
    record aa-admit-deny
    collection-interval 10
    to file 5
    no shutdown
  exit
exit
```

```
A:7750-1>config>app-assure>group#
  statistics
    aa-admit-deny
      accounting-policy 5
      collect-stats
      gtp-filter-stats
      session-filter-stats
      policer-stats-resources
      policer-stats
    exit
  protocol
    shutdown
  exit
```

Configuring threshold crossing alerts

As well as admit-deny statistics, the operator can optionally enable the FW to generate Threshold Crossing Alerts (TCAs) against the collected statistics.

**Note:**

NSP also supports TCAs. The operator has a choice to enable TCAs on both the FW and/or NSP. The advantage of TCAs generated directly from the FW is that they tend to be more real-time relative to NSP TCAs. However, NSP supports a larger TCA scale than the FW.

The operator needs to set the low- and high-water marks according to the conditions of their networks. The following values are for illustration purposes only.

```
A:7750-1> config>app-assure>group>stats#
  threshold-crossing-alert
    gtp-sanity-drop direction from-sub create
      high-wmark 100 low-wmark 60
    exit
  threshold-crossing-alert
    gtp-sanity-drop direction from-sub create
      high-wmark 100 low-wmark 60
    exit
  gtp-filter "gtp-filter-partner1"
    validate-gtp-tunnels direction from-sub create
      high-wmark 100 low-wmark 60
    exit
    validate-sequence-number direction from-sub create
      high-wmark 100 low-wmark 60
    exit
    validate-src-ip-addr direction from-sub create
      high-wmark 100 low-wmark 60
    exit
    missing-mandatory-ie direction from-sub create
      high-wmark 100 low-wmark 60
    exit
    tunnel-resource-limit direction from-sub create
      high-wmark 100 low-wmark 60
    exit
    tunnel-endpoint-limit direction from-sub create
      high-wmark 100 low-wmark 60
    exit
  message-type
    default-action direction from-sub create
      high-wmark 100 low-wmark 60
    exit
    header-sanity direction from-sub create
      high-wmark 100 low-wmark 60
    exit
  exit
  message-type-gtpv2
    entry 524 direction from-sub create
      high-wmark 100 low-wmark 60
    exit
    entry 525 direction from-sub create
      high-wmark 100 low-wmark 60
    exit
  exit
  imsi-apn-filter
    default-action direction from-sub create
      high-wmark 100 low-wmark 60
    exit
  exit
  exit
  gtp-filter "gtp-filter-partner2"
    missing-mandatory-ie direction from-sub create
      high-wmark 100 low-wmark 60
    exit
```

```
tunnel-resource-limit direction from-sub create
    high-wmark 100 low-wmark 60
exit
tunnel-endpoint-limit direction from-sub create
    high-wmark 100 low-wmark 60
exit
imsi-apn-filter
    default-action direction from-sub create
        high-wmark 100 low-wmark 60
    exit
exit
exit
exit
```

Configuring GTP and GTP-c applications

By configuring AA app-filters to define GTP-u and GTP-c applications, the operator can gain further visibility into the volume of traffic of these applications, as follows:

```
A:7750-1>config> config>app-assure>group>
policy
    begin
    application "GTP_c" create
    exit
    application "GTP_other" create
    exit
    application "GTP_u" create
    exit
    app-filter
        entry 40000 create
            protocol eq "gtp"
            server-port eq 2152
            application "GTP_u"
            no shutdown
        exit
        entry 40010 create
            protocol eq "gtp"
            server-port eq 2123
            application "GTP_c"
            no shutdown
        exit
        entry 40020 create
            protocol eq "gtp"
            application "GTP_other"
            no shutdown
        exit
    exit
exit
commit
```

The export and display of statistics related to these applications use standard AA per application per partition or per application per subscriber XML records and "show" routines.

Relevant debug routines

CLI show routines

```
*A:7750-1>config>app-assure>group>gtp>gtp-fltr# show application-assurance
group 1:1 session-filter "restricted_access_partner1"
```

```
=====
AA Session Filter Instance "restricted_access_partner1"
=====
```

```
Description      : SGWs_allowed_partner1
Default Action   : deny
  Event Log      : (Not Specified)
AQP Entries      :
  510
```

```
-----
Filter Match Criteria
-----
```

```
Entry           : 10
Description      : allow GTP-u from authorized subnets
IP Protocol      : udp
Source IP List   : Roaming1_ALL_SGWs
Dest Port        : eq 2152
Action           : permit
  Event Log      : (Not Specified)
Hits             : 1 flows
```

```
Entry           : 11
Description      : allow GTP-c from authorized subnets
IP Protocol      : udp
Source IP List   : Roaming1_ALL_SGWs
Dest Port        : eq 2123
Action           : permit
  Event Log      : (Not Specified)
Hits             : 2 flows
```

```
Entry           : 20
Description      : allow DNS
IP Protocol      : *
Source IP List   : Roaming1_ALL_SGWs
Dest Port        : eq 53
Action           : permit
  Event Log      : (Not Specified)
Hits             : 0 flows
```

```
-----
No. of entries   : 3
=====
```

```
*A:7750-1>show application-assurance group 1:1 gtp gtp-filter "gtp-filter-partner1"
```

```
=====
Application Assurance Group 1:1 GTP Filter "gtp-filter-partner1"
=====
```

```
Description      : gtp-filter for partner 1
Maximum payload length : 1250
Event log         : FW_events_log
Event log action  : deny
Default action    : deny
Default GTPv2 action : permit
Default IMSI-APN action : deny
GTP in GTP action : deny
Validate GTP tunnels : enabled
Validate sequence number : enabled
Validate source IP address : enabled
GTP tunnel endpoint limit : 400
Configured messages : 6
Configured GTPv2 messages : 2
```

```

Configured IMSI-APN filters : 1
Packets arrived             : 18
Packets denied
  Payload length            : 0
  Message type              : 0
  GTPv2 message type       : 2
  IMSI-APN filter          : 0
  Mandatory header         : 0
  Extension header         : 0
  Information element       : 0
  Invalid TEID              : 0
  Invalid sequence number   : 0
  Invalid source IP address : 0
  Missing mandatory IE     : 0
  GTP in GTP               : 0
  No tunnel resource       : 0
  Tunnel endpoint limit    : 0
Packets permitted          : 16
=====
    
```

```
*A:7750-1>show application-assurance group 1:1 gtp
```

```
=====
Application Assurance Group 1:1 GTP
=====
```

```

Admin status      : Up
Event log         : FW_events_log
Event log action  : deny
Mode              : filtering
GTP-C inspection : Enabled
    
```

```

-----
GTP Statistics                                     sub-to-net          net-to-sub
-----
Incoming packets                                 9                   9
Packets denied
  UDP packet length                             0                   0
  GTP message length                           0                   0
  GTP version                                  0                   0
-----
Packets permitted                               9                   9
-----
    
```

```

-----
GTP Policing Statistics                           sub-to-net          net-to-sub
-----
Packets arrived                                 9                   9
Packets denied
  gtp-traffic flow-count policer               0                   0
  Other                                         0                   0
-----
Packets permitted                               9                   9
-----
    
```

```

-----
GTP Filter Statistics                             sub-to-net          net-to-sub
-----
Packets arrived                                 9                   9
Packets denied                                  1                   1
Packets permitted
  gtp-filter                                    8                   8
  no gtp-filter                                0                   0
-----
Total GTP packets permitted                     8                   8
    
```



```

=====
*A:7750-1>show application-assurance group 1 aa-sub-list
=====
Application-Assurance Subscriber List for Group 1
=====
type      aa-sub                ISA      App-Profile      divert
          assigned
-----
group 1:1
-----
sap       1/2/1                 1/1      strict-FW        Yes
sap       1/2/3                 1/1      default          Yes
transit   1_10.10.68.1/32      1/1      strict-FW        Yes
-----
Number of aa-subs found in group 1:1      : 3
Total number of aa-subs found             : 3
=====

```

CLI tools dump routines

```

*A:7750-1>tools dump application-assurance group 1:1 flow-record-search isa 1/1
=====
Application-Assurance flow record search
Search Start Time:    "05/19/2020 19:04:37" (UTC)
Search Criteria:
  group[:partition]:  1:1
  isa:                 1/1
  protocol name:      none specified
  application name:   none specified
  app-group name:     none specified
  flow-status:        none specified
  start-flowId:       none specified
  classified:         none specified
  server-ip:          none specified
  server-port:        none specified
  client-ip:          none specified
  bytes-tx:           none specified
  flow-duration:     none specified
  max-count:          none specified
  flow-modified:     none specified
  search-type:        default
=====
FlowId   Init  Src-ip      Dst-ip      Ip-prot  Src-prt  Dst-prt
Protocol
Pkts-tx  Bytes-tx   Application  Pkts-disc  Bytes-disc
Time-ofp(UTC)  Time-olp(UTC)
2        no    10.10.68.3  10.10.68.1  udp      2123     2123
"gtp"
  1          46          0          0
"05/19/2020 18:54:50" "05/19/2020 18:54:50"
3        yes   10.10.68.1  10.10.68.3  udp      2123     2123
"gtp"
  2          359         0          0
"05/19/2020 18:54:50" "05/19/2020 18:54:50"
4        yes   10.10.68.3  10.10.68.1  udp      2123     2123
"gtp"
  2          326         1          51
"05/19/2020 18:54:50" "05/19/2020 18:54:55"

```

```

7      yes  10.10.68.1  10.10.68.3  udp  63760  2152
  "gtp"
    5      500      "GTP_u"      0      0
    "05/19/2020 18:54:50" "05/19/2020 18:54:50"
8      yes  10.10.68.3  10.10.68.1  udp  64784  2152
  "gtp"
    5      500      "GTP_u"      0      0
    "05/19/2020 18:54:50" "05/19/2020 18:54:50"
11     yes  10.10.68.1  10.10.68.3  udp  2123   2123
  "gtp"
    1      136      "GTP_c"      1      91
    "05/19/2020 18:54:50" "05/19/2020 18:54:50"
SEARCH COMPLETED.
Search End Time:    "05/19/2020 19:04:37" (UTC)
Total Records:     6
=====

```

```
*A:7750-1>tools dump application-assurance group 1:1 admit-deny-stats
```

```
=====
Application-Assurance Group 1:1 Admit-Deny Statistics
=====
```

```
-----
Admitted Sub-To-Net   Denied Sub-To-Net   Admitted Net-To-Sub   Denied Net-To-Sub
Packet Validation Statistics
(Packets)             (Packets)           (Packets)             (Packets)
-----
Error
0                     0                   0                     0
Fragments: Out-Of-Order
0                     0                   0                     0
Fragments: All
0                     0                   0                     0
Overload
N/A                   0                   N/A                   0
GTP Sanity
9                     0                   9                     0
-----
```

```
-----
Admitted Sub-To-Net   Denied Sub-To-Net   Admitted Net-To-Sub   Denied Net-To-Sub
GTP Filter Statistics
(Packets)             (Packets)           (Packets)             (Packets)
-----
GTP Filter: gtp-filter-partner1
Entry: 1  echo-request
0         0                   0                     0
Entry: 2  echo-response
0         0                   0                     0
Entry: 3  error-indication
0         0                   0                     0
Entry: 4  supported-extension-headers-notification
0         0                   0                     0
Entry: 5  end-marker
0         0                   0                     0
Entry: 6  g-pdu
5         0                   5                     0
Message Type Default Action

```

| | | | |
|-----------------------------------|----------------------------|---------------------|-------------------|
| 0 | 0 | 0 | 0 |
| GTPv2 Entry: 524 | trace-session-activation | | |
| 0 | 1 | 0 | 0 |
| GTPv2 Entry: 525 | trace-session-deactivation | | |
| 0 | 0 | 0 | 1 |
| GTPv2 Message Type Default Action | | | |
| 3 | 0 | 3 | 0 |
| IMSI-APN Entry: 1031 | | | |
| N/A | 0 | N/A | 0 |
| IMSI-APN Filter Default Action | | | |
| N/A | 0 | N/A | 0 |
| Max Payload Length | | | |
| N/A | 0 | N/A | 0 |
| Message Type Header Sanity | | | |
| N/A | 0 | N/A | 0 |
| Invalid TEID | | | |
| N/A | 0 | N/A | 0 |
| Invalid Sequence Number | | | |
| N/A | 0 | N/A | 0 |
| Invalid Source IP Address | | | |
| N/A | 0 | N/A | 0 |
| Missing Mandatory IEs | | | |
| N/A | 0 | N/A | 0 |
| GTP in GTP Action | | | |
| N/A | 0 | N/A | 0 |
| GTP Tunnel DB Resource | | | |
| N/A | 0 | N/A | 0 |
| Tunnel Endpoint Limit | | | |
| N/A | 0 | N/A | 0 |
| GTP Filter: gtp-filter-partner2 | | | |
| Message Type Default Action | | | |
| 0 | 0 | 0 | 0 |
| GTPv2 Message Type Default Action | | | |
| 0 | 0 | 0 | 0 |
| IMSI-APN Entry: 1040 | | | |
| N/A | 0 | N/A | 0 |
| IMSI-APN Entry: 1041 | | | |
| N/A | 0 | N/A | 0 |
| IMSI-APN Filter Default Action | | | |
| N/A | 0 | N/A | 0 |
| Max Payload Length | | | |
| N/A | 0 | N/A | 0 |
| Message Type Header Sanity | | | |
| N/A | 0 | N/A | 0 |
| Invalid TEID | | | |
| N/A | 0 | N/A | 0 |
| Invalid Sequence Number | | | |
| N/A | 0 | N/A | 0 |
| Invalid Source IP Address | | | |
| N/A | 0 | N/A | 0 |
| Missing Mandatory IEs | | | |
| N/A | 0 | N/A | 0 |
| GTP in GTP Action | | | |
| N/A | 0 | N/A | 0 |
| GTP Tunnel DB Resource | | | |
| N/A | 0 | N/A | 0 |
| Tunnel Endpoint Limit | | | |
| N/A | 0 | N/A | 0 |
| ----- | | | |
| ----- | | | |
| Admitted Sub-To-Net | Denied Sub-To-Net | Admitted Net-To-Sub | Denied Net-To-Sub |
| Session Filter Statistics | | | |
| (Sessions) | (Packets) | (Sessions) | (Packets) |

```

-----
Session Filter: restricted_access_partner1
Entry: 10
1          0          0          0
Entry: 11
2          0          0          0
Entry: 20
0          0          0          0
Default Action
0          0          0          0
Session Filter: restricted_access_partner2
Entry: 10
0          0          0          0
Entry: 11
0          0          0          0
Entry: 20
0          0          0          0
Default Action
0          0          0          0
Session Filter: restricted_downstream_partner1
Entry: 10
0          0          1          0
Entry: 11
0          0          1          0
Entry: 20
0          0          0          0
Default Action
0          0          0          0
Session Filter: restricted_downstream_partner2
Entry: 10
0          0          0          0
Entry: 11
0          0          0          0
Default Action
0          0          0          0
-----

Admitted Sub-To-Net   Denied Sub-To-Net   Admitted Net-To-Sub   Denied Net-To-Sub
Flow Policer Statistics
(Flows)              (Flows)              (Flows)              (Flows)
-----

Subscriber Flow Count Policers
  limit_roamingSGW_Flows
0                    0                    0                    0
  limit_roamingSGWs_total_Flows
0                    0                    0                    0
-----

```

Conclusion

A 3GPP roaming interface using GTP presents a security risk to mobile access networks. The AA GTP stateful firewall protects the network infrastructure from untrusted roaming partner networks.

Application Assurance — HTTP and HTTPS Redirect

This chapter provides information about Application Assurance - HTTP and HTTPS Redirect.

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.R5.

Overview

This chapter is intended for network architects and engineers working with Application Assurance (AA). The user should have a basic understanding of AA policies (application service options, application filters, applications, application QoS policies), HTTP, and DNS.

Nokia recommends using the AppDB when configuring Application QoS Policies (AQP) for traffic control policies. The AppDB is a default configuration file that can be obtained through Nokia's support organization, and contains information to classify all supported applications, including all relevant application ("app-") groups, applications, and app-filters.

Operators have scenarios where they need to allow some restricted web access to subscribers (while generally blocking all other Internet access). The allowed sites may be authentication sites, such as payment sites. Alternatively, the sites may be portals owned by the operators, which may offer app-store services or user account management.

There are several ways that HTTP redirection can be implemented in an AA policy:

- HTTP policy redirect with (or without) allowed URLs: redirects HTTP traffic to a web portal based on a policy decision
- HTTPS captive redirect with (or without) DNS IP cache allow-list: redirects all HTTP/HTTPS traffic to a web portal, except for traffic toward domains configured in a DNS IP cache
- URL filter triggered HTTP redirect or HTTPS captive redirect (SR OS Release 16.0): redirects all HTTP/HTTPS traffic, based on a URL filter policy decision, where on block decision, the request can be redirected to a message portal

AA is a preferred choice for the operator to configure HTTP redirect. As opposed to CPM-based redirect (offered as part of subscriber management), AA can provide high performance with no system (or ISA) limits on the number of HTTP sessions per second that can be redirected. ISA2 can support more than 5000 redirects per second. CPM-based redirect is quite limited and, in many customer deployments, not able to support the required redirect rates (CPM-based redirect can support a maximum of 250 simultaneous connections). Even if only a simple subscriber redirect is needed (without URL allow-lists or captive redirect), AA-based redirect should be preferred over CPM-based redirect.

The following sections describe the different methods of implementing HTTP and HTTPS redirection with URL allow-lists, why and when to use each one, provide the configuration, and describe how to troubleshoot and correct common errors.

HTTP redirection based on URL filtering is not included in this chapter. For URL filtering, see chapter [Application Assurance — Local URL List Filtering](#).

Configuration

The following two use cases are described in this section:

- [AA HTTP redirect with URL allow-lists](#)
- [HTTPS redirect using session filter captive redirect](#)

AA HTTP redirect with URL allow-lists

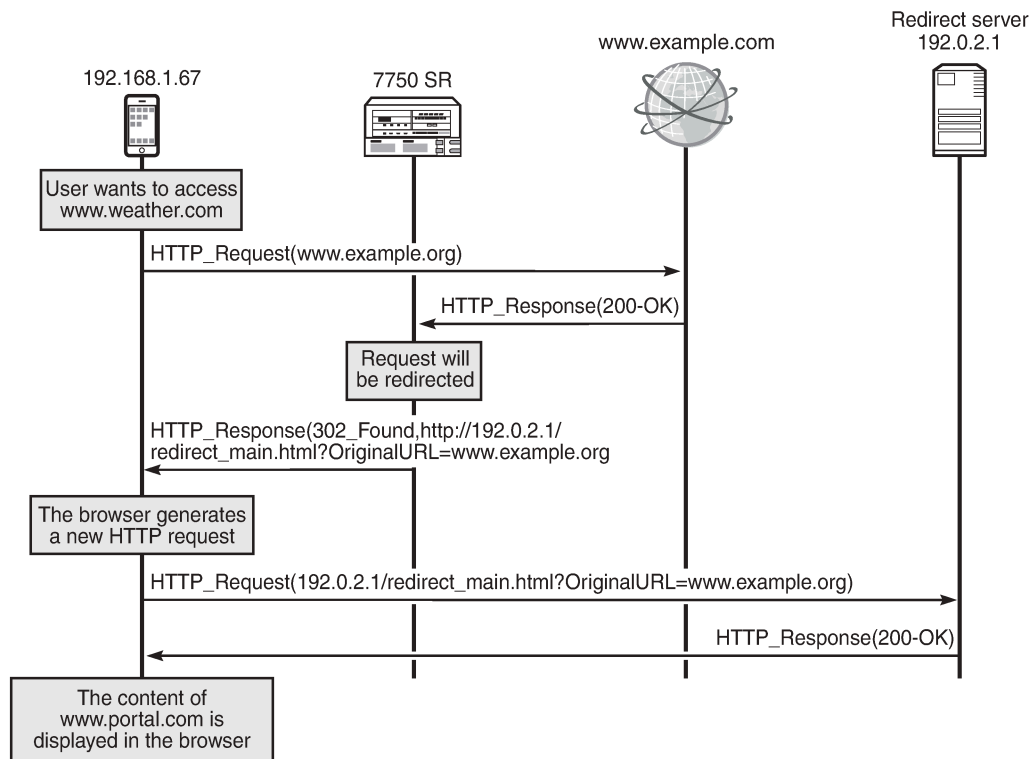
Operators have several use cases where they need to redirect a subscriber to a specific web page. Redirection may be the result of a policy (for example, the subscription does not allow certain content to be viewed) or for informative reasons (for example, the subscriber is not authenticated).

Operators may still need to allow the subscriber to access certain sites (for example, portals provided by the operator).

AA has the capability to perform HTTP redirection by using the "HTTP 302 Found" response code. This works as follows:

The subscriber generates an HTTP request to view a specific page. When AA receives the HTTP response, it classifies the flow and performs the redirect action by modifying that response. The web server replies with a "200 OK" response containing the page content, and AA modifies the HTTP response to a "302 Redirect" and sends it to the subscriber. The response contains the new URL that the subscriber should be redirected to. The subscriber's browser, upon receiving the response, automatically generates a new HTTP request for the URL received, as shown in [Figure 20: HTTP redirect](#).

Figure 20: HTTP redirect



27928

The redirect is shown in more detail in the following HTTP traces:

1. The browser sends an HTTP request for www.example.org. The HTTP packet contains the following information:

```

Hypertext Transfer Protocol
GET / HTTP/1.1\r\n
... ..
Host: www.example.org\r\n
... ..
    
```

- "GET" is followed by the path requested. In this example, the subscriber requested the root path (therefore "/").
 - "Host" contains the hostname (domain) requested.
2. The web server replies with a "200 OK" message.
 3. The SR OS node receives the server response and classifies the flow. The router determines that it will redirect the request. It replies with a message indicating the new URL along with the original URL. The HTTP packet contains the following information:

```

Hypertext Transfer Protocol
HTTP/1.1 302 Found\r\n
Location: http://192.0.2.1/redirect_main.html?originalURL=www.example.org/\r\n
... ..
    
```


- The HTTP response status code "302" indicates that the web browser should generate a new HTTP request.
 - "Location" contains the URL that the subscriber should be redirected to.
4. The subscriber's browser receives the redirect request and automatically generates a new HTTP request for the new URL. The HTTP packet sent by the browser is as follows:

```
Hypertext Transfer Protocol
GET /redirect_main.html?originalURL=www.example.org/ HTTP/1.1
Host: 192.0.2.1\r\n
...
```

- As in step 1, the GET request contains the path.
- Similarly, "Host" contains the hostname (domain) requested. "Host" can contain either a domain or an IP address.

Figure 21: Example setup shows the example setup containing a 7750 SR with ISA-AA. The setup requires a web client, a web server (which will be the redirect server), and Internet access through the router. A default configuration has been loaded from the AppDB.

Figure 21: Example setup



27929

The following sections describe the configuration of the following objects:

- The redirect policy to configure the redirect URL, the information to be appended to the URL, plus how to handle non-HTTP traffic
- An application and app-filter for the redirect portal
- A charging group for the URLs in the allow-list
- An app-filter to define the URL in the allow-list
- An ASO that will be used as a policy decision to perform redirection
- An AQP to perform redirection

Creating the redirect policy

The redirect policy is configured as follows:

```
*A:Dut-C>config>app-assure>group# info
-----
http-redirect "policy-redirect" create
```



```

description "redirects certain http traffic to a portal"
template 5
tcp-client-reset
redirect-url "http://192.0.2.1/redirect_main.html?OriginalURL=$URL"
no shutdown
exit
-----

```

The preceding configuration defines:

- The template to be used ("template 5" means that macro substitution will be used). This allows the operator to append additional information to the redirect information, which can be used by the redirect portal to offer more details about why the subscriber was redirected or to offer more personalized services.
- Traffic that cannot be redirected will be TCP reset. This provides immediate feedback to the subscriber that the service requested is not available. The client application will not have to wait for the TCP session to expire.
- The redirect URL. The original URL will be appended.

Creating the application and app-filter for the redirect portal

The application "Captive Portal" is created and used in the app-filter that contains the redirect server address, as follows:

```

*A:Dut-C>config>app-assure>group# info
-----
policy
begin
application "Captive Portal" create
description "HTTP Redirect portal"
app-group "Web"
exit

app-filter
entry 200 create
server-address eq 192.0.2.1/32
application "Captive Portal"
no shutdown
exit
exit
commit
exit
-----

```

The preceding configuration defines:

- An application called "Captive Portal"
- An app-filter using destination address 192.0.2.1 (the address of the redirect server) and linking it to the application "Captive Portal"

Creating a charging group for the URLs in the allow-list

The following configuration defines a charging group and an application. The charging group created is assigned to the application.

```
*A:Dut-C>config>app-assure>group# info
-----
policy
begin
charging-group "allow-list" create
description "Charging group for redirect portal and URLs in the allow-list"
exit

application "Captive Portal" create
description "HTTP Redirect Portal"
app-group "Web"
charging-group "allow-list"
exit
commit
-----
```

Nokia recommends defining a charging group for the URLs in the allow-list. Even though an application (or several applications) can be configured, defining a charging group provides a simpler approach; AA application group mappings are not affected, and it is easier to configure AQP. The redirect pages are usually free-of-charge, so configuring a charging group is the preferred way to aggregate these URLs.

As well as the redirect portal, DNS traffic should also be assigned to the "allow-list" charging group. Before performing an HTTP request, the subscriber's browser will first make a DNS request to resolve the hostname. Therefore, DNS traffic should be allowed.

This example assumes that the operator has loaded an AppDB configuration in the system. The AppDB contains the configuration (app-group, application, app-filter) for DNS traffic. The default application for DNS is as follows:

```
application "DNS" create
description "Domain Name System"
app-group "Network Infrastructure"
exit
```

This preceding configuration should be modified as follows:

```
*A:Dut-C>config>app-assure>group# info
-----
policy
begin
application "DNS" create
description "Domain Name System"
app-group "Network Infrastructure"
charging-group "allow-list"
exit
commit
-----
```

Creating an application for the URLs in the allow-list

The following configuration creates a new application that contains all the allowed URLs. The application is assigned to the "allow-list" charging group.

```
*A:Dut-C>config>app-assure>group# info
-----
policy
begin
  application "URL-allow-list" create
  description "URLs to be allowed"
  app-group "Web"
  charging-group "allow-list"
exit
commit
-----
```

Creating app-filters for the allowed URLs

The following configuration creates an app-filter for the URLs in the allow-list. The configuration contains the URL and assigns it to an application.

```
*A:Dut-C>config>app-assure>group# info
-----
policy
begin
  app-filter
  entry 61001 create
  description "dummy allowed domain"
  expression 1 http-host eq "/*.domain.com$"
  application "URL-allow-list"
  no shutdown
exit
commit
-----
```

The operator can extend the preceding configuration and add more URLs to the allow-list.

Configuring the ASO and AQP

The following configuration defines:

- An ASO that will be used as a policy criterion to decide whether redirection will occur. The ASO is called "redirect" and has two values: "no" and "yes". The default value will be "no" (do not redirect).
- An AQP to perform redirection. The match criteria will be the preceding ASO and the value of the charging group. If the match criteria are met, the action will be to drop and redirect the flow using the "policy-redirect" defined.

```
*A:Dut-C>config>app-assure>group# info
-----
policy
begin
  app-service-options
  characteristic "redirect" create
  value "no"
  value "yes"
-----
```

```

        default-value "no"
    exit
exit
app-qos-policy
  entry 400 create
  match
    charging-group neq "allow-list"
    characteristic "redirect" eq "yes"
  exit
  action
    drop
    http-redirect "policy-redirect" flow-type dropped-flows
  exit
  no shutdown
  exit
exit
commit
-----

```

Having the charging group as match criteria is a future-proof solution: the configuration can easily be adapted in the future (if, for example, additional URLs should be allowed).



Note:

The use of ASOs is optional and shown here for completeness. It is used only in some policy-dependent redirect cases, where the operator may need to provide the redirect policy only to some subscribers. In most cases, all subscribers get the same redirect policy and the ASOs are not used. App-profiles may also be used to implement a match condition for a policy-based redirect.

Troubleshooting and debugging

The following show commands can be used to verify the configuration and correct possible errors.

The following command shows the redirect policy:

```

*A:Dut-C# show application-assurance group 1 http-redirect "policy-redirect" detail
=====
Application Assurance Group 1 HTTP Redirect policy-redirect
=====
Description      : redirects certain http traffic to a portal
Template         : 5
                  : Redirect supporting macro substitution using HTTP 302
Redirect URL     : http://10.99.99.40/redirect_main.html?originalURL=$URL
Captive Redirect: No
  Redirect HTTPS: No
  VLAN ID       : N/A
Admin Status    : Up
AQP Ref         : Yes
-----
Group 1
-----
Redirects Sent   : 0
Client Resets Sent : 0
Redirects Not Sent :
  Out of Resources : 0
  Config Errors    : 0
  Other Errors     : 0

```

```

-----
Total
-----
Redirects Sent      : 0
Client Resets Sent : 0
Redirects Not Sent :
  Out of Resources : 0
  Config Errors    : 0
  Other Errors     : 0
=====

```

The preceding command provides the following information:

- The template used, along with a definition of that template
- The URL that the subscriber will be redirected to
- Whether this is a captive redirect and whether HTTPS will be redirected
- The admin status of the policy. An admin status of "Down" indicates that the policy is not in use.
- Whether this policy is referenced by an AQP. If it is not referenced, it is not used.
- "Redirects Sent" indicates the number of times that AA has sent a redirect message.
- "Client Resets Sent" indicates the number of times that TCP reset was sent.
- "Out of Resources" indicates the number of times that a redirect was not sent. The reasons relate to exhaustion of ISA resources (for example, memory).
- "Config Errors": for example, a command is in shutdown state
- "Other Errors": errors that are not in any of the preceding two categories

The following command shows the "Captive Portal" application:

```

*A:Dut-C# show application-assurance group 1 policy application "Captive Portal"

=====
Application Instance "Captive Portal"
=====
Description      : HTTP Redirect portal
App Group        : Web
Export Id        : none
Charging Group   : allow-list

References
-----
AA Sub Stat      : none
Cflowd           : none
AQP Entries      :
App Filters      :
                  200

Reference Counts
-----
No. of cflowd references : 0
No. of AQP entry references : 0
No. of app filter references : 1
=====

```

The preceding command provides the following information:

- The charging group assigned

- The app-filters linked (200)

The following command shows the app-filter for the redirect portal:

```
*A:Dut-C# show application-assurance group 1 policy app-filter 200

app-filter
  entry 200 create (0 flows, 0 B)
    server-address eq 10.99.99.40/32
    application "Captive Portal"
    no shutdown
  exit
exit
```

Using the preceding command, the operator can verify the number of flows and bytes matched and, therefore, determine if traffic toward the redirect portal was classified.

Show the AQP for redirection:

```
*A:Dut-C# show application-assurance group 1 policy app-qos-policy 400

=====
Application QoS Policy Entry 400 (Application Based Policy)
=====
Description : (Not Specified)
Admin State : in-service
Hits       : 0 flows
Conflicts  : 0 flows

Match :
  Charging Group       : neq allow-list
  ASO Characteristics :
  redirect             : eq yes

Action :
  Drop                 : yes
  HTTP Redirect       : policy-redirect flow-type dropped-flows
=====
```

Using the preceding display, the operator can verify whether flows are matched to the AQP.

Operational considerations

- QUIC is a protocol originally developed by Google to reduce end-to-end latency. It is used by Google services and browsers (Chrome). Because QUIC cannot be redirected, traffic that would have been redirected may be dropped instead. This can be achieved in two ways:
 1. Allow UDP traffic on port 53 and block all other UDP traffic.

QUIC is over UDP. Therefore, by allowing UDP on port 53 and blocking all other traffic, the operator can perform DNS requests, but all other UDP traffic (including QUIC) will be blocked.
 2. Create an application to block UDP traffic on port 443 at the first packet.

QUIC is running on port 443. Therefore, the following app-filter will block QUIC traffic (UDP traffic on port 443). The "first-packet-validate" option is used for applications using a well-known TCP/UDP port and ensures that the policy will be applied from the first packet, while allowing AA to detect an unexpected application to the well-known port.

```
policy
```

```


begin
application "UDP 443 Validate" create
  description "First Packet Verify application for any UDP on port 443
  such as QUIC"
  app-group "Web"
exit
app-filter
  entry 700 create
    ip-protocol-num eq udp
    server-port eq 443 first-packet-validate
    application "UDP 443 Validate"
    no shutdown
  exit
exit
commit

```

- Some mobile OSs (for example, Android) generate an HTTP request to a predefined "test URL" to verify that Internet access exists. The HTTP request is sent immediately after the device connects to a WiFi network. Even though the HTTP requests are generated in the background, redirecting these URLs will result in the redirect response to be presented in the web browser.

Depending on the operator's use case, these URLs may be redirected. Redirecting these URLs may affect the DNS cache of the mobile device. The device thinks that it does not have full Internet access. On Android devices, the subscriber will see a message similar to the following:

Figure 22: Connected device without full Internet access



The app-filter for the redirect server was configured using the server IP address (`server-address eq 192.0.2.1/32`). The app-filter for the redirect server can also be configured using a host (or any other HTTP-related) expression. An example configuration follows:

```

app-filter
  entry 200 create
    expression 1 http-host eq "*.redirect_domain.com$"
    application "Captive Portal"
    no shutdown
  exit

```

```
exit
```

The preceding configuration will classify HTTP traffic when the hostname is `*.redirect_domain.com$`.

If this configuration approach is chosen, the operator must be certain that the app-filter will classify all the traffic toward the redirect portal. There may be cases where some content may be fetched from additional locations too (for example, `*.content_domain.com$`); therefore, additional app-filters should be configured.

To ensure that the configuration is correct and complete, a trace can be analyzed.

[Table 8: AA redirect errors](#) describes what each AA redirect error means.

Table 8: AA redirect errors

| Redirect not sent | Description |
|-------------------|--|
| Out of resources | No buffer resources |
| Config errors | Disabled redirect configuration, VLAN ID missing, session is not TCP |
| Other errors | Fragmented response packet |

HTTPS redirect using session filter captive redirect

A second use case is HTTP redirect using session filter captive redirect. Captive redirect is used to redirect HTTP or HTTPS flows, without sending any traffic to the Internet. In the case of HTTP traffic, it is achieved by terminating the TCP sessions in AA. HTTP flows will then be redirected to a predefined URL, while non-HTTP flows will be TCP reset.

However, the majority of web traffic is HTTPS. Simply blocking (and TCP resetting) HTTPS traffic is not friendly to the end user: they are not given any explanation as to why access was not allowed, might think there was some different error, and will probably try again. AA can also redirect HTTPS traffic to a different page and provide information to the subscriber about why access to the site is not allowed.

Operators may still need to allow some HTTP traffic. Captive redirect can be combined with a DNS IP cache to allow (not redirect) certain HTTP traffic. The operator may configure a DNS IP cache for the allowed domains. This section will provide an example configuration to achieve this.

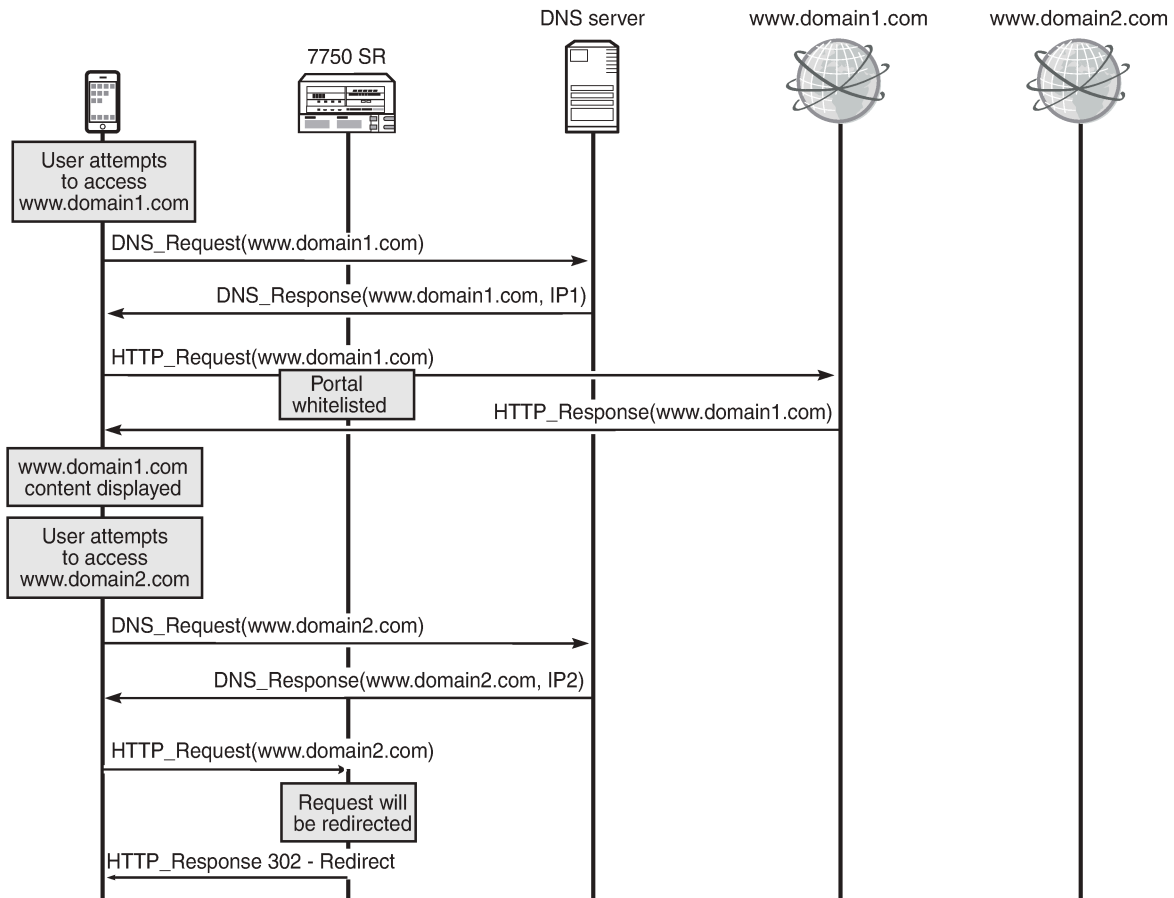
A DNS IP cache is used to ensure that traffic will be correctly classified and charged. After the feature has been enabled, AA will sniff subscribers' DNS responses and populate a cache containing the HTTP domain and the IP addresses that the DNS server provided. Traffic will only be classified into applications or app-groups if the hostname configured matches the one found, and the IP address of the packet matches an entry in the DNS cache.

Additional information about DNS IP cache (including configuration instructions) is available in the [Application Assurance — DNS IP Cache](#) chapter.

Captive redirect is mostly used when the operator wants to redirect HTTPS traffic. Even though this chapter also provides the configuration for HTTP captive redirect, it is only shown for completeness. Captive redirect of HTTP traffic is hardly ever used, adds more complexity, and provides no additional benefit over the HTTP redirect described in the previous section.

[Figure 23: HTTP redirect using session filter captive redirect](#) shows the high-level message flow of HTTP redirect using session filter captive redirect.

Figure 23: HTTP redirect using session filter captive redirect



27930

Consider an operator that needs to allow traffic to www.domain1.com and redirect all other HTTP traffic, see [Figure 23: HTTP redirect using session filter captive redirect](#). Therefore, the operator will configure the DNS IP cache to sniff DNS traffic and populate the cache with the IP addresses of that allowed domain.

- If the subscriber needs to access www.domain1.com, the DNS traffic will be analyzed by SR and the DNS IP cache will be updated. When the subscriber sends the HTTP request toward www.domain1.com, the traffic will be allowed because the destination IP address of the traffic will match the IP address contained in the cache.
- If the subscriber attempts to access www.domain2.com, the request will not be allowed because the destination IP address will not be present in the cache. This request will be redirected. The DNS IP cache will not be populated with the IP address of www.domain2.com.

Creating a DNS IP cache for the allowed domain

The following configuration will create and populate a DNS IP cache:

- DNS responses for the host `"*.domain1.com$"`. Traffic to www.domain.com will be allowed.

- responses only from the DNS server with IP address 8.1.17.21 (which is considered a trusted DNS server)

```
*A:Dut-C>config>app-assure>group# info
-----
  dns-ip-cache "domain1-cache" create
  description "dns cache for domain1"
  dns-match
    domain "domain1" expression "*.domain1.com$"
    server-address 8.1.17.21
  exit
  no shutdown
exit
-----
```

Creating a default subscriber policy

The following configuration will create a default subscriber policy to snoop DNS responses. This AQP is used to populate the cache.

No match criteria are needed. AA will only analyze the DNS traffic to populate the cache.

```
*A:Dut-C>config>app-assure>group# info
-----
policy
  begin
  app-qos-policy
    entry 50 create
    action
      dns-ip-cache "domain1-cache"
    exit
    no shutdown
  exit
  exit
  commit
exit
-----
```

Configuring captive redirect

The following will configure a captive redirect policy for the redirect portal. The configuration defines the redirect URL. Using macro substitution, the original URL that the subscriber requested will be appended in the redirect URL. Non-HTTP traffic will be TCP reset. The template specifies that the redirect template be used: "template 5" defines macro substitution using HTTP 302. HTTPS traffic will also be redirected.

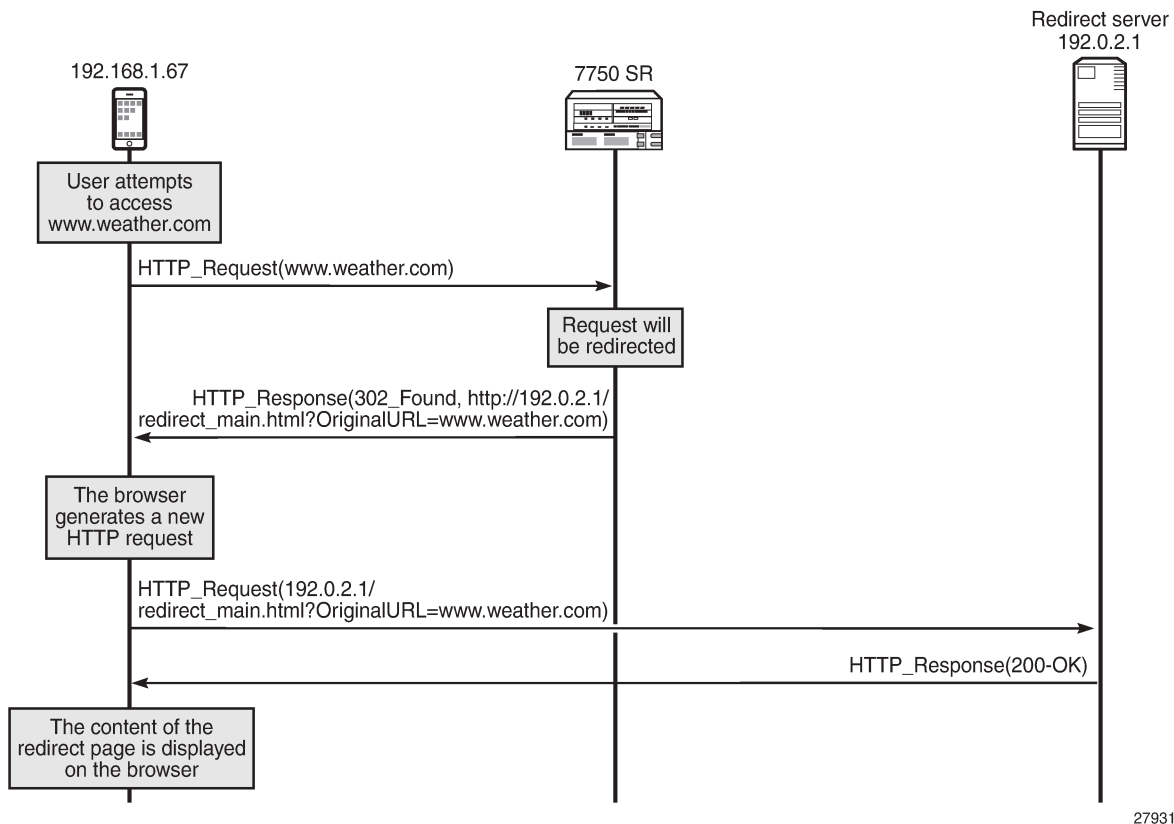
```
*A:Dut-C>config>app-assure>group# info
-----
http-redirect "capt-redirect" create
  template 5
  tcp-client-reset
  redirect-url "http:// 192.0.2.1/redirect_main.html?OriginalURL=$URL"
  captive-redirect
    redirect-https
  exit
  no shutdown
exit
```

The preceding configuration did not define the VLAN ID. The VLAN ID refers to an AA IP aa-interface that the operators need to create within the subscriber service (IES or VPRN) for ESM subscribers, and is used by the ISA to respond to the subscriber. Operators need to create one interface per AA ISA card.

The VLAN ID does not need to be defined if the ISA group aa-sub-scale mode is configured for DSM.

Figure 24: HTTP redirection shows the high-level message flow of HTTP redirection.

Figure 24: HTTP redirection




As shown in Figure 24: HTTP redirection:

- The subscriber attempts to access a site.
- The SR OS node determines that traffic will be redirected. No traffic will be allowed to pass.
- The SR OS node returns an HTTP response message indicating the redirect.
- The browser automatically generates a new request toward the redirect portal.
- The content of the redirect page is shown to the subscriber.

Figure 25: Access to domain included in the allow-list versus HTTPS redirect shows how HTTP access to a domain included in the allow-list is granted, while access to an HTTPS site is redirected.

Figure 25: Access to domain included in the allow-list versus HTTPS redirect



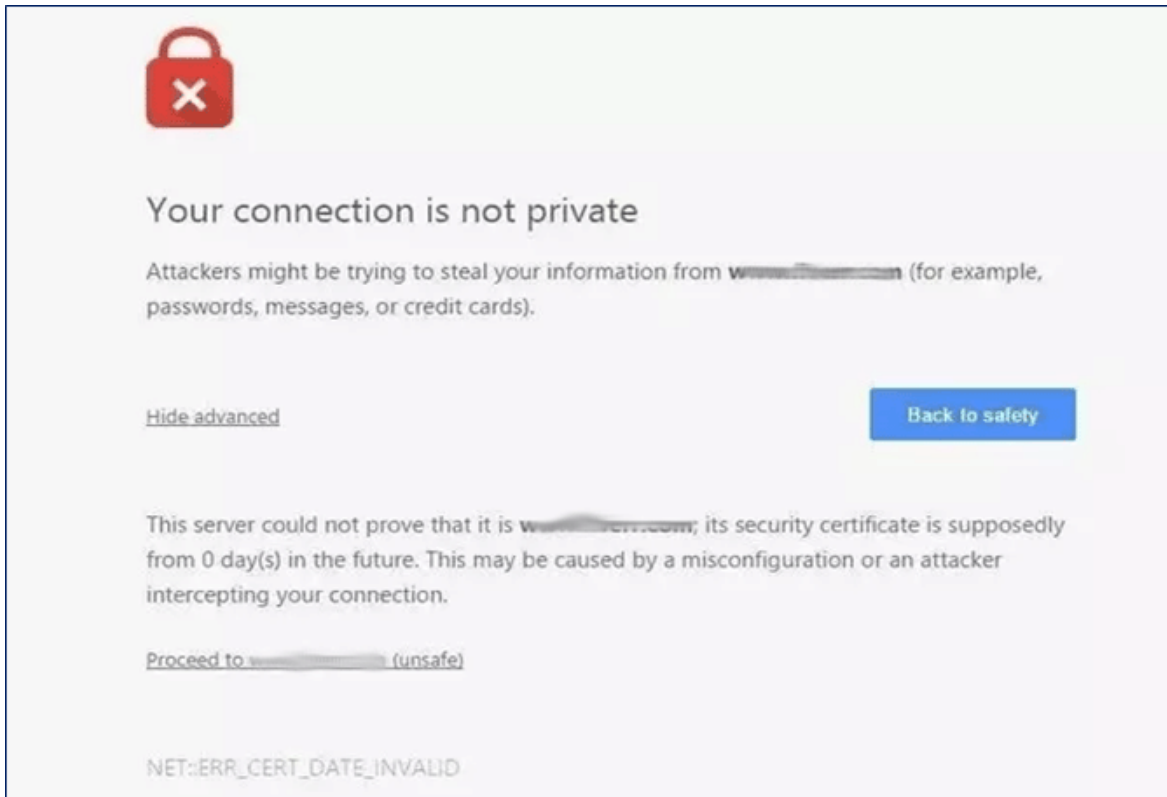
27932

As shown in [Figure 25: Access to domain included in the allow-list versus HTTPS redirect](#):

- The subscriber first attempts to access `www.domain1.com`. This site has been configured as a allowed site in SR OS.
- The subscriber next attempts to access the HTTPS site `www.portal.com`. This portal is not an allowed one. Therefore, the system will perform HTTPS redirection.

The system will return a self-signed Nokia certificate, similar to the following.

Figure 26: Subscriber receives warning



- If the subscriber accepts the new certificate, a new SSL tunnel will be established and a new request toward the redirect server will be generated.
- Finally, the subscriber will see the contents of the redirect page.

Configuring a session filter

The following session filter is configured to:

- Allow DNS traffic (UDP traffic on port 53)
- Allow domains included in the DNS IP cache (traffic to www.domain1.com)
- Allow traffic to the redirect portal (portal IP address: 192.0.2.1)
- Redirect all other TCP traffic to the redirect portal

```
*A:Dut-C>config>app-assure>group# info
-----
session-filter "capt_redirect" create
  description "Session filter for captive redirect policy"
  default-action deny
  entry 5 create
    match
      ip-protocol-num udp
      dst-port eq 53
    exit
```

```

        action permit
    exit
    entry 10 create
        match
            dst-ip dns-ip-cache "domain1-cache"
        exit
        action permit
    exit
    entry 15 create
        description "Allow traffic to the redirect landing page server"
        match
            ip-protocol-num tcp
            dst-ip 192.0.2.1/32
            dst-port eq 80
        exit
        action permit
    exit
    entry 20 create
        match
            ip-protocol-num tcp
        exit
        action http-redirect "capt-redirect"
    exit
exit
-----

```

Troubleshooting and debugging: show commands

Operators can verify their configuration as follows.

The following command shows the session filter together with the number of flows matched:

```

*A:Dut-C# show application-assurance group 1 session-filter "capt_redirect"

=====
AA Session Filter Instance "capt_redirect"
=====
Description      : Session filter for captive redirect policy
Default Action   : deny
Event Log        : (Not Specified)
AQP Entries      :
-----
Filter Match Criteria
-----
Entry           : 5
Description      : (Not Specified)
IP Protocol      : udp
Dest Port        : eq 53
Action           : permit
Event Log        : (Not Specified)
Hits             : 0 flows
-----
Entry           : 10
Description      : (Not Specified)
IP Protocol      : none
Dest IP Cache    : domain1-cache
Action           : permit
Event Log        : (Not Specified)
Hits             : 0 flows
-----
Entry           : 15

```

```

Description      : Allow traffic to the redirect landing page server
IP Protocol      : tcp
Dest IP          : 192.0.2.1/32
Dest Port        : eq 80
Action           : permit
  Event Log      : (Not Specified)
Hits             : 0 flows
-----
Entry            : 20
Description      : (Not Specified)
IP Protocol      : tcp
Action           : http-redirect capt-redirect
  Event Log      : (Not Specified)
Hits             : 0 flows
-----
No. of entries   : 4
=====
    
```

The preceding show output provides the following information:

- The "Default Action" specifies the action to take if a packet does not match any of the filters defined in the session filter (deny).
- The filters defined are:
 - Allowing DNS traffic (UDP traffic on port 53, entry 5)
 - Allowing traffic matching the DNS IP cache (entry 10)
 - Allowing traffic to the redirect portal (entry 15)
 - Redirecting all other traffic (entry 20)
- The "Hits" counter on each filter provides information about how many flows matched the corresponding filter.

The following command shows the DNS IP cache:

```

*A:Dut-C# show application-assurance group 1 dns-ip-cache "domain1-cache"

=====
Application Assurance Group 1 dns-ip-cache "domain1-cache"
=====
Admin Status           : Up
AQP Ref                : Yes
Domain expressions     : 1 (out of 32)
Server addresses       : 1 (out of 64)
High watermark         : 90%
Low watermark          : 80%
Cache size              : 10
-----
ISA                    Usage          (%)          Alarm State          Hit Count
-----
1/2                    0          0.00          clear                0
3/1                    0          0.00          clear                0
=====
    
```

The preceding command shows:

- The administrative status of the DNS IP cache ("Up"). An administrative status of "Down" indicates that the DNS IP cache is not in use.
- Whether the DNS IP cache created is referenced by an AQP

- The number of configured domain expressions (in the example: "*.domain1.com\$") and the maximum value (shown in brackets)
- The number of server addresses and the maximum value (shown in brackets)
- The number of times a cache lookup was successful ("Hit Count"). A hit count of 0 indicates a possible configuration error.

The following command shows information about the HTTP redirect:

```
*A:Dut-C# show application-assurance group 1 http-redirect "capt-redirect" detail

=====
Application Assurance Group 1 HTTP Redirect capt-redirect
=====
Description      : (Not Specified)
Template         : 5
                 : Redirect supporting macro substitution using HTTP 302
Redirect URL     : http://192.0.2.1/redirect_main.html?originalURL=$URL
Captive Redirect : Yes
  Redirect HTTPS : Yes
  VLAN ID       : N/A
Admin Status    : Up
AQP Ref        : No

-----
Group 1
-----
Redirects Sent      : 0
Client Resets Sent : 0
Redirects Not Sent :
  Out of Resources : 0
  Config Errors    : 0
  Other Errors     : 0

-----
Total
-----
Redirects Sent      : 0
Client Resets Sent : 0
Redirects Not Sent :
  Out of Resources : 0
  Config Errors    : 0
  Other Errors     : 0
=====
```

The preceding output provides the following information:

- The template used and description about the template
- The redirect URL
- Whether this is a captive redirect and whether HTTPS will be redirected also
- The VLAN ID. Captive redirect uses the provisioned VLAN ID to send the HTTP response to the subscribers; therefore, this VLAN ID must be properly assigned in the same VPN as the subscriber (not needed for DSM).
- The admin status of the redirect. An admin status of "Down" indicates that the policy is not in use.
- "Redirects Sent" indicates the number of times that AA has sent a redirect message.
- "Client Resets Sent" indicates the number of times that TCP reset was sent.

- "Out of Resources" indicates the number of times that a redirect was not sent. The reasons relate to exhaustion of resources in ISA (such as memory).
- "Config Errors": for example, a command is in shutdown state
- "Other Errors": errors that are not in any of the preceding two categories

Troubleshooting and debugging: tools commands

Based on an FTP server with IP address 192.168.1.100, the operator can dump the contents of the DNS IP cache, as follows:

```
*A:Dut-C# tools dump application-assurance group 1 dns-ip-cache "domain1-cache" url ftp://
username:password@192.168.1.100/path/to/mylog.log
```

The preceding command will generate a text file, containing the following:

```
=====
Application-Assurance dns-ip-cache "domain1-cache"
Current Time:         "11/08/2017 12:45:32" (UTC)
  group:              1
  isa:                3/2
  admin state:       no shutdown
  max-entries:       10
=====
ip-address  creationTime(UTC)  lastUpdated(sec)  numDNSResponses  lastMatchTime(UTC)  numTimes
Matched
Total entries in-use:  0
=====
```

The log file shows the number of times that a cache entry was matched (numDNSResponses) and the number of times that traffic to this address was sent (numTimesMarches).

Operational considerations

HTTPS redirect is not available in Releases earlier than SR OS 14.0 (the session will simply be TCP reset).

Some browsers store certificates of popular sites. If the subscriber tries to access one of those sites, the browser will not accept the AA-signed certificate and the redirect page will never be displayed. The subscriber does not have an option to accept the certificate.

If the operator wants to configure an HTTPS redirect (either on VSR since Rel. 14.0 or in CMG since Rel.10.0), the FW ASL License is mandatory. It is needed to enable the use of the session filter by the captive redirect.

When using the standard HTTP redirect, there is no need for an aa-interface. However, when using the captive redirect, the operator needs one aa-interface per VPN service, not one interface per subscriber VLAN. To associate the redirect policy with the subscriber in the correct VPN for captive redirect, the operator needs to use ASO/ASO override, which will point to a redirect policy per VPN-this redirect policy will point to an aa-interface in the correct VPN service.

Captive redirect is special in the sense that no traffic is allowed from the Internet (apart from the traffic to the redirect portal).

HTTPS captive redirect cannot be used with ESM subscribers on L2-aware NAT. The forwarding model used for L2-aware NAT does not allow the use of HTTPS captive redirect. If configured, the session filter will be hit, but no TCP message will be sent to the subscriber. The AA statistics will show discarded packets.

Conclusion

This chapter shows how an operator can configure HTTP redirection. The chapter describes the different use cases, how HTTP redirection can be combined with DNS IP cache, and how HTTPS traffic can be redirected.

Application Assurance — HTTP In Browser Notification

This chapter provides information about Application Assurance HTTP in browser notification.

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.R2.

There are no specific prerequisites for this example.

Overview

Using the SR OS nodes and application assurance (AA), subscribers connected to an operator network can be sent fully customizable on-screen notification messages displayed in a non-disruptive and cost-effective manner through their web browser.

This chapter describes the different options for the operator to customize the notification messages returned to the subscriber using either different HTTP-notification policies or using the flexible HTTP-URL-PARAM VSA mechanism.


This chapter also describes the additional configuration required with the introduction in SR OS 13.0.R1 of the Notification status monitoring capability allowing the system to notify the subscriber at the next candidate flow instead of waiting for the next notification interval in case the previous notification did not result in a success.

Configuration

The setup comprises of the following elements, see [Figure 27: HTTP notification – setup](#):


- SR node with ISA-AA.
- Apache web server (delivering notification Javascript and content).
- Subscriber (desktop/tablet/smartphone).
- Authentication, authorization and accounting (AAA) for subscriber authentication and policy modification.
- Internet access.

Figure 27: HTTP notification – setup



This example describes how to configure HTTP notification to display different notification messages. It demonstrates a simple example in the context of a residential deployment where a message is displayed when the subscriber reaches 80% or 100% of their maximum allowed volume (usage cap).

Figure 28: Notification Message Example – Quota 80%



In this context the operator has two options:

- Use a dedicated http-notification policy per message type.
- Use a common http-notification policy for any message type together with the newly introduced Http-Url-Param RADIUS VSA.

This example provides configuration examples for both options.

HTTP Notification Policy per Message Type

In this option a dedicated http-notification policy for each notification message is required.

HTTP Notification Policy

Two dedicated HTTP notification policies are used to return a different message to the subscribers when reaching 80% and 100% of their usage cap, the interval in between notifications is set to 15 minutes.

```
configure
  application-assurance group 1
    http-notification "notification-quota-100" create
      description "100% Usage Cap Notification"
      script-url "http://192.168.150.251/In-Browser-Notification/script/quota-100.js"
      template 1
      interval 15
      no shutdown
    exit

configure
  application-assurance group 1
    http-notification "notification-quota-80" create
      description "80% Usage Cap Notification"
      script-url "http://192.168.150.251/In-Browser-Notification/script/quota-80.js"
      template 1
      interval 15
      no shutdown
    exit
```

Notification Status Monitoring

The operator then needs to enable the `http-match-all-req` feature for any HTTP request sent to the messaging server in order to monitor HTTP notification success and failures. This is done by creating a new application and enabling `http-match-all-req` within the `app-filter`.

Success and failure notifications include a specific HTTP encoded URI automatically interpreted as a success or a failure by Application Assurance on a per subscriber basis. If a failure is detected, the system will automatically attempt to notify a new candidate flow instead of waiting for the next notification interval.

```
configure
  application-assurance group 1:1 policy
    application "IBN Messaging Server" create
    app-group "Web"
  exit
  app-filter
    entry 100 create
      expression 1 http-host eq "^192.168.150.251$"
      http-match-all-req
      application "IBN Messaging Server"
      no shutdown
    exit
  exit
```

App-Profiles and App-Service-Options

Event based HTTP notifications is enabled by a policy modification triggered via RADIUS or Gx by modifying the subscriber app-profile or using the Application Service Option (ASO) override.

In this implementation of the HTTP notification policy per message type, the following ASO configuration is used:

```
configure
  application-assurance group 1:1 policy
    app-service-option
      characteristic "quota-message-notification" create
        value "100"
        value "80"
        value "disabled"
        default-value "disabled"
    exit
  exit
  app-profile "1-1/Default" create
    divert
  exit
```

The ASO characteristic *quota-message-notification* values of 100 and 80 enable the App-Qos-Policy (AQP) *notification-quota-100* and *notification-quota-80* as defined below:

```
configure
  application-assurance group 1:1 policy app-qos-policy
    entry 1000 create
      match
        characteristic "quota-message-notification" eq "100"
        application neq "Advertising Statistics"
      exit
      action
        http-notification "notification-quota-100"
      exit
      no shutdown
    exit
    entry 1100 create
      match
        characteristic "quota-message-notification" eq "80"
        application neq "Advertising Statistics"
      exit
      action
        http-notification "notification-quota-80"
      exit
      no shutdown
    exit
```

RADIUS Policy

The following RADIUS CoA message is used to override the ASO characteristic of a residential subscriber so that a notification message can be returned to the subscriber when they reach 80% of their usage cap:

```
NAS-Port-Id = "1/1/5:4088"
Framed-IP-Address = 192.168.211.30
Alc-AA-App-Service-Options = "quota-message-notification=80"
Alc-App-Prof-Str = "1-1/Default"
```

Show Commands

Before the subscriber usage cap limit is reached, and before the RADIUS CoA message is received, the subscriber ASO parameter flag *quota-message-notification* is set to its default value *disabled* and therefore no App QoS Policy is triggered.

```
A:PE# show application-assurance group 1:1 aa-sub esm "sub1" summary
=====
Application-Assurance Subscriber Summary (realtime)
=====
AA-Subscriber      : sub1 (esm)
ISA assigned       : 1/2
App-Profile        : 1-1/Default
App-Profile divert : Yes
Capacity cost      : 1
Aarp Instance Id   : N/A
HTTP URL Parameters : (Not Specified)
Last HTTP Notified Time : N/A
-----
Traffic            Octets           Packets           Flows
-----
... ..
-----
Application Service Options (ASO)
-----
Characteristic      Value              Derived from
-----
quota-message-notification  disabled          default
=====
```

After the RADIUS CoA message is sent, the subscriber ASO characteristic *quota-message-notification* value is set to *80*, the subscriber-related App QoS Policy entry 1100 now matches for this subscriber:

```
A:PE# show application-assurance group 1:1 aa-sub esm "sub1" summary
=====
Application-Assurance Subscriber Summary (realtime)
=====
AA-Subscriber      : sub1 (esm)
ISA assigned       : 1/2
App-Profile        : 1-1/Default
App-Profile divert : Yes
Capacity cost      : 1
Aarp Instance Id   : N/A
HTTP URL Parameters : (Not Specified)
Last HTTP Notified Time : N/A
-----
Traffic            Octets           Packets           Flows
-----
... ..
-----
Application Service Options (ASO)
-----
Characteristic      Value              Derived from
-----
quota-message-notification  80                dyn-override
=====
```

The same command can be used to identify when the last successful subscriber notification occurred, see the Last HTTP Notified Time field:

```
A:PE# show application-assurance group 1:1 aa-sub esm "sub1" summary
```

```

=====
Application-Assurance Subscriber Summary (realtime)
=====
AA-Subscriber      : sub1 (esm)
ISA assigned       : 1/2
App-Profile        : 1-1/Default
App-Profile divert : Yes
Capacity cost      : 1
Aarp Instance Id   : N/A
HTTP URL Parameters : (Not Specified)
Last HTTP Notified Time : 2014/06/24 15:35:49
-----
    
```

The operator can also identify how many notifications have been sent per http-notification policy per partition:

```

A:PE# show application-assurance group 1 http-notification "notification-quota-80"
=====
Application Assurance Group 1 HTTP Notification "notification-quota-80"
=====
Description  : 80% Usage Cap Notification
Template     : 1 - Javascript-url with subId and optional Http-Url-Param
Script URL   : http://192.168.150.251/In-Browser-Notification/script/quota-80.
              js
Admin Status : Up
AQP Ref      : Yes
Interval     : 15 minutes
-----
Group 1:1 Statistics
-----
Notified      : 2                Succeeded : 2
Criteria Not Matched : 5          Failed      : 0
=====
    
```

The counter Criteria Not Matched is the number of HTTP flows which did not meet the AA ISA flow selection criteria for In Browser Notification. HTTP flow selection is constrained so that only HTTP web pages flows originating from a web browser are targeted, HTTP requests for content such as video or images are not candidate for notification.

HTTP Notification Customization using RADIUS VSA

Instead of using a dedicated HTTP notification policy for every single message type, the operator can return a RADIUS Http-Url-Param VSA at subscriber creation time or via CoA to customize the notification URL using a single policy. This VSA string is automatically appended to the end of the HTTP notification script-url by the 7x50 which can then be used by the web server to decide which notification message to return to the subscriber.

SR OS Release supports 1 active HTTP Notification policy per subscriber, 8 different HTTP notification policies per AA ISA group and 1500 different values for the Http-Url-Param VSA. Therefore, using the Http-Url-Param VSA for the customization of the notification is the recommended model to scale the number of notification messages.

For example:

- RADIUS VSA (Alc-AA-Sub-Http-Url-Param): &message=quota80"
- 7750 HTTP Notification configured script-url: http://1.1.1.1/notification.js

- Subscriber HTTP request to the messaging server:

```
http://1.1.1.1/notification.js?SubId=sub1&var=&message=quota80
```

HTTP Notification Policy

A single HTTP notification policy is used to return different notification messages. The interval between notifications is set to 15 minutes.

```
configure
  application-assurance group 1
    http-notification "in-browser-notification" create
      description "Default HTTP Notification Policy"
      script-url "http://192.168.150.251/In-Browser-Notification/script/
                notification-select.php"

      template 1
      interval 15
      no shutdown
  exit
```



Note:

This example does not describe the content of the *notification-select.php* file used to parse the URL parameters.

Notification Status Monitoring

The operator then needs to enable the `http-match-all-req` feature for any HTTP request sent to the messaging server in order to monitor HTTP notification success and failures. This is done by creating a new application and enabling `http-match-all-req` within the `app-filter`.

Success and failure notifications include a specific HTTP encoded URI automatically interpreted as a success or a failure by Application Assurance on a per subscriber basis. If a failure is detected, the system will automatically attempt to notify a new candidate flow, instead of waiting for the next notification interval.

```
configure
  application-assurance group 1:1 policy
    application "IBN Messaging Server" create
      app-group "Web"
  exit
  app-filter
    entry 100 create
      expression 1 http-host eq "^192.168.150.251$"
      http-match-all-req
      application "IBN Messaging Server"
      no shutdown
  exit
exit
```

App-Profile and App-Service-Options

Similar to the previous example, HTTP notifications are enabled per subscriber using RADIUS or Gx by modifying the subscriber `app-profile` or using ASO override.

The following ASO configuration is used:

```
configure
  application-assurance group 1:1 policy
    app-service-option
      characteristic "in-browser-notification"
        value "enabled"
        value "disabled"
        default-value "disabled"
  exit
```

The ASO characteristic *in-browser-notification* value *enabled* is used to enable the app-qos-policy matching the http-notification policy *in-browser-notification* as shown below:

```
configure
  application-assurance group 1:1 policy app-qos-policy
    entry 1300 create
      match
        characteristic "in-browser-notification" eq "enabled"
        application neq "Advertising Statistics"
      exit
      action
        http-notification "in-browser-notification"
      exit
    no shutdown
```

RADIUS Policy

The following RADIUS CoA message is used to modify the ASO characteristic of a residential subscriber and assign a specific Http-Url-Param VSA. The *in-browser-notification* ASO characteristic with value *enabled* is dynamically assigned to the subscriber along with the **Http-Url-Param &message=quota80**:

```
NAS-Port-Id = "1/1/5:4088"
Framed-IP-Address = 192.168.211.30
Alc-AA-App-Service-Options = "in-browser-notification=enabled"
Alc-AA-Sub-Http-Url-Param = "&message=quota80"
Alc-App-Prof-Str = "1-1/Default"
```

The subscriber HTTP request to the messaging server has the following format and includes the Http-Url-Param value as an argument of the URL:

```
http://192.168.150.251/In-Browser-Notification/script/notification-select.php?SubId=sub1&var=
&message=quota80
```

The web server can now use the parameter value to make a decision to return a suitable notification message related to the subscriber usage cap.

Show Commands

Both the *in-browser-notification* ASO characteristic with value *enabled* and the HTTP-Url-Param VSA can be shown as follows:

```
A:PE# show application-assurance group 1:1 aa-sub esm "sub1" summary
```

```
=====
Application-Assurance Subscriber Summary (realtime)
```

```

=====
AA-Subscriber      : sub1 (esm)
ISA assigned       : 1/2
App-Profile        : 1-1/Default
App-Profile divert : Yes
Capacity cost      : 1
Aarp Instance Id   : N/A
HTTP URL Parameters : &message=quota80
Last HTTP Notified Time : N/A
-----
Traffic            Octets            Packets            Flows
-----
... ..
-----
Application Service Options (ASO)
-----
Characteristic      Value            Derived from
-----
in-browser-notification    enabled        dyn-override
quota-message-notification disabled        default
=====

```

The operator can also display the HTTP URL parameters VSA currently in use, per AA ISA group:

```

A:PE## tools dump application-assurance group 1 http-url-param-list
-----
Application-Assurance Subscriber HTTP URL parameters for Group 1:
-----
=====
Http Url Parameter      Sub Usage
-----
"&message=quota80"      1
=====
Total entries displayed 1

```

Conclusion

This chapter, intended for Application Assurance (AA) network architects and engineers, provides two implementation options for configuring and deploying HTTP In Browser Notification. It also explains how to take advantage of the Http-Url-Param RADIUS VSA to flexibly define various messaging campaigns using a common AA notification policy.

Application Assurance — Local URL List Filtering

This chapter provides information about the Application Assurance local URL list filtering.

Topics in this chapter include:

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

Applicability

This chapter is applicable to all 7x50 systems supporting Application Assurance and was initially based on SR OS Release 13.0.R2. The chapter has been updated with CLI and functionality changes introduced in SR OS Release 19.10.R1, but the CLI in the sections [URL-list update](#) and [Cron job](#) is based on SR OS Release 16.0.

There are no specific prerequisites for this feature.

Overview

The local URL-list filtering capability provided by Application Assurance offers the following:

- configuration of a list of URLs which are globally allowed (allow-list)
- configuration of a list of restricted URLs subscribers should be prevented from accessing (deny-list)

The lists are stored locally in the 7x50 system compact flash (CF).

This chapter provides a guide for configuring deny-lists.

Deny-lists assist service providers to comply with regulatory requirements for network-wide URL filtering policies, such as:

- Court-ordered URL takedown
- Child protection
- Government-mandated URL takedown list

The 7x50 uses the Application Assurance capabilities to extract the URL from a subscriber HTTP/HTTPS request and compare it to the list of URLs contained in the local file. If a match occurs, the subscriber request is redirected to a preconfigured web server landing page, typically describing why the access to this resource was denied.

The system supports both unencrypted and OpenSSL Triple Data Encryption Standard (3DES) encrypted file formats to protect the content of the list.

HTTP/HTTPS filtering


Each HTTP request within a TCP flow is analyzed and filtered. For HTTPS traffic, the system extracts the domain name information contained in the Transport Layer Security (TLS) Server Name Indication (SNI).

Setup details

The setup consists of the following elements, as shown in [Figure 29: Local URL-list filtering setup](#):

- 7x50 + ISA-AA
- Web server (redirect landing page)
- FTP server (source for the URL-list file)
- Subscriber (desktop/laptop/tablet/smart phone)
- Internet access
- Optional: AAA for subscriber authentication and policy modification

Figure 29: Local URL-list filtering setup



This chapter is written in the context of a residential or WiFi deployment. However, local URL-list filtering is also applicable to business VPN services.

Configuration

To configure the system for local URL-list filtering, the operator needs to:

- Create a URL-list policy referencing a valid URL-list file located on the system compact flash
- Create a URL-filter policy for local filtering by referencing the URL-list policy previously created
- Create an App-QoS-Policy (AQP) to apply this url-filter policy
- Optionally configure a cron job to automatically fetch a new list file and upgrade the URL list.

URL-list policy and URL-filter policy

In the following example, two dedicated URL-list and URL-filter policies show URL filtering based on a plain text file and an encrypted file:

```
configure
  application-assurance group 1
    url-list "denylist1-encrypted" create
      description "Demo URL Filtering List - Encrypted File"
      decrypt-key "ON3HU2GFPmp0HwWbSGw/zdM4iuxzySpqS7pw/u3qIcuG4mABmrhc."
                                                    hash2
      file "cf3:\aa-url-list\url-list1.encrypted"
      no shutdown
    exit
  url-filter "local-filter-list1-encrypted" create
    local-filtering
      deny-list "denylist1-encrypted"
      default-action allow
      http-redirect "redirect-denylist"
    exit
  no shutdown
exit

configure
  application-assurance group 1
    url-list "denylist1-plaintext" create
      description "Demo URL Filtering List - Plaintext File"
      file "cf3:\aa-url-list\url-list1-plaintext.txt"
      no shutdown
    exit
  url-filter "local-filter-list1-txt" create
    local-filtering
      deny-list "denylist1-plaintext"
      default-action allow
      http-redirect "redirect-denylist"
    exit
  no shutdown
exit
```

In the preceding example, both URL-filter policies are defined using **default-action allow**. The default action is used in case the file could not be loaded by the system, either at boot time or the first time the URL-list file was configured in the system. Possible causes are, for example:

- File corrupted, compact flash corrupted
- Incorrect file encryption format or password
- Wrong URL format in the file
- Too many URLs in the file

Operators should always use **default-action allow** when configuring the URL-filter policy associated with a URL-list file because the file or the CF may be corrupted, in which case the system logs an error and a trap is raised.

Note that if a valid URL-list file was previously in use, and an invalid file is uploaded and the URL-list policy upgraded using this file, then the system will continue using the previous list.

HTTP-redirect policy

Both URL-filter policies defined in the preceding example refer to the following http-redirect policy; subscribers accessing a URL from the URL-list file are redirected to the following landing page:

```
configure
  application-assurance group 1
    http-redirect "redirect-denylist" create
      description "Redirect for Local List URL Filtering"
      template 5
      tcp-client-reset
      redirect-url "http://172.16.70.100/Redirect/redirect-denylist.html?
                  Request edURL=$URL"
    no shutdown
  exit
```

URL-list file

File format

A URL-list file may contain either hostnames or URLs.

To create a URL-list containing hostnames, set **expression-match** in the url-list configuration:

```
config>app-assure>group# url-list url-list-name [create]
      description <description-string>
      no description
      decrypt-key key | hash-key | hash2-key [hash | hash2]
      no decrypt-key
      file file-url
      no file
      size url-list-size
      expression-match
      [no] shutdown
```

A URL-list with hostnames only (using **expression-match**) may contain the following wildcards:

- Head anchors character set [^ *]
- Tail anchors character set [\$ *]
- Mid expression character set [\d \. \. **]
- Hex escaped characters [\x00 - \xFF]

Note that when **expression-match** is enabled, the list should contain hostnames only (with optional wildcards).

When configuring a URL-list with **expression-match** disabled (default), the system supports the following format for the URLs contained in the URL-list file:

- URLs without the HTTP keyword. For example:

```
www.domain.com/path
```

- URLs with the HTTP keyword. For example:

```
http://www.domain.com/path
```

In all cases, the following is supported:

- Comment lines starting with the number sign character (#). For example:

```
# This is a comment line
```

- Printable ASCII characters. URLs using non-printable ASCII characters are percent-encoded by the web browser automatically and, therefore, need to be percent-encoded in the URL-list file.

File encryption

OpenSSL triple DES -nosalt is the supported encryption format. Files can be encrypted offline on a server using the following command:

```
openssl des3 -nosalt -in <input.txt> -out <output.enc>
```

List upgrade

The URL-list file can be upgraded using the **admin** command:

```
A:BNG# admin application-assurance group 1 url-list "denylist1-plaintext" upgrade
```

The upgrade result is logged in the system log-id 99:

```
A:BNG# show log log-id 99
=====
Event Log 99
=====
Description : Default System Log
Memory Log contents [size=500 next event=72 (not wrapped)]
71 2015/07/07 13:09:25.01 EST MINOR: APPLICATION_ASSURANCE #4446 Base url-list success
"URL list "denylist1-plaintext" in ISA-AA group 1 has been updated. There are 3 entries in the
URL list."
```

App-profiles and app-service-options

Application Assurance policies can be selectively applied to specific AA subscribers by modifying the app-profile assigned to the subscriber or using Application Service Option (ASO) override. See SR OS User Guide for more information about modifying the app-profile or ASO assigned to AA subscribers (RADIUS, Gx, Override).

In this example, the following ASO configuration is used:

```
configure
  application-assurance group 1:1 policy
    app-service-options
      characteristic "local-list-filtering" create
        value "no"
        value "yes-encrypted"
```



```

        value "yes-plaintext"
        default-value "no"
    exit
exit
app-profile "1-1/Default" create
divert
exit

```

The ASO characteristic local-list-filtering value of yes-encrypted and yes-plaintext enable the AQP entry 210 and 220 in the example:

```

configure
  application-assurance group 1:1 policy app-qos-policy
    entry 210 create
      match
        characteristic "local-list-filtering" eq "yes-encrypted"
      exit
      action
        url-filter "local-filter-list1-encrypted"
      exit
      no shutdown
    exit
  entry 220 create
    match
      characteristic "local-list-filtering" eq "yes-plaintext"
    exit
    action
      url-filter "local-filter-list1-txt"
    exit
    no shutdown
  exit

```

If the url-filter policy needs to be applied to 100% of the subscribers in the network, it is also possible to remove the ASO match criteria.

URL-list update

The system supports a flexible mechanism to upgrade the URL list automatically, using either cron or the NSP, to comply with the regulatory requirements for list upgrade frequency.

To configure a crontab job to periodically upgrade the URL list, the operator needs to:

- Generate the file to be periodically executed and store it to compact flash. As an example, create a file with filename "fetch.txt" with the following content:

```

file
copy ftp://user:pwd@192.168.1.99/home/cmg/test/list.txt . force
exit
admin application-assurance group 1 url-list "test" upgrade

```

The preceding commands will fetch a file from an ftp server and store it in compact flash. Assuming the operator has configured a local url-list "test" containing the file "list.txt", the url list will be upgraded.

- Configure the script policy:

```

>config>system>script-control# info
-----
    script "bring"
      location "cf3:/fetch.txt"

```

```

        no shutdown
    exit
    script-policy "test"
        results "cf3:/results.txt"
        script "bring"
        no shutdown
    exit

```

- Configure the crontab job:

```

>config>system>cron# info
-----
    schedule "bring_list"
        interval 60
        script-policy "test"
        no shutdown
    exit
-----

```

The preceding configuration will execute the commands stored in the file "fetch.txt" every 60 seconds. A value of 60 seconds was chosen for the test. In a real deployment, a list would be typically updated every 12 – 24 hours. A log file (results "results.txt") will also be created.

The end result will be that the system will automatically fetch from an ftp server a new list file, store it to compact flash and upgrade the URL list.

Show commands

url-list

The status of the URL list can be shown in the CLI. The url-list show command provides basic admin and operational status, as well as the number of URLs in the list. The command also provides reasons for any possible issue related to loading the list, as well as the last successfully deployed file and the last upgrade attempt. Therefore, the operator can determine whether the latest version of the file is currently in use or if an error occurred when trying to upgrade the list.

Show command output:

| Label | Description |
|----------------------|--|
| Admin | Status [Up Down] - Administrative status of the url-list |
| Oper Status | [Up Down] - Operational status of the url-list |
| Oper Flags | [admin-down file-does-not-exist invalid-file-format too-many-urls switch-over-error] |
| File Deployed to ISA | [Yes No] - This flag describes if the file located in the compact flash is the one deployed in the ISA, in the event the file is overwritten and before the admin upgrade command is used this flag will display "No". |
| Upgrade Statistics | |
| Last Success | Last time the list was successfully upgraded |
| File Name | File name for the last successful upgrade |
| URL Entries | Number of URLs loaded at the last success |
| Blank/Comment Lines | Number of blank or commented out lines |
| Last Attempt | Last time the operator tried to upgrade the list |
| Result | Success Failure. Result of the last upgrade |

| | |
|-----------|--|
| File Name | File name for the last upgrade attempt |
|-----------|--|

```
*A:Dut-C# show application-assurance group 1 url-list "Deny List1"
```

```
=====
Application Assurance Group 1 url-list "Deny List1"
=====
```

```
Description : (Not Specified)
Size : standard
Host Expressions : disabled
Admin Status : Up
Oper Status : Up
Oper Flags : <none>
File deployed to ISAs : Yes
```

```
-----
Upgrade Statistics
-----
```

```
Last Success : 11/02/2020 15:07:40
Deployed
File Name : cf1:/host.txt
URL Entries : 1 ( 0.01% full)
URL Characters : 6 (~0.00% full)
URL Host Expr Entries: 1 ( 0.01% full)
Blank/Comment Lines : 1
Last Attempt : 11/02/2020 15:07:40
Result : Success
File Name : cf1:/host.txt
=====
```

url-filter

The **url-filter** show command provides its operational and admin status, as well as actions taken, such as the number of redirects. With URL list filtering, using a **default-action** set to **allow**, the only counters increasing are **allow**, **redirect**, and **default**.

```
*A:Dut-C>config>app-assure>group>url-filter# show application-assurance group 1 url-filter "Url Filter1"
```

```
=====
Application Assurance Group 1 URL Filter "Url Filter1"
=====
```

```
Description          : (Not Specified)
Admin Status         : Up
Oper Status          : Up
Oper Flags           : <none>
HTTP Request Filtering : all
AQP Referenced       : No
```

```
-----
URL Stats Summary
-----
```

```
Total Requests : 0          Default Action : 0
Requests Allowed: 0          Reqs Block/Redir: 0
```

```
-----
Local Filter
-----
```

```
deny-list           : Deny List1
Admin Status        : Up
Oper Status         : Up
Oper Flags          : <none>
Number of URLs      : 1
Default Action      : block-all
```

```

HTTP Redirect      : (Not Specified)
URL-List Lookups  : 0
  Match           : 0
  Miss            : 0
  Default Action  : 0
=====
    
```

http-redirect

The **http-redirect** show command provides more information about how the traffic was blocked; for example, it differentiates TCP client reset used for HTTPS from regular redirect used for HTTP traffic.

```

A:BNG# show application-assurance group 1 http-redirect "redirect-denylist"
=====
Application Assurance Group 1 HTTP Redirect redirect-denylist
=====
Description      : Redirect for Local List URL Filtering
Template         : 5
                 : Redirect supporting macro substitution using HTTP 302
Redirect URL     : http://172.16.70.100/Redirect/redirect-denylist.
                 : html?RequestedURL=$URL
Admin Status    : Up
AQP Ref         : No
-----
Summary Statistics
-----
Grp:Part          Redirects      Client Resets      Redirects
                  Sent                Sent                Not Sent
-----
1:1                2                    1                    0
-----
Total              2                    1                    0
=====
    
```

Cron job

This chapter provides details on how to verify that the cron job executes successfully. The operator can see the location and execution interval of the file which is executed periodically, the time it was last executed and any possible errors which may have occurred during execution using the following command:

```

# show system cron schedule "bring_list"
=====
CRON Schedule Information
=====
Schedule          : bring_list
Schedule owner    : TiMOS CLI
Description       : none
Administrative status : enabled
Operational status : enabled
Script Policy     : test
Script Policy Owner : TiMOS CLI
Script            : bring
Script Owner      : TiMOS CLI
Script source location : cf3:/fetch.txt
Script results location : cf3:/results.txt
    
```

```

Schedule type      : periodic
Interval          : 0d 00:01:00 (60 seconds)
Repeat count      : infinite
Next scheduled run : 0d 00:00:51
End time          : none
Weekday           : none
Month             : none
Day of month      : none
Hour              : none
Minute            : none
Number of schedule runs : 12
Last schedule run  : 2019/09/13 11:28:25 EEST
Number of schedule failures : 0
Last schedule failure : no error
Last failure time  : never
    
```

The following command provides more information on the execution cycle:

```
*A:4LS_CloudMG# show system script-control script-policy "test"
```

```
=====
Script-policy Information
=====
```

```

Script-policy      : test
Script-policy Owner : TiMOS CLI
Administrative status : enabled
Operational status : enabled
Script             : bring
Script owner       : TiMOS CLI
Script source location : cf3:/fetch.txt
Script results location : cf3:/results.txt
Max running allowed : 1
Max completed run histories : 1
Max lifetime allowed : 0d 01:00:00 (3600 seconds)
Completed run histories : 1
Executing run histories : 0
Initializing run histories : 0
Max time run history saved : 0d 01:00:00 (3600 seconds)
Script start error : N/A
Last change        : 2019/09/13 10:31:20 EEST
Max row expire time : never
Last application   : cron
Last auth. user account : not-specified
    
```

```
=====
Script Run History Status Information
-----
```

```
Script Run #25
-----
```

```

Start time      : 2019/09/13 11:29:25 EEST
End time        : 2019/09/13 11:29:26 EEST
Elapsed time    : 0d 00:00:01           Lifetime      : 0d 00:00:00
State           : terminated           Run exit code : noError
Result time     : 2019/09/13 11:29:26 EEST
Keep history    : 0d 00:59:41
Error time      : never
Results file    : cf3:/results.txt_20190913-082924-UTC.646420.out
Run exit        : Success
Error           : N/A
Application     : cron                 Auth. user ac*: not-specified
* indicates that the corresponding row element may have been truncated.
    
```

Every time a crontab job is executed, a log file is generated. The filename will be "results.txt_<timestamp>.out", where:

- results.txt: configured in the script-policy
- <timestamp>:file timestamp. As an example: 20190913-082324-UTC.646436.

An example log file will be: results.txt_20190913-082324-UTC.646436.out and the log file contents is as follows:

```
Pre-processing configuration file (V0v0)...
Completed processing 4 lines in 0.0 seconds
*A:4LS_CloudMG# file
*A:4LS_CloudMG# copy ftp://user:pwd@192.168.1.99/home/cmg/test/list.txt . force
Copying file ftp://user:pwd@192.168.1.99/home/cmg/test/list.txt ... OK
1 file copied.
*A:4LS_CloudMG# exit
*A:4LS_CloudMG# admin application-assurance group 1 url-list "test" upgrade
Executed 4 lines in 0.7 seconds from file cf3:\fetch.txt
```

Finally, using the following command, the operator may check when the list was last upgraded and verify that the cron job runs as intended:

```
# show application-assurance group 1 url-list "test"

=====
Application Assurance Group 1 url-list "test"
=====
Description          : (Not Specified)
Size                 : standard
Admin Status         : Up
Oper Status          : Up
Oper Flags           : <none>
File deployed to ISAs : Yes

-----
Upgrade Statistics
-----
Last Success          : 09/13/2019 11:49:57
  Deployed
  File Name           : cf3:\list.txt
  URL Entries         : 8 ( 0.05% full)
  URL Characters      : 103 (~0.00% full)
  Blank/Comment Lines : 0
Last Attempt          : 09/13/2019 11:49:57
  Result              : Success
  File Name           : cf3:\list.txt
=====
```

Conclusion

This chapter, intended for Application Assurance (AA) network architects and engineers, provides two examples for deploying URL-list filtering, upgrading the list and displaying its statistics, as well as configuring a cron job so that the system will periodically fetch a new URL list file and upgrade the list automatically.

Application Assurance — Security Gateway Stateful Firewall

This chapter provides information about Application Assurance Security gateway stateful firewall.

Topics in this chapter include:

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

Applicability

This chapter is applicable to all 7750 SR/SR-c and 7450 ESS chassis supporting Application Assurance (AA).

The configuration was tested on Release 13.0.R2.


Overview

The SR OS 13.0.R1 AA stateful firewall feature runs on AA-ISA and extends application-level analysis to provide an in-line stateful service, integrated within the Security Gateway (SeGW). The feature provides protection for mobile infrastructure; Mobility Management Entities (MMEs), Serving Gateways (SGWs), and Network Management Systems (NMSs), against attacks from compromised base stations, evolved NodeBs (eNBs), or Femto Access Points (FAPs). AA stateful packet filtering, combined with AA L7 classification and control, provides advanced, next-generation firewall functionality. Using stateful packet filtering, the AA FW not only inspects packets at layers 3 to 7, but also monitors the connection state.

AA FW deployed within a SeGW in ultra-broadband access networks (3G/4G/Femto) provides back-end core network security protection, as per Figure 1. AA FW offers protection for the following 3rd Generation Partnership Project (3GPP) defined interfaces:


1. S1-MME
2. S1-U
3. Operations, Administration and Maintenance (OAM)

Figure 30: LTE SeGW Firewall Deployment



Similarly, the SeGW architecture for Femto deployment is based on two 7750 SR systems terminating the mobile backhaul side (front-end and connecting to, for example, a 9365 Base Station Router Gateway (BSR-GW) and other network elements of the packet core (back-end), as per [Figure 31: SeGW in Small Cells Architecture](#):

Figure 31: SeGW in Small Cells Architecture



The two SeGWs run in stateful redundant mode: upon partial or total failure of the active SeGW for a set of IPsec tunnels, the other SeGW takes over without terminating the IPsec tunnels, providing hitless failover. In addition to MS-ISA hardware dedicated to the IPsec function, each SeGW supports one or more additional MS-ISAs running AA to provide firewall capabilities. The firewall rules protect the BSR as well as the BSR-GW and packet core network elements (NEs) from malicious attacks or unauthorized traffic.

The objective of this chapter is to describe the required configuration within AA-ISA in order to enable AA FW and protection for S1-MME, S1-U, and OAM traffic. Basic knowledge of AA-ISA diversion configuration is assumed.

S1-MME Traffic Protection

The purpose of AA FW in this deployment is to protect the MME infrastructure against an attack from a compromised eNB or FAP. Network flooding attacks, malformed packets, and port scans are examples of denial of service (DoS) attacks that can be carried out using a compromised eNB or FAP.

AA FW provides inspection of the Stream Control Transmission Protocol (SCTP) used to communicate to the MME. Such inspection includes checking for SCTP payload protocol IDs (PPIDs), source /destination ports, SCTP chunk validation, and malformed SCTP packets (such as checksum validation). In addition, the operator can configure DoS flooding rules, such as policers to limit the bandwidth and/or flow counts of SCTP traffic.

S1-U Traffic Protection

The purpose of AA FW in this deployment is to protect the SGW infrastructure against an attack from a compromised eNB or FAP. AA FW supports protection against:

- malformed GPRS Tunneling Protocol User plane (GTP-U) packet attacks
Checking packet sanity, which include GTP-U mandatory, optional, and extension header checks, as well as checks for invalid reserved information elements (IE) and missing IEs.
- unsupported GTP messages
Filtering messages based on message type and/or message length.
- flooding attacks
Shaping GTP traffic bandwidth, which limits the GTP-U bandwidth that a FAP can send to the core (SGW).
Limiting GTP tunnels, which limits the number of concurrent GTP tunnels and/or setup rate of these tunnels from a FAP to the core network.
To prevent the shared resources of bandwidth and the SGW processor from being consumed by an attacker, Nokia recommends the GTP flow rate limiting configuration.
- IP fragmentation-based attacks
Applying various drop rules for IP fragmentation of GTP messages.

OAM Traffic Protection

The purpose of AA FW protection in this deployment is to protect against any abuse of OAM network resources, such as NMS.

Network flooding attacks, malformed packets, and port scans are examples of such attacks that can be carried out using a compromised eNB or FAP.

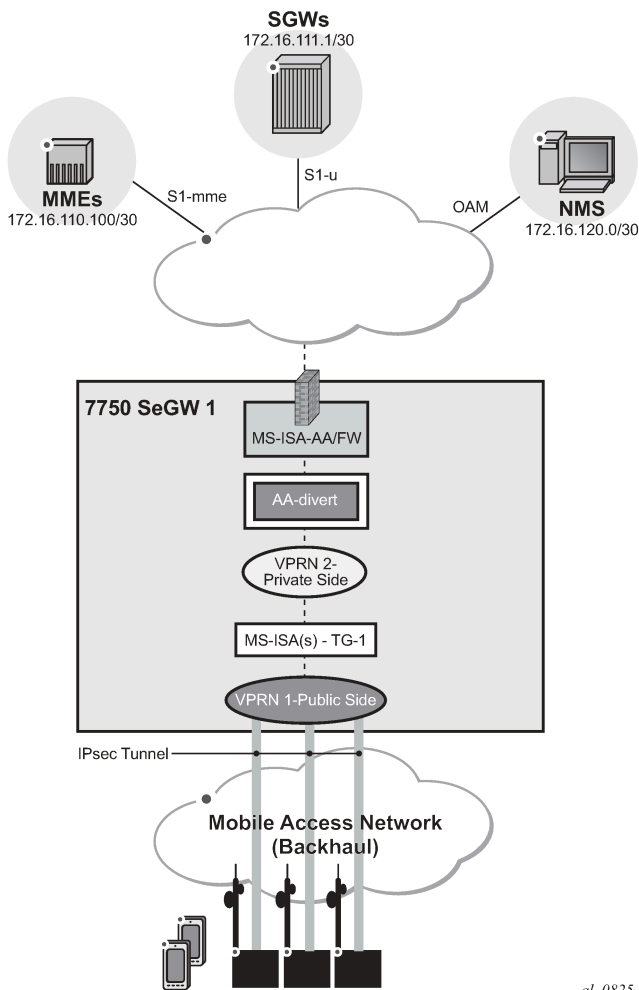
See the configuration described in the [Application Assurance — Stateful Firewall](#) chapter for this context of OAM protection in SeGW.

Configuration

AA-ISA Application QoS Policies (AQPs) are enhanced in Release 13.0.R1 with several new AQP actions that provide SCTP and GTP filtering functionality. As with all AQPs, these actions have partition-level scope, which allows different FW policies to be implemented by using AA partition concepts within the same AA-ISA.

The configuration topology in [Figure 32: Configuration Topology](#) shows the SeGW FW functionality of S1-U and S1-MME interfaces. Geo-redundancy, which is a very common deployment option, is not described in this here because it is described in the [Multi-Chassis IPsec Redundancy](#) chapter.

Figure 32: Configuration Topology



al_0825

Pre-Setup Requirements

Configure tunnel ISAs with optional multi-chassis redundancy. See the [Multi-Chassis IPsec Redundancy](#) chapter for more information.

1. Divert AA traffic and apply basic firewall rules.

Step 1.1. Divert private VPRN traffic into AA-ISA with AA multi-chassis redundancy.

This step is required for any of the configurations in Steps 2, 3 or 4.

There is no dependency between Steps 2, 3 or 4.

In this example, one private VPRN is used for all traffic to/from eNBs. In some small cell deployments, eNB traffic is split into three different VPRNs: one for control (S1-MME), one for management (OAM), and one for bearer traffic (S1-U GTP-U). In that case, each of these VPRNs needs to be diverted into AA-ISA in order to provide firewall protection.

First, define an application profile and transit IP policy, such as:

```
*A:7750-1>config>app-assure>group$ info
-----
    policy
      begin
      app-profile "default" create
        description " App profile that applies to the whole SAP"
        divert
      exit
      commit
    exit
  transit-ip-policy 1 create
    description "Per eNB-IP Sub policy"
    detect-seen-ip
    transit-auto-create
      no shutdown
    exit
```

Then, apply these policies to the SAP on the private side of the IPsec tunnel ISA:

```
*A:7750-1>config>service>vprn>if# info
-----
    sap tunnel-1.private:1 create
      transit-policy ip 1
      app-profile "default"
    exit
```

This configuration achieves:

- Traffic to/from the IPsec tunnel ISA private SAP is diverted to AA-ISA for the purpose of FW protection
- Within AA-ISA, the diverted SAP is treated as a parent SAP. That is, instead of treating the whole SAP as a single subscriber, subscribers are auto-created within this SAP based on the IP address of the eNBs

Step 1.2. Protect against malformed packets.

In firewall deployments, it is recommended that overload-drop, error-drop, and fragment-drop (all) are enabled within the default sub-policy, as shown in the example below:

- **overload-drop** ensures that AA-ISA, when overloaded, drops the excess traffic instead of allowing it through, without applying firewall rules.
- **error-drop** ensures that AA-ISA drops malformed IP packets.

- **fragment-drop (all)** because many network DoS attacks use IP fragmentation to initiate attacks, the operator has the option to drop all fragmented traffic, drop out-of-order fragments only, or allow fragments through. Allowing fragments through is not recommended for firewall deployments.

```
*A:7750-1>config>app-assure>group>policy# app-qos-policy
*A:7750-1>config>app-assure>group>policy>aqp# info
-----
        entry 500 create
            description "apply SeGW session filter rules"
            match
                traffic-direction subscriber-to-network
            exit
            action
                overload-drop
                error-drop
                fragment-drop all
            exit
            no shutdown
        exit
    exit
exit
-----
*A:7750-1>config>app-assure>group>policy#
```

Step 1.3. Limit total traffic from any eNB.

It is recommended that a total limit be placed on how much bandwidth and how many flows an eNB or FAP can generate toward the network, regardless of the type of traffic.

The limit values are a function of the number of end devices that are served by the eNB or FAP, plus some additional margin:

```
*A:7750-1>config>app-assure>group# info
-----
        policer "limit_eNBs_total_Flows" type flow-count-limit granularity subscriber
            create
            flow-count 1000
        exit
        policer "limit_eNBs_total_bw" type single-bucket-bandwidth granularity sub
            sscriber create
            rate 500
            mbs 500
        exit
-----
*A:7750-1>config>app-assure>group#
```



Note:

If the traffic from eNB or FAP is separated into different private SAPs, based on traffic type (S1-AP, S1-U, or OAM), as with some deployment topologies, then the policing limit value is dependent on the SAP traffic type as well as the number of end devices. See policing limit settings in Steps 2 and 3.

Apply the configured policers as actions from within the default sub-policy AQP entry:

```
*A:7750-1>config>app-assure>group>policy# app-qos-policy entry 500
*A:7750-1>config>app-assure>group>policy>aqp>entry>action# flow-count-limit "limit_eNBs_
total_Flows"
*A:7750-1>config>app-assure>group>policy>aqp>entry>action# bandwidth-policer "limit_eNBs_
total_bw"
*A:7750-1>config>app-assure>group>policy>aqp>entry# info
```

```

-----
description "apply SeGW session filter rules"
match
    traffic-direction subscriber-to-network
exit
action
    bandwidth-policer "limit_eNBs_total_bw"
    flow-count-limit "limit_eNBs_total_Flows"
    session-filter "SeGW_FW"
    overload-drop
    error-drop
    fragment-drop all
exit
no shutdown
-----
*A:7750-1>config>app-assure>group>policy>aqp>entry#

```

**Note:**

All of the above listed actions use the traffic direction of subscriber-to-network. That is, they are not applied to traffic in the other direction (downstream) because the purpose of the firewall is to protect the network resources from upstream traffic coming from compromised eNBs or FAPs.

2. Configure AA-ISA to provide firewall protection to protect MMEs (S1-AP traffic).

Step 2.1. Create IP AA lists.

First, create an AA IP prefix list that contains eNB IP addresses or range of addresses:

```

*A:7750-1>config>app-assure# group 1:1
*A:7750-1>config>app-assure>group# ip-prefix-list "ALL_eNBs" create
*A:7750-1>config>app-assure>group>pfx>$ description "eNodeB subnet"
*A:7750-1>config>app-assure>group>pfx>$ prefix 172.16.100.0/24
*A:7750-1>config>app-assure>group>pfx>$ exit

```

Next, optionally create an AA IP list that contains MME IP addresses (in case there are more than one):

```

*A:7750-1>config>app-assure>group# ip-prefix-list "MMEs" create
*A:7750-1>config>app-assure>group>pfx>$ description "MME(s) subnet"
*A:7750-1>config>app-assure>group>pfx>$ prefix 172.16.110.100/30
*A:7750-1>config>app-assure>group>pfx>$ exit

```

After the above lists are created, they can be referenced and used in AA FW rules using session filters and AQPs.

Step 2.2. Allow only SCTP traffic towards MMEs — No port scanning.

A basic setup creates session-filter rules that will only allow SCTP traffic between eNBs and MMEs.

```

*A:7750-1>config>app-assure>group# session-filter "SeGW_FW" create
*A:7750-1>config>app-assure>group>sess-fltr$ default-action deny
*A:7750-1>config>app-assure>group>sess-fltr$ entry 1 create
*A:7750-1>config>app-assure>group>sess-fltr>entry$ description "allow SCTP to MM Es"
*A:7750-1>config>app-assure>group>sess-fltr>entry$ match
*A:7750-1>config>app-assure>group>sess-fltr>entry>match$ ip-protocol-num "sctp"
*A:7750-1>config>app-assure>group>sess-fltr>entry>match$ src-ip ip-prefix-list "ALL_eNBs"
*A:7750-1>config>app-assure>group>sess-fltr>entry>match$ dst-ip ip-prefix-list "MMEs"
*A:7750-1>config>app-assure>group>sess-fltr>entry>match$ dst-port eq 6005
*A:7750-1>config>app-assure>group>sess-fltr>entry>match$ exit
*A:7750-1>config>app-assure>group>sess-fltr>entry$ action permit

```

```
*A:7750-1>config>app-assure>group>sess-fltr>entry$ exit
```

**Note:**

In the above configuration, SCTP traffic on MMEs is assumed to be running on port 6005.

Next, the newly created session filter needs to be referenced from a default sub-policy AQP action, as follows:

```
*A:7750-1>config>app-assure>group>policy>aqp# info
-----
      entry 500 create
        description "apply SeGW session filter rules"
        match
          traffic-direction subscriber-to-network
        exit
        action
          session-filter "SeGW_FW"
          overload-drop
          error-drop
          fragment-drop all
        exit
        no shutdown
      exit
    exit
  exit
-----
*A:7750-1>config>app-assure>group>policy#
```

Using traffic direction **subscriber-to-network** in the above AQP entry achieves two objectives:

- a. Protects MMEs by allowing only SCTP traffic to be initiated from eNB subnets toward MMEs. Port scanning toward MME is blocked.
- b. Allows MMEs to have full access to eNBs.

**Note:**

It is important that an AQP, containing a session-filter action, does not contain any matching condition other than ASOs, traffic direction, or subscriber ID. Subscriber ID is not applicable in this deployment use-case.

Step 2.3. DoS protection: Limit the number of SCTP flows from eNBs.

In this step, the operator configures a flow count policer to limit the number of SCTP flows that an eNB can generate toward the MMEs. This protects the MMEs against a compromised eNB trying to set up many SCTP flows.

```
*A:7750-1# configure application-assurance group 1
*A:7750-1>config>app-assure>group# policer "sctp_flow_count" type flow-count-limit
granularity subscriber create
*A:7750-1>config>app-assure>group>policer$ flow-count 2
*A:7750-1>config>app-assure>group>policer$ exit
```

In the above configuration, an eNB or FAP can have up to two flows at a time. In practice, there should only be one SCTP session, one flow in each direction, per eNB-MME pair. The above example uses two flows to leave a margin in case a second, backup, MME needs to communicate with the eNB, while still providing enough protection.

Add the defined policer as a **flow-count-limit** as an AQP action, as follows:

```
A:7750-1>config>app-assure>group>policy>aqp# entry 100
```

```
A:7750-1>config>app-assure>group>policy>aqp>entry$ info
-----
description "limit SCTP traffic"
match
    traffic-direction subscriber-to-network
    ip-protocol-num eq sctp
exit
action
    flow-count-limit "sctp_flow_count"
exit
no shutdown
-----
A:7750-1>config>app-assure>group>policy>aqp>entry$
```

Step 2.3.1. Configure an AA FW events log.

It is sometimes advisable to configure a log that captures events related to various AA FW actions. Due to the limited size of the log and the large amount of traffic AA can handle, consider the usefulness of the information in the log when:

- debugging a configuration
- testing a configuration in a staged environment
- capturing infrequent actions

To configure a log:

```
*A:7750-1# configure application-assurance group 1:1
*A:7750-1>config>app-assure>group# event-log "FW_drops_log" create
*A:7750-1>config>app-assure>group>evt-log$ buffer-type circular
*A:7750-1>config>app-assure>group>evt-log$ max-entries 100000
*A:7750-1>config>app-assure>group>evt-log$ no shutdown
*A:7750-1>config>app-assure>group>evt-log$ exit
*A:7750-1>config>app-assure>group# info
-----
---snip---
event-log "FW_drops_log" create
buffer-type circular
max-entries 100000
no shutdown
exit
```

To reference the configured log from within the deny action of the session filter:

```
*A:7750-1>config>app-assure>group>sess-fltr# info
-----
default-action deny event-log "FW_drops_log"
entry 1 create
description "allow SCTP to MMEs"
match
    ip-protocol-num sctp
    src-ip ip-prefix-list "ALL_eNBs"
    dst-ip ip-prefix-list "MMEs"
exit
action permit
exit
-----
*A:7750-1>config>app-assure>group>sess-fltr#
```

To view the log:

```
*A:7750-1# tools dump application-assurance group 1:1 event-log "FW_drops_log" isa 1/2
=====
Application-Assurance event-log "FW_drops_log"
Current Time:          "06/10/2015 22:45:30" (UTC)
  group[:partition]:   1:1
  isa:                 1/2
  admin state:         no shutdown
  buffer-type:         circular
  max-entries:         100000
=====
Event-source
Src-ip          Action          SubType      SubName      Direction
Src-port Dst-port Timestamp      Dst-ip
Total Records:   0
=====
*A:7750-1#
```

To clear all the entries within the specified log:

```
*A:7750-1# clear application-assurance group 1:1 event-log "FW_drops_log"
```

Step 2.4. DoS protection: Limit the SCTP bandwidth from eNB

Similar to the previous step, the operator configures a flow bandwidth policer to limit the amount of SCTP traffic that an eNB can generate toward the MMEs. This protects the MMEs against a compromised eNB trying to flood the MMEs.

```
*A:7750-1# configure application-assurance group 1
*A:7750-1>config>app-assure>group# info
-----
---snip---
      policer "sctp_bw_limit" type single-bucket-bandwidth granularity subscriber
create
      rate 30
      mbs 10
      exit
---snip---
      exit
-----
*A:7750-1>config>app-assure>group#
```

In the above example, a single leaky-bucket policer is configured with a rate set to 30 kb/s and maximum burst size of 10 kbytes. This provides enough bandwidth to ensure normal operations, while still providing a ceiling limit of how much traffic any eNB can send toward the MMEs.

The value for this policer is a function of the amount of user equipment (UEs) served by the eNB/FAP. For example, in a small cell deployment, with 32 active users per 9962 FAP, the S1-MME bandwidth is estimated to be:

Uplink — toward MME : 2.7 kb/s

Downlink — from MME toward FAP : 28 kb/s

Add the defined policer as a subscriber policy, as follows:

```
A:7750-1>config>app-assure>group>policy>aqp# entry 100
```



```
A:7750-1>config>app-assure>group>policy>aqp>entry$ info
-----
description "limit SCTP traffic"
match
  traffic-direction subscriber-to-network
  ip-protocol-num eq sctp
exit
action
  bandwidth-policer "sctp_bw_limit"
  flow-count-limit "sctp_flow_count"
exit
no shutdown
-----
A:7750-1>config>app-assure>group>policy>aqp>entry$
```

Step 2.4.1 Configure additional limits for all traffic to MMEs.

To further protect the MMEs from a distributed attack, whereby a number of eNBs or FAPs are compromised, an AA FW can be configured to limit total traffic, not just from a single eNB as outlined in previous sections, but from all eNBs toward the MMEs.

It is recommended to configure the following three protection limits:

- a. total bandwidth of SCTP toward MMEs
- b. total number of flows toward MMEs
- c. flow setup rate toward the MMEs

The configuration is shown below:

```
*A:7750-1>config>app-assure>group# info
-----
create policer "limit_total_sctp_bw" type single-bucket-bandwidth granularity system
      rate 1200
      mbs 100
      exit
      policer "limit_total_sctp_flows" type flow-count-limit granularity system create
      flow-count 400
      exit
      policer "limit_total_sctp_flows_rate" type flow-rate-limit granularity system
create
      rate 100
      mbs 100
      exit
-----
*A:7750-1>config>app-assure>group#
```



Note:

- The policers are of type **system** and not **subscriber** in order to be applied to all eNBs or FAPs, as is the case when auto-transit subscribers are created (see Step 1).
- The actual limits of these policers are a function of the total number of eNBs served by the SeGW. In the above configuration, it is assumed that there are 400 eNBs. Therefore, the total limit is 400 flows of SCTP traffic.
- A flow setup rate limit of 100 is set to protect MMEs from a storm of new SCTP flows.

The policers are then referenced from within the appropriate AQP entry that matches the MMEs traffic and SCTP:

```
*A:7750-1>config>app-assure>group>policy>aqp# info
-----
---snip---
        entry 110 create
            description " limit system traffic towards MMEs"
            match
                traffic-direction subscriber-to-network
                src-ip eq ip-prefix-list "ALL_eNBs"
                dst-ip eq ip-prefix-list "MMEs"
            exit
            action
                bandwidth-policer "limit_total_sctp_bw"
                flow-rate-limit "limit_total_sctp_flows_rate"
                flow-count-limit "limit_total_sctp_flows"
            exit
            no shutdown
        exit
*A:7750-1>config>app-assure>group>policy>aqp#
```



Note:

It is possible, but redundant, to add the **ip-protocol eq sctp** command as a match condition, because the configured session filter already ensures that only SCTP traffic can flow between eNBs and MMEs.

Step 2.5. Allow only specified SCTP PPIDs toward the MMEs.

In this step, the operator blocks all except the specified SCTP messages that contain configured PPIDs, using an AA SCTP filter configuration:

```
*A:7750-1>config>app-assure>group# sctp-filter
- no sctp-filter <sctp-filter-name>
- sctp-filter <sctp-filter-name> [create]

<sctp-filter-name> : [32 chars max]
<create>          : keyword

[no] description  - Configure a description of the SCTP filter
[no] event-log    - Configure an event log for packets dropped by the SCTP
                  filter
      ppid         + Configure actions for specific or default PPIDs
                  (Payload Protocol Identifiers)
[no] ppid-range   - Configure the range of allowable PPIDs for the SCTP
                  filter
```

The filter specifies either a range of PPIDs or individual PPIDs:

```
*A:7750-1>config>app-assure>group>sctp-fltr>ppid$ entry 1
- entry <entry-id> value <ppid-value> action {permit|deny}
- no entry <entry-id>

<entry-id>       : [1..255]
<ppid-value>     : [0..4294967295]D | [256 chars max]
<permit|deny>    : permit|deny
```

The PPIDs can be specified either by their values or by names. Names are specified in RFC 4960. See [Table 9: SCTP PPIDs](#).

Table 9: SCTP PPIDs

| Value | SCTP PPID | Value | SCTP PPID |
|-------|---|-------|--|
| 0 | Reserved by SCTP | 31 | Service Area Broadcast Protocol (SABP) |
| 1 | IUA | 32 | Fractal Generator Protocol (FGP) |
| 2 | M2UA | 33 | Ping Pong Protocol (PPP) |
| 3 | M3UA | 34 | CalcApp Protocol (CALCAPP) |
| 4 | SUA | 35 | Scripting Service Protocol (SSP) |
| 5 | M2PA | 36 | NetPerfMeter Protocol Control Channel (NPMP-CONTROL) |
| 6 | V5UA | 37 | NetPerfMeter Protocol Data Channel (NPMP-DATA) |
| 7 | H.248 | 38 | Echo (ECHO) |
| 8 | BICC/Q.2150.3 | 39 | Discard (DISCARD) |
| 9 | TALI | 40 | Daytime (DAYTIME) |
| 10 | DUA | 41 | Character Generator (CHARGEN) |
| 11 | ASAP | 42 | 3GPP RNA |
| 12 | ENRP | 43 | 3GPP M2AP |
| 13 | H.323 | 44 | 3GPP M3AP |
| 14 | Q.IPC/Q.2150.3 | 45 | SSH over SCTP |
| 15 | SIMCO <draft-kiesel-midcom-simco-sctp-00.txt> | 46 | Diameter in a SCTP DATA chunk |
| 16 | DDP Segment Chunk | 47 | Diameter in a DTLS/SCTP DATA chunk |
| 17 | DDP Stream Session Control | 48 | R14P. BER Encoded ASN.1 over SCTP |
| 18 | S1 Application Protocol (S1AP) | 49 | Unassigned |
| 19 | RUA | 50 | WebRTC DCEP |
| 20 | HNBAP | 51 | WebRTC String |
| 21 | ForCES-HP | 52 | WebRTC Binary Partial (deprecated) |
| 22 | ForCES-MP | 53 | WebRTC Binary |
| 23 | ForCES-LP | 54 | WebRTC String Partial (deprecated) |

| Value | SCTP PPID | Value | SCTP PPID |
|-------|---|---------------|---------------------|
| 24 | SBc-AP | 55 | 3GPP PUA |
| 25 | NBAP | 56 | WebRTC String Empty |
| 26 | Unassigned | 57 | WebRTC Binary Empty |
| 27 | X2AP | 58-4294967295 | Unassigned |
| 28 | IRCP - Inter Router Capability Protocol | | |
| 29 | LCS-AP | | |
| 30 | MPICH2 | | |

It is recommended to limit the SCTP traffic to only those packets with S1 AP PPID. The SCTP filter can be configured to deny all by default and only allow PPID S1 AP (by value = 18 or by name: *s1-application-protocol*) as follows:

```
*A:7750-1# configure application-assurance group 1:1
  sctp-filter "SCTP-PPID-Filter" create
  description "Allow only S1AP PPID"
  event-log "FW_drops_log"
  ppid
  default-action deny
  entry 1 value "s1-application-protocol" action permit
  exit
```

This configured SCTP filter is then referenced as an action from within an AQP entry:

```
*A:7750-1>config>app-assure>group>policy>aqp# info
-----
  entry 100 create
  description "limit SCTP traffic"
  match
    traffic-direction subscriber-to-network
    ip-protocol-num eq sctp
  exit
  action
    bandwidth-policer "sctp_bw_limit"
    flow-count-limit "sctp_flow_count"
    sctp-filter "SCTP-PPID-Filter"
  exit
  no shutdown
  exit
-----
A:7750-1>config>app-assure>group>policy>aqp#
```

To view the packets allowed or denied by the configured SCTP filter:

```
*A:7750-1# show application-assurance group 1:1 sctp-filter "SCTP-PPID-Filter"
=====
Application Assurance Group 1:1 SCTP Filter "SCTP-PPID-Filter"
=====
Description      : Allow only S1AP PPID
Maximum PPID     : 4294967295
```

```

Minimum PPID      : 0
Default action    : deny
Configured PPIDs  : 1

Packets arrived   : 0
Packets denied
  Malformed packet : 0
  PPID out of range : 0
  PPID denied      : 0
Packets permitted : 0
=====
*A:7750-1#
    
```



Note:

The SCTP malformed packet counter shown above increments when an AA SCTP filter encounters an SCTP packet that is malformed, such as:

- IP packet is too small to contain a common SCTP header
- SCTP chunk LEN < 4 bytes: each SCTP chunk header is 4 bytes, so the SCTP chunk cannot be smaller than this
- remaining space in the IP packet is too small to contain a chunk header (for example, your packet has 2 chunks and the 2nd chunk length goes beyond the IP length advertised)
- IP packet is too small to contain the chunk

Currently, the SCTP filter statistics cannot be reset on the fly without shutting down the SCTP filter.

Another way to view the effect of the configured SCTP filter is to check the firewall log, if configured:

```
*A:7750-1# tools dump application-assurance group 1:1 event-log "FW_drops_log" isa 1/2
```

3. Configure AA-ISA to protect SGW (GTP-U traffic).

The steps to configure the AA-ISA in an SeGW to protect against attacks toward the SGW are similar to the steps for SCTP traffic. While GTP filtering is very different from SCTP filtering, configuration to limit the flow counts, bandwidth, and session filter are similar.

Step 3.1. Create an AA IP list for SGWs.

In addition to the lists configured in step 2.1, the operator can optionally configure a list that contains the SGW IP addresses that are served by the SeGW, in case there is more than one.

```

*A:7750-1# configure application-assurance group 1:1 ip-prefix-list "SGWs" create
*A:7750-1>config>app-assure>group>pfx>$ description "Serving Gateways IPs"
*A:7750-1>config>app-assure>group>pfx>$ prefix 172.16.111.1/32
*A:7750-1>config>app-assure>group>pfx>$ prefix 172.16.111.2/32
*A:7750-1>config>app-assure>group>pfx>$ exit
    
```

Step 3.2. Allow only GTP-U traffic toward SGWs — No port scanning.

Similar to Step 2.2, create an GTP filter to allow only GTP traffic to/from eNBs to SGWs:

```

*A:7750-1>config>app-assure>group>sess-fltr# info
-----
      default-action deny event-log "FW_drops_log"
      ---snip---
      entry 2 create
        description "allow GTP-u to SGWs"
        match
          ip-protocol-num udp
    
```

```

src-ip ip-prefix-list "ALL_eNBs"
dst-ip ip-prefix-list "SGWs"
dst-port eq 2152
exit
action permit
exit
-----
*A:7750-1>config>app-assure>group>sess-fltr#

```

The following session filter needs to be added to the default sub-policy AQP, similar to Step 2.2:

```

*A:7750-1>config>app-assure>group>policy>aqp# info
-----
entry 500 create
description "apply SeGW session filter rules"
match
traffic-direction subscriber-to-network
exit
action
session-filter "SeGW_FW"
overload-drop
error-drop
fragment-drop all
exit
no shutdown
exit
exit
-----

```

For AA to recognize GTP traffic and perform sanity packet checking, configure a GTP filter at the group:partition level:

```

*A:7750-1# configure application-assurance group 1:1 gtp no shutdown

```

Step 3.3. DoS protection — Limit the number of GTP-U flows from eNBs.

AA can be configured to limit the number of GTP flows from an eNB. A GTP-U flow is defined by GTP-U packet destination IP + tunnel ID (TEID).

AA allows the operator to configure two limits: one that applies to the each eNB and one that applies for all GTP-U traffic from all eNBs:

```

*A:7750-1>config>app-assure>group# info
-----
policer "GTPu-Flow-count-limit" type flow-count-limit granularity subscriber
create
flow-count 800
gtp-traffic
exit

```

The actual value of the flow count limit is a function of the number of UEs or devices served by an eNB or FAP. In the above case, it is assumed that there are 100 devices with a maximum of 8 GTP-U flows per device. For FAP, the number is typically around 32 devices per FAP. Note: By 3GPP standards, the maximum number of GTP-U tunnels per device is 16.

Assuming that there are 1000 eNBs or FAPs that are served by the SeGW, then to limit the total number of GTP-U flows, the operator can apply the following system policer:

```

*A:7750-1>config>app-assure>group# info
-----

```

```

create      policer "limit_total_GTPU_Flow_count" type flow-count-limit granularity system
            flow-count 800000
            gtp-traffic
            exit
    
```

Configure AQPs to execute the policers:

```

*A:7750-1>config>app-assure>group>policy>aqp# info
-----
---snip---
        entry 120 create
            description "limit GTP-U traffic"
            match
                traffic-direction subscriber-to-network
            exit
            action
                flow-count-limit "GTPu-Flow-count-limit"
            exit
            no shutdown
        exit
        entry 130 create
            description "limit TOTAL GTPU towards SGWs"
            match
                traffic-direction subscriber-to-network
            exit
            action
                flow-count-limit "limit_total_GTPU_Flow_count"
            exit
            no shutdown
        exit
    
```

For GTP-U flow count policing, it is important that **aqp-initial-lockup** is enabled:

```

*A:7750-1# configure application-assurance group 1:1 aqp-initial-lockup
    
```

The above configured limits are applied only to upstream traffic, to protect the network. No limit is placed on the downstream traffic toward the eNBs.



Note:

For small cell deployments, the number of GTP-U tunnels per FAP is a function of:

- a. deployment mode:
 - i. residential = 32 (9962 MSEC-MS-MCI Enterprise) UEs
 - ii. enterprise = 8 (9961 MSHC) UEs.
- b. number of guaranteed bit rate (GBR) tunnels (max 8) and non-GBR tunnels (max 8) per UE.

Therefore, the GTP-U tunnel limit per FAP should be set to $32 \times 8 = 256$ for residential deployments or $8 \times 8 = 64$ for enterprise deployments.

The operator can view the effect of the configured policers on GTP traffic by running the following show routine:

```

*A:7750-1>show>app-assure>group# gtp
=====
Application Assurance Group 1:1 GTP
=====
    
```

```

Admin status : Up
Event log    : (Not Specified)
-----
GTP Statistics                               sub-to-net          net-to-sub
-----
Incoming packets                            0                   0
Packets denied
  UDP packet length                          0                   0
  GTP message length                         0                   0
  GTP version                                0                   0
-----
Packets permitted                            0                   0
-----
GTP Policing Statistics                     sub-to-net          net-to-sub
-----
Packets arrived                              0                   0
Packets denied
  gtp-traffic flow-count policer              0                   0
  Other                                        0                   0
-----
Packets permitted                            0                   0
-----
GTP Filter Statistics                       sub-to-net          net-to-sub
-----
Packets arrived                              0                   0
Packets denied (gtp-filter)                  0                   0
Packets permitted
  gtp-filter                                  0                   0
  no gtp-filter                               0                   0
-----
Total GTP packets permitted                  0                   0
=====
*A:7750-1>show>app-assure>group#

```

In the last section shown above, GTP filter statistics are related to GTP filters that are discussed and configured later in Step 3.5 of this chapter.

Step 3.4. DoS protection: Limit the GTP-U bandwidth from eNBs.

This step is similar to Step 3.3, but instead of configuring a flow count policer, the operator configures bandwidth policers:

```

*A:7750-1>config>app-assure>group# info
-----
create      policer "GTPU_bw_limit" type single-bucket-bandwidth granularity subscriber
            rate 5000
            mbs 100
            exit

create      policer "limit_total_GTPU_bw" type single-bucket-bandwidth granularity system
            rate 2000000
            mbs 2000
            exit

*A:7750-1>config>app-assure>group>policy>aqp# info
-----
---snip---
            entry 120 create
            description "limit GTP-U traffic"

```



```

match
    traffic-direction subscriber-to-network
exit
action
    bandwidth-policer "GTPU_bw_limit"
    flow-count-limit "GTPu-Flow-count-limit"
exit
no shutdown
exit
entry 130 create
description "limit TOTAL GTPU towards SGWs"
match
    traffic-direction subscriber-to-network
exit
action
    bandwidth-policer "limit_total_GTPU_bw"
    flow-count-limit "limit_total_GTPU_Flow_count"
exit
no shutdown
exit

```

The above configured limits are applied only to upstream traffic, to protect the network. No limit is placed on downstream traffic toward the eNB.

As a debugging tool, the operator can use the AA **flow-record-search** command to check the status of GTP flows through the ISA:

```

*A:7750-1# tools dump application-assurance group 1:1 flow-record-search isa 1/2 flow-
status active protocol "gtp"
=====
Application-Assurance flow record search, Version 1.0
Search Start Time: "06/16/2015 20:38:09" (UTC)
Search Criteria:
  group[:partition]: 1:1
  isa: 1/2
  protocol name: "gtp"
  application name: none specified
  app-group name: none specified
  flow-status: active
  start-flowId: none specified
  classified: none specified
  server-ip: none specified
  server-port: none specified
  client-ip: none specified
  bytes-tx: none specified
  flow-duration: none specified
  max-count: none specified
  search-type: default
=====
FlowId  Init  Src-ip                               Dst-ip
      Ip-prot  Src-prt Dst-prt Protocol           Application           Pkts-tx
Bytes-tx          Pkts-disc Bytes-disc Time-ofp(UTC)       Time-olp(UTC)
SEARCH COMPLETED.
Search End Time: "06/16/2015 20:38:09" (UTC)
Total Records: 0
=====
*A:7750-1#

```

GTP flows that are to be denied by the previous AA configurations should not appear in the search results.

Step 3.5. Further GTP filtering and validation.

AA allows the operator to configure a GTP filter to enforce which GTP message types are allowed/denied, as well as the maximum allowed GTP message length:

```
*A:7750-1>config>app-assure>group>gtp>gtp-fltr#
[no] description      - Configure a description of the GTP filter
[no] event-log        - Configure an event log for packets dropped by the GTP filter
[no] max-payload-le*  - Configure the maximum payload length of the GTP filter
      message-type    + Configure actions for specific or default messages

*A:7750-1>config>app-assure>group>gtp>gtp-fltr#
```



Note:

An AA GTP filter allows the operator to configure a maximum payload size for the GTP traffic. However, in this configuration example, no maximum payload size is configured.

The list of GTP message types are defined by 3GPP standard 3GPP TS 29.281 as per [Table 10: GTP Messages](#).

Table 10: GTP Messages

| Message Type Value (Decimal) | Message | Message Type Value (Decimal) | Message |
|------------------------------|---|------------------------------|---|
| 1 | "echo-request" | 55 | "forward-relocation-complete" |
| 2 | "echo-response" | 56 | "relocation-cancel-request" |
| 3 | "version-not-supported:" | 57 | "relocation-cancel-response" |
| 4 | "node-alive-request" | 58 | "forward-sms-context" |
| 5 | "node-alive-response" | 59 | "forward-relocation-complete-acknowledge" |
| 6 | "redirection-request" | 60 | "forward-sms-context-acknowledge" |
| 7 | "redirection-response" | 70 | "ran-information-relay" |
| 16 | "create-pdp-context-request" | 96 | "mbms-notification-request" |
| 17 | "create-pdp-context-response" | 97 | "mbms-notification-response" |
| 18 | "update-pdp-context-request" | 98 | "mbms-notification-reject-request" |
| 19 | "update-pdp-context-response" | 99 | "mbms-notification-reject-response" |
| 20 | "delete-pdp-context-request" | 100 | "create-mbms-context-request" |
| 21 | "delete-pdp-context-response" | 101 | "create-mbms-context-response" |
| 22 | "initiate-pdp-context-activation-request" | 102 | "update-mbms-context-request" |

| Message Type Value (Decimal) | Message | Message Type Value (Decimal) | Message |
|------------------------------|--|------------------------------|--|
| 23 | "initiate-pdp-context-activation-response" | 103 | "update-mbms-context-response" |
| 26 | "error-indication" | 104 | "delete-mbms-context-request" |
| 27 | "pdu-notification-request" | 105 | "delete-mbms-context-response" |
| 28 | "pdu-notification-response" | 112 | "mbms-registration-request" |
| 29 | "pdu-notification-reject-request" | 113 | "mbms-registration-response" |
| 30 | "pdu-notification-reject-response" | 114 | "mbms-de-registration-request" |
| 31 | "supported-extension-headers-notification" | 115 | "mbms-de-registration-response" |
| 32 | "send-routing-information-for-gprs-request" | 116 | "mbms-session-start-request" |
| 33 | "send-routing-information-for-gprs-response" | 117 | "mbms-session-start-response" |
| 34 | "Failure-report-request" | 118 | "mbms-session-stop-request" |
| 35 | "failure-report-request" | 119 | "mbms-session-stop-response" |
| 36 | "note-ms-gprs-present-request" | 120 | "mbms-session-update-request" |
| 37 | "note-ms-gprs-present-response" | 121 | "mbms-session-update-response" |
| 48 | "identification-request" | 128 | "ms-info-change-notification-request" |
| 49 | "identification-response" | 129 | "ms-info-change-notification-response" |
| 50 | "sgsn-context-response" | 240 | "data-record-transfer-request" |
| 51 | "sgsn-context-request" | 241 | "data-record-transfer-response" |
| 52 | "sgsn-context-acknowledge" | 254 | "end-marker" |
| 53 | "forward-relocation-request" | 255 | "g-pdu" |
| 54 | "forward-relocation-response" | | |

Of the 67 GTP message types shown above, only 6 are allowed, by the standards, for GTP-U:

echo-request echo-response error-indication
g-pdu end-marker supported-extension-headers-notification

If these message types are permitted by the configured GTP filter, AA performs extensive GTP-U header checking on these six types.



Note:

If no GTP filter is configured, no extensive GTP-U header checks are performed. For example, if the operator wants to allow all GTP-U packets and perform all GTP header sanity checks, then a GTP filter that permits all message types needs to be configured, with the default action of permit and with no values, such as:

```
gtp-filter "allow-all" create
  message-type
  default-action permit
```

Because AA FW in an SeGW is protecting an S1-U interface running GTP-U, the GTP filter only needs to allow the six GTP messages that are permitted for GTP-U:

```
*A:7750-1>config>app-assure>group>gtp# info
-----
      gtp-filter "filter-gtp-msgs" create
      description "allow only certain msg types"
      message-type
      default-action deny
      entry 1 value "echo-request" action permit
      entry 2 value "echo-response" action permit
      entry 3 value "error-indication" action permit
      entry 4 value "supported-extension-headers-notification" action
permit
      entry 5 value "end-marker" action permit
      entry 6 value "g-pdu" action permit
      exit
    exit
  no shutdown
-----
*A:7750-1>config>app-assure>group>gtp#
```

This GTP filter is then referenced from within an AQP entry action, as follows, in order for it to take effect:

```
*A:7750-1>config>app-assure>group>policy>aqp# info
-----
      entry 120 create
      description "limit GTP-U traffic"
      match
        traffic-direction subscriber-to-network
        dst-ip eq ip-prefix-list "SGWs"
      exit
      action
        bandwidth-policer "GTPU_bw_limit"
        flow-count-limit "GTPu-Flow-count-limit"
        gtp-filter "filter-gtp-msgs"
      exit
    no shutdown
  exit
-----
```

The operator can view the effect of the configured GTP filter on S1-U traffic using the following show routine:

```
*A:7750-1>show>app-assure>group# gtp gtp-filter "filter-gtp-msgs"
=====
Application Assurance Group 1:1 GTP Filter "filter-gtp-msgs"
=====
```

```

Description      : allow only certain msg types
Maximum payload length : (Not Specified)
Default action   : deny
Configured messages : 6

Packets arrived   : 0
Packets denied
  Payload length  : 0
  Message type    : 0
  Mandatory header : 0
  Extension header : 0
  Information element : 0
Packets permitted : 0
=====
*A:7750-1>show>app-assure>group#
    
```

The above output is in addition to the information provided by the overall GTP show command:

```
*A:7750-1>show>app-assure>group# gtp
```

4. Configure AA-ISA to protect NMS (OAM Traffic).

Step 4.1. Create an IP AA list that contains the NMS server IPs.

```

*A:7750-1# configure application-assurance group 1:1
*A:7750-1>config>app-assure>group# info
-----
      ip-prefix-list "NMSs" create
      description "Network Management-OAM subnet"
      prefix 172.16.120.0/30
      exit
    
```



Note:

In the case of small cell deployments, different NMS servers need to be configured.

Step 4.2. Allow eNBs to initiate FTP- and ICMP-only traffic toward NMS, block port scanning.

```

*A:7750-1>config>app-assure>group>sess-fltr# info
-----
      default-action deny event-log "FW_drops_log"

      entry 3 create
      description "allow FTP to NMS"
      match
        ip-protocol-num tcp
        src-ip ip-prefix-list "ALL_eNBs"
        dst-ip ip-prefix-list "NMSs"
        dst-port eq 22
      exit
      action permit
    exit
  entry 4 create
  description "allow ICMP to NMS"
  match
    ip-protocol-num icmp
    src-ip ip-prefix-list "ALL_eNBs"
    dst-ip ip-prefix-list "NMSs"
  exit
  action permit
  exit
-----
    
```

```
*A:7750-1>config>app-assure>group>sess-fltr#
```

The operator can view the effect of the session filter on traffic, in terms of how many times it is applied, using the following show routine:

```
*A:7750-1>show>app-assure>group# session-filter
=====
AA Session Filter Table
=====
Name                               Default Action   Referenced       Entries
-----
SeGW_FW                             deny             aqp               4
-----
No. of session filters: 1
=====
*A:7750-1>show>app-assure>group#
*A:7750-1>show>app-assure>group# session-filter "SeGW_FW"
=====
AA Session Filter Instance "SeGW_FW"
=====
Description      : (Not Specified)
Default Action   : deny
  Event Log      : FW_drops_log
AQP Entries      :
  500
-----
Filter Match Criteria
-----
Entry           : 1
Description     : allow SCTP to MMEs
IP Protocol     : sctp
Source IP List  : ALL_eNBs
Dest IP List    : MMEs
Action          : permit
  Event Log     : (Not Specified)
Hits           : 0 flows
-----
Entry           : 2
Description     : allow GTP-u to SGWs
IP Protocol     : udp
Source IP List  : ALL_eNBs
Dest IP List    : SGWs
Dest Port      : eq 2152
Action          : permit
  Event Log     : (Not Specified)
Hits           : 0 flows
-----
Entry           : 3
Description     : allow FTP to NMS
IP Protocol     : tcp
Source IP List  : ALL_eNBs
Dest IP List    : NMSs
Dest Port      : eq 22
Action          : permit
  Event Log     : (Not Specified)
Hits           : 0 flows
-----
Entry           : 4
Description     : allow ICMP to NMS
IP Protocol     : icmp
Source IP List  : ALL_eNBs
Dest IP List    : NMSs
Action          : permit
```

```
Event Log : (Not Specified)
Hits      : 0 flows
-----
No. of entries : 4
=====
*A:7750-1>show>app-assure>group#
```



Note:

The above configuration is generic and may need to be modified to suit the deployment requirements. For example, in the case of small cell SeGW deployment, traffic on other ports needs to be allowed to/from different NMS type servers, such as allowing TCP port 7003 and port 7013 to HDM servers. This can be accomplished by configuring additional entries in the above session filter.



Note:

By allowing port 22 for FTP, the AA FW automatically opens and closes the associated data channel ports. For more information about AA FW capabilities, with regard to OAM FW protection, see Application Assurance Stateful Firewall.

Conclusion

The SR OS AA stateful firewall feature runs on AA-ISA and extends application-level analysis to provide an in-line stateful service, integrated within the Security Gateway (SeGW).

AA stateful packet filtering, combined with AA Layer 7 classification and control, provides advanced, next-generation firewall functionality, protecting mobile network core infrastructure, such as MMEs, SGWs, and NMSs.

Application Assurance — Stateful Firewall

This chapter describes Application Assurance stateful firewall (FW) configurations for protecting residential and WiFi subscribers.

Topics in this chapter include:

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

Applicability

Initially, this chapter was written for SR OS Release 11.0.R1. The TCP validation section was added for SR OS Release 14.0.R4.

Overview

The AA SR OS 11.0.R1 stateful FW feature extends AA-ISA application level analysis to provide an in-line integrated stateful service that protects subscribers from malicious attacks. AA stateful packet filtering combined with AA L7 classification and control, empowers operators with advanced, next generation firewall functionality that is integrated within the Service Router. The AA stateful firewall (FW) and application firewall runs on AA-ISA. Using stateful inspection, the AA firewall not only inspects packets at layers 3-7, but also monitors and keeps track of the connection's state. If the operator configures a **deny** action within a session filter, then the matching packets (matching both the AA Application QoS policy (AQP) and associated session filter match conditions) are dropped and no flow session state/context is created.

AA FW can be used in all deployments of AA-ISA; mobile (MG OS) and fixed (SR OS), however the configurations examples here, while still very applicable (and almost 100% identical in mobile deployments) are focused on AA-ISA deployments in fixed networks.

The AA-ISA FW enabled solution provides:

- Stateful (and stateless) packet filtering and inspection with application-level gateway (ALG) support
- DoS attack protection

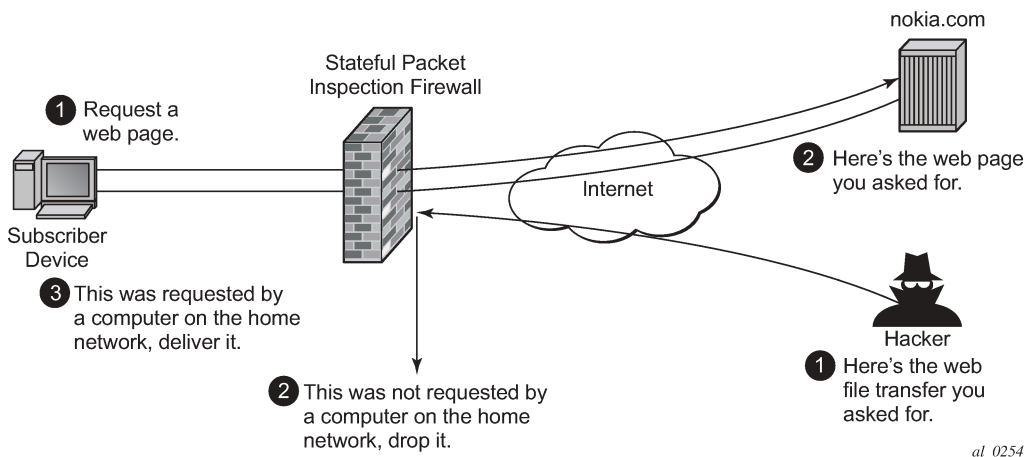
In SR OS Release 14.0, additional firewall functionalities were added, such as TCP-validation, threshold crossing alerts, syslog and statistics related to firewall actions.

The objective of this chapter is to describe the required configuration within AA-ISA (divert to AA-ISA basic knowledge is assumed) in order to enable AA FW and protect AA subscribers from attacks (Unsolicited attacks and DoS attacks), while still allowing pin-holing through the firewall, so that applications like peer to peer gaming and various ALGs (such as FTP) are not affected.

Stateful Filtering

By performing stateful inspection, AA-ISA takes into account which side initiated a session and acts accordingly. Stateful flow processing and inspection utilizes IP layers 3/ 4 header information to build a state of the flow within AA-ISA. Layer 7 inspection is used in order to provide ALG support. Stateful flow/session processing takes note of the originator of the session and hence can allow traffic to be initiated from the subscriber, while denying (when configured) traffic originating from the network. Packets received from the network are inspected against the session filter and only those that are part of a subscriber initiated session are allowed.

Figure 33: Block Unsolicited Traffic



To support the example shown in [Figure 33: Block Unsolicited Traffic](#), AA is configured with an action to block unsolicited traffic; traffic that is not originated/initiated from the subscriber. The direction field in match criteria of AQPs is utilized to enable this functionality.

Figure 34: SFW — Allow Gaming





Figure 34: SFW — Allow Gaming shows a similar concept. It is used to allow UDP traffic for peer to peer applications (such as gaming). Once the traffic from one peer is seen by AA-ISA, a pin-hole is opened in the reverse direction to allow for the corresponding UDP traffic from the peer.

Stateless packet filtering on the other hand does not take note of the session initiator. It discards or allows packets independently of any previous packets. In addition to AA-ISA's support for stateless (and stateful) filtering, stateless packet filtering can be performed in the system using line card ACLs (and/or MGISM PCC rules in mobile gateway deployments).

Application Layer Gateway Filtering

Figure 35: ALG Support Example — FTP



al_0256

AA FW inspection of packets at Layer 7 offers Application Layer Gateway functionality for a large list of applications (for example, FTP, SIP, RTSP, PPTP, IRC, etc.). These applications make use of control channels or flows that spawn other flows. AA FW inspects the payload of these control flows so it can open a pinhole in advance for unspawned data flows. [Figure 35: ALG Support Example — FTP](#) depicts an example of AA ALG support for FTP traffic.

Denial Of Service (DOS) Protection

DoS attacks work by consuming network and system resources, making them unavailable for legitimate network applications. Network flooding attacks, malformed packets and port scans are examples of such DoS attacks.

The aim of AA FW DOS protection is to protect subscribers and prevent any abuse of network resources.

Using AA FW stateful session filters, operators can protect their subscribers from any port scan scheme. This can be done by configuring the session filters to disallow any traffic that is initiated from the network.

Furthermore, AA ISA provides configurable flow policers. These policers, once configured, prevent a wide range of flooding attacks (such as ICMP PING flooding, UDP flooding, SYN Flood Attack...etc.). These policers provide protection at multiple levels; per system per application/application groups and per subscriber per applications/applications groups.

There are two types of AA ISA flow policers; flow setup rate policers and flow count policers. Flow setup rate policers limit the number of new flows, while flow count policers limit the total number of active flows.

In order to protect hosts and network resources, AA FW validates/checks different fields in the packet's header (checksum, TCP Flag, etc.) and if any fails it declares the packet to be invalid. This complements the 7x50 subscriber management enhanced security features, such as IP (or MAC) anti-spoofing protection

(such as protecting against LAND attacks) and network protocol DoS protections. The cut-through-drop AQP action must be configured in order to drop these types of invalid packets.

Virtual FW/Zone-Based FW

AA FW can provide up to 128 virtual FWs, each with its own FW policies. This is achieved through the use of AA-partitions.

In addition, AA subscribers within the same AA partition can have different application profiles with different Application Service Options (ASO) values. This provides a further control mechanism to enable/disable firewall rules.

For example, the operator may want to have some subscribers possess full firewall protection, while other subscribers not subscribed to this service to have a partial firewall protection that focuses on protecting network resources, rather than network and subscribers resources.

Configuration

AA-ISA AQPs were enhanced in R11.0.R1 with several AQP actions that provide session filtering functionality. As is the case of all AQPs, these have partition level scope, which allows different FW policies to be implemented by utilizing AA partitions concepts within the same AA-ISA group. Hence, multiple virtual AA FW instances can be realized, without the need for multiple physical instances of FWs to implement different FW policies.

The AA FW stateful session filter consists of multiple entries (similar to ACLs) with a match an action per entry. Actions are **deny** or **permit**. A **deny** action results in packets being discarded without creating a session/flow context. Match conditions include IP protocol types, source and destination IP addresses and ports. An overall default action is also configurable in case of no match to any session filter entry.

AQPs with session filter actions need to have — as a matching condition — traffic direction, ASOs, and/or a subscriber name. These AQP match rules cannot have any references to applications and/or application groups.


An AQP action to drop malformed/errored packets is also available.

Statistics are incremented when packets are dropped by a session filter. These are accounted against:

- protocol = denied by default policy
- application= unknown
- application group = unknown.

The configuration topology is shown in [Figure 36: Configuration Topology](#).

Figure 36: Configuration Topology



1. Application Profile configuration:

There is nothing new introduced in application profiles in order to support FW. This section explains how to configure the application profile to allow differentiated FW services for different subscribers. In a nut shell, the AA common building construct/attribute for differentiated policy is ASO.

To configure an ASO for FW protection:

```
configure application-assurance group 1:1 policy
begin
app-service-options
characteristic "FW-Protection" create
value "None"
value "ON"
default-value "None"
exit
characteristic "ISP-Protection" create
value "None"
value "ON"
default-value "None"
exit
characteristic "DOS-Protection" create
value "None"
value "ON"
default-value "None"
exit
```

```
exit
```

In the preceding example:

- ASO FW protection allows the operator to select if the subscriber is FW protected or not.
- ASO DOS protection refers to if the subscriber is protected from DOS attacks.
- ASO ISP protection is different from the preceding two as it protects the ISP resources by (in the example that follows) not allowing unsolicited traffic. This should be ON for all subscribers (it is then arguable if someone needs it to be defined in the ASO list, instead of merely configuring an AQP to protect ISP resources all the time).

These ASOs are referenced in appProfiles (and later in AQPs) as follows:

```
configure application-assurance group 1:1 policy
begin
  app-profile "Protected" create
  divert
  characteristic "FW-Protection" value "ON"
  characteristic "ISP-Protection" value "ON"
  characteristic "DOS-Protection" value "ON"
exit
```

The preceding application profile Protected is assigned to subscribers who opted/subscribed to the firewall protection service; for example sub 1 and sub 2 in the example shown in [Figure 36: Configuration Topology](#).

Subscribers who are not protected (for example sub 3 in [Figure 36: Configuration Topology](#)) are assigned a different profile:

```
configure application-assurance group 1:1 policy
begin
  app-profile "unProtected" create
  divert
  characteristic "FW-Protection" value "None"
  characteristic "ISP-Protection" value "ON"
  characteristic "DOS-Protection" value "None"
exit
```

An alternative method to using application profiles/ASOs to provide differentiated services is to configure multiple partitions with different AQPs/ session filters. One partition for example will be for subscribers who are provided with firewall protection, while another is used for subscribers who are not protected. This configuration is simpler and provides statistics per partition. This example however covers the more complex case using ASOs/appProfiles.

2. Flow count policer configuration:

```
configure application-assurance group 1 policer Dos_police_Flow_count type flow-count-limit
granularity subscriber create
  flow-count 500
exit
```

The preceding configuration limits the number of flows a subscriber can have at any time to 500. This is done to protect against DoS attacks. The value 500 is arbitrary and requires tuning for each deployment.

```
configure application-assurance group 1 policer Dos_Police_ICMPFlows type flow-count-limit
granularity system create
flow-count 5000
exit
```

This configuration limits the total number of flows that matches the configured AQP matching condition. It is used for ICMP applications to prevent mass port scanning.

3. TCP Protocol Validation configuration

```
configure application-assurance group 1:1 tcp-validate TCP_protect create
```

This simple configuration allows the operator to call TCP_protect policy from within an AQP action entry.

The operator can also configure the policy to be "strict", in which case the AA checks for valid sequence and acknowledgements numbers. In this example, the "strict" option is not used.

4. Application configuration

The following configuration is standard with AppDB. It is shown here for reference.

```
configure application-assurance group 1:1 policy begin
application ICMP create
exit
app-filter
entry 1540 create
protocol eq "non_tcp_udp"
ip-protocol-num eq icmp
application "ICMP"
no shutdown
exit
entry 35500 create
protocol eq "non_tcp_udp"
ip-protocol-num eq ipv6-icmp
application "ICMP"
no shutdown
exit
```

5. Session-Filter

The following displays session-filter configuration commands to be used in Step 6 later.

```
configure application-assurance group 1:1 session-filter <name> create
description <description>
default-action permit|deny # default=deny
entry n create
description <entry-description>
match
ip-protocol-num <ip-protocol-number>
no src-ip <ip4_or_v6-address/mask>
no dst-ip <ip4_or_v6-address/mask>
no src-port {eq|gt|lt} <port-num> #or
range <start-port-num> <end-port-num>
no dst-port {eq|gt|lt} <port-num> #or
range <start-port-num> <end-port-num>
exit
action permit|deny
```

```
exit
entry m create
---snip---
```

Parameters

- **entry** *n* — A session filter can have multiple match-action rules, each of these match-action rules represent an entry within the session-filter. The entries are executed in order. If a match is found, within one entry, the subsequent entries within the session-filter are skipped (not evaluated).
- **default-action** [**permit** | **deny**] — This action is performed if no match is found for any of the configured entries within the session-filter. Default is deny.
 - A **deny** action will drop the packet and will not allow a flow record to be allocated for that flow. A **drop** action within AA AQP will drop the packet but it will still create flow record.
 - A **permit** action will allow the packet to flow through the system. A flow record is also allocated. The packet may get dropped by other configured AQP actions (due to header check failures).
- **description** *description-string*
This configures a text string, up to 80 characters, which can be used to describe the use of the session-filter.
- **match** — Keywords to perform the action specified under the **action** keyword only if the conditions in the match section are met.
 - **ip-protocol** *ip-protocol-number*
ip-protocol-number — 1..255
 - Decimal, hexadecimal or binary representation
 - Supported IANA IP protocol names:
 - crtp, crudp, egp, eigrp, encap, ether-ip, gre, icmp, idrp, igmp, igp, ip, ipv6, ipv6-frag, ipv6-icmp, ipv6-no-nxt, ipv6-opts, ipv6-route, isis, iso-ip, l2tp, ospf-igp, pim, pnni, ptp, rdp, rsvp, sctp, stp, tcp, udp, vrrp
 - **src-ip>/dst-ip** *ipv4-address/mask src-ip/dst-ip ipv6-address/mask*
 - Source/destination IP address within the packet header.
 - IPv4 or IPv6 formats are allowed, with prefixes masks.
 - **src-port** *src-port-numbers*
src-port {**eq** | **gt** | **lt**} *port-num*
eq — equal, exact match
gt — match port numbers that are greater than the one specified.
lt — match port numbers that are smaller than the one specified.
port-num — 0..65535 (Applicable to TCP, UDP and SCTP protocols only.)
 - **src-port range** *start-port-num end-port-num*
range — Keyword- that match port numbers within the specified range:
start-port-num — 0..65535
end-port-num — 0..65535
 - **dst-port** *dst-port-number*
 - Same as source port number explained above, but applied against destination port number.

- **action deny | permit**
 - **deny** or **permit** action is only executed if a match is found.
 - **deny** action will drop the packet and will not create a flow record.
 - **permit** action will allow the packet to go through (unless another different action is found that causes it to be dropped).
- **no entry *entry-id***
 - Causes the entry to be deleted.
- **no session-filter *session-filter-name***
 - Causes the session filter to be deleted.

```
config application-assurance group 1:1
  session-filter " denyUnsolicitedwMgmtCntrl" create
  description "S-FW opted-in sub – allow ISP access"
  default-action deny
  entry 10 create
    description "allow ICMP access from ISP LAN1"
    match
      ip-protocol-num icmp
      src-ip 10.10.8.0/24
    exit
    action permit
  exit
  entry 20 create
    description "allow ICMP access from ISP LAN2"
    match
      ip-protocol-num icmp
      src-ip 192.168.0.0/24
    exit
    action permit
  exit
  entry 30 create
    description "allow all TCP (e.g. FTP/telnet)access from ISP LAN2"
    match
      ip-protocol-num tcp
      src-ip 192.168.0.0/24
    exit
    action permit
  entry 40 create
    description "allow TCP on port 80 /HTTP access from ISP LAN1"
    match
      ip-protocol-num tcp
      src-ip 10.10.8.0/24
      dst-port eq 80
    exit
    action permit
  exit
```

This session filter is used to protect systems located in LAN2. It drops all unsolicited traffic except for FTP coming from LAN1.

```
configure application-assurance group 1:1
  session-filter "protectISPlan2" create
  description "S-FW to deny all unsolicited requests to LAN2"
  default-action deny
  entry 10 create
```

```
        description "allow ftp access from ISP LAN1"
        match
            ip-protocol-num tcp
            src-ip 10.10.8.0/24
            dst-port eq 21
        exit
        action permit
    exit
exit
```

6. AQP configuration:

```
configure application-assurance group 1:1 policy
begin
app-qos-policy

entry 100 create
description "Protecting ISP1 from DoS attacks from subs"
match
    traffic-direction subscriber-to-network
    characteristic "ISP-Protection" eq "ON"
    dst-ip eq 10.10.8.0/24
exit
action
    flow-count-limit Dos_police_Flow_count
exit
no shutdown
exit

entry 105 create
description "Protecting ISP2 from DoS attacks from subs"
match
    traffic-direction subscriber-to-network
    characteristic "ISP-Protection" eq "ON"
    dst-ip eq 192.168.0.0/24
exit
action
    flow-count-limit Dos_police_Flow_count
exit
no shutdown
exit
```

These AQPs protect the ISP network by limiting the number of concurrent flows. Dropping malformed packets is done by entry 130 (later).

To guard against ICMP flooding attacks, a flow count policer (defined earlier) is used as follows:

```
configure application-assurance group 1:1 policy
begin
app-qos-policy entry 107 create
match
    application eq "ICMP"
    traffic-direction subscriber-to-network
exit
action
    flow-count-limit Dos_Police_ICMPFlows
exit
no shutdown
exit
```

To guard against attacks exploiting TCP handshake mechanisms, TCP validate policy (defined earlier) is used as follows:

```
configure application-assurance group 1:1 policy
begin
  app-qos-policy
  entry 108 create
  match
    characteristic "ISP-Protection" eq "ON"
  exit
  action
    tcp-validate "TCP_protect"
  exit
  no shutdown
exit
entry 109 create
match
  characteristic "FW-Protection" eq "ON"
exit
action
  tcp-validate "TCP_protect"
exit
no shutdown
exit
```

TCP validation works on both direction and needs to be called in from within a sub-default AQP entry. Therefore, this AQP action cannot be restricted to one ISP versus another because no destination IP can be specified. The configuration shown runs TCP validation policy when ISP-Protection or FW-protection ASOs are enabled.

The preceding configuration will ensure, for example, that no TCP session starts without the proper handshake message exchanges.

In order to protect ISP LAN2 from all incoming traffic (unsolicited), the operator configures entry 120.

```
entry 120 create
match
  traffic-direction subscriber-to-network
  characteristic "ISP-Protection" eq "ON"
exit
action
  session-filter "protectISPLan2"
exit
no shutdown
exit
```

ProtectISPLan2 session filter drops all unsolicited traffic to LAN2 (highly secure) except for access to FTP services coming from ISP LAN1. Details of these configurations are shown in Session-Filter (step 5).

To enable stateful protection for opted-in subscribers:

```
configure application-assurance group 1:1 policy
begin
  app-qos-policy
  entry 110 create
  description "FW for managed opted-in subs"
  match
    traffic-direction network-to-subscriber
    characteristic "FW-Protection" eq "ON"
```

```

exit
action
    session-filter "denyUnsolicitedwMgmtCntrl"
exit
no shutdown
exit

```

The preceding AQP protects opt-in subscribers from unsolicited traffic but still allows unsolicited traffic from ISP subnets to manage the subscriber's network.

Dropping malformed/illegal packets and protecting against DOS attacks is done via the following entry 130 and 131.

```

entry 130 create
    match
        traffic-direction subscriber-to-network

        characteristic "DOS-Protection" eq "ON"
    exit
    action
        flow-count-limit Dos_police_Flow_count
    exit
    no shutdown
exit
entry 131 create
    match
        characteristic "DOS-Protection" eq "ON"
    exit
    action
        error-drop
        overload-drop
        fragment-drop all
    exit
    no shutdown
exit

```

7. Configuration of Threshold Crossing Alerts (TCA).

Operators can configure AA to generate TCAs for various firewall related parameters, such as error-drop, session-filter hits, TCP-validate, fragment-drop-all etc. as well as flow count policers. An example of a TCA used for TCP_validation policy is as follows:

```

configure application-assurance group 1:1 statistics threshold-crossing-alert
    tcp-validate "TCP_protect" direction from-sub create
    high-wmark 50 low-wmark 40
exit

```

Unlike the other TCAs, in order to configure TCAs for flow count policers, operators need first to configure AA admit-deny to allocate ISA resources to record, such as:

```

configure application-assurance group 1:1 statistics aa-admit-deny policer-stats-resources

```

Then, a TCA can be configured for any flow based policer in the system, such as:

```

configure application-assurance group 1:1 statistics threshold-crossing-alert
    policer "Dos_police_Flow_count" direction from-sub create
    high-wmark 300 low-wmark 199
exit

```

The system allows the various AA-admit-deny statistics to be exported via XML according to the configured accounting policy on the system. SAM-A can then use these statistics to generate the right reports / alerts.

As a prerequisite, an accounting policy is configured for aa-admit-deny statistics:

```
configure log accounting-policy 5 record aa-admit-deny
```

Then, the operator can configure AA to export the statistics related to various firewall functions configured in the system, such as:

```
configure application-assurance group 1:1 statistics aa-admit-deny
    accounting-policy 5
    collect-stats
    session-filter-stats
    policer-stats-resources
    tcp-validate-stats
exit
```

GTP and STCP admit deny stats are related to firewall deployment within a SeGW, which is not covered within the scope of this chapter.

Show Routine — AQP:

```
*A:PE-1# show application-assurance group 1:1 policy app-qos-policy 110
```

```
=====
Application QoS Policy Entry 110 (Default Subscriber Policy)
=====
```

```
Description : FW for managed opted-in subs
Admin State : in-service
Hits:       : 0 flows
Conflicts   : 0 flows
```

```
Match :
  Traffic Direction      : network-to-subscriber
  ASO Characteristics    :
  FW-Protection          : eq ON
```

```
Action :
  Session Filter         : denyUnsolicitedwMgntCtrl
=====
```

Show Routines — Session Filter:

```
*A:PE-1# show application-assurance group 1:1 session-filter
"denyUnsolicitedwMgntCtrl"
```

```
=====
AA Session Filter Instance "denyUnsolicitedwMgntCtrl"
=====
```

```
Description      : (Not Specified)
Default Action   : deny
  Event Log      : (Not Specified)
AQP Entries:     :
  110
```

```
-----
Filter Match Criteria
-----
```

```

Entry      : 10
Description : allow ICMP access from ISP LAN1
IP Protocol : icmp
Source IP   : 10.10.8.0/24
Action      : permit
  Event Log : (Not Specified)
Hits:       : 0 flows
-----
Entry      : 20
Description : allow ICMP access from ISP LAN2
IP Protocol : icmp
Source IP   : 192.168.0.0/24
Action      : permit
  Event Log : (Not Specified)
Hits:       : 0 flows
-----
Entry      : 30
Description : allow all TCP (e.g. FTP/telnet)access from ISP LAN2
IP Protocol : tcp
Source IP   : 192.168.0.113/320/24
Action      : permit
  Event Log : (Not Specified)
Hits:       : 0 flows
-----
Entry      : 40
Description : allow TCP on port 80 /HTTP access from ISP LAN1
IP Protocol : tcp
Source IP   : 10.10.8.0/24
SourceDest Port : eq 80
Action      : permit
  Event Log : (Not Specified)
Hits:       : 0 flows
-----
No. of entries : 4
=====

```

Show Routines — TCP Validation:

```

*A:PE-1# show application-ass group 1:1 tcp-validate "TCP_protect"
=====
Application Assurance Group 1:1 tcp-validate "TCP_protect"
=====
Description      : (Not Specified)
Event log        : (Not Specified)
Strict Validation: No
AQP referenced   : Yes
-----
Decision Statistics          sub-to-net          net-to-sub
-----
Total
-----
Allowed
  Octets              0              0
  Packets             0              0
Dropped
  Octets              0              0
  Packets             0              0

```

| Dropped Reason | | | |
|------------------------------|---|--|---|
| ----- | | | |
| Bad Flags | | | |
| Octets | 0 | | 0 |
| Packets | 0 | | 0 |
| Bad Options | | | |
| Octets | 0 | | 0 |
| Packets | 0 | | 0 |
| Bad Sequence Number | | | |
| Octets | 0 | | 0 |
| Packets | 0 | | 0 |
| Bad Acknowledgment Number | | | |
| Octets | 0 | | 0 |
| Packets | 0 | | 0 |
| No Establishment | | | |
| Octets | 0 | | 0 |
| Packets | 0 | | 0 |
| SYN After Conn Establishment | | | |
| Octets | 0 | | 0 |
| Packets | 0 | | 0 |
| Asymmetric Traffic | | | |
| Octets | 0 | | 0 |
| Packets | 0 | | 0 |
| Traffic After Reset (RST) | | | |
| Octets | 0 | | 0 |
| Packets | 0 | | 0 |
| Fragmented | | | |
| Octets | 0 | | 0 |
| Packets | 0 | | 0 |

```

*A:PE-1# show application-assurance threshold-crossing-alert detail
=====
Application Assurance Threshold Crossing Alerts
=====
-----
policer "Dos_police_flow_count" from-sub
-----
Group:Part      : 1:1                Trigger on      : denied-traffic
High watermark  : 300                Low watermark   : 199
Last raised     : N/A                Last cleared    : N/A
State           : cleared
-----
tcp-validate "TCP_protect" from-sub
-----
Group:Part      : 1:1                Trigger on      : denied-traffic
High watermark  : 50                 Low watermark   : 40
Last raised     : N/A                Last cleared    : N/A
State           : cleared
No. of TCAs    : 2
=====
*A:PE-1#
*A:PE-1>tools>dump>app-assure>group# admit-deny-stats
=====
Application-Assurance Group 1:1 Admit-Deny Statistics
=====
-----
Admitted Sub-To-Net   Denied Sub-To-Net   Admitted Net-To-Sub   Denied Net-To-Sub
Packet Validation Statistics
(Packets)             (Packets)           (Packets)             (Packets)
-----
Error

```

```

0
Fragments: Out-Of-Order 0 0 0 0
Fragments: All 0 0 0 0
Overload 0 0 0 0
N/A 0 N/A 0
-----
Admitted Sub-To-Net Denied Sub-To-Net Admitted Net-To-Sub Denied Net-To-Sub
Session Filter Statistics
(Sessions) (Packets) (Sessions) (Packets)
-----
Session Filter: test
Entry: 1 0
Default Action 0 0
-----
Admitted Sub-To-Net Denied Sub-To-Net Admitted Net-To-Sub Denied Net-To-Sub
TCP Validation Statistics
(Packets) (Packets) (Packets) (Packets)
-----
test 0
TCP_protect 0
TCP_protect_ISP1 0
-----
*A:PE-1>tools>dump>app-assure>group#

```

Conclusion

The AA stateful packet filtering feature combined with AA Layer 7 classification and control empowers operators with an advanced, next generation firewall functionality that is integrated within SR OS. This chapter focused on traditional stateful and stateless session firewall functionality.

Application Assurance — Usage Monitoring and Policy Control via Diameter Gx Protocol

This chapter provides information about the diameter (Gx) control feature.

Topics in this chapter include:

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

Applicability

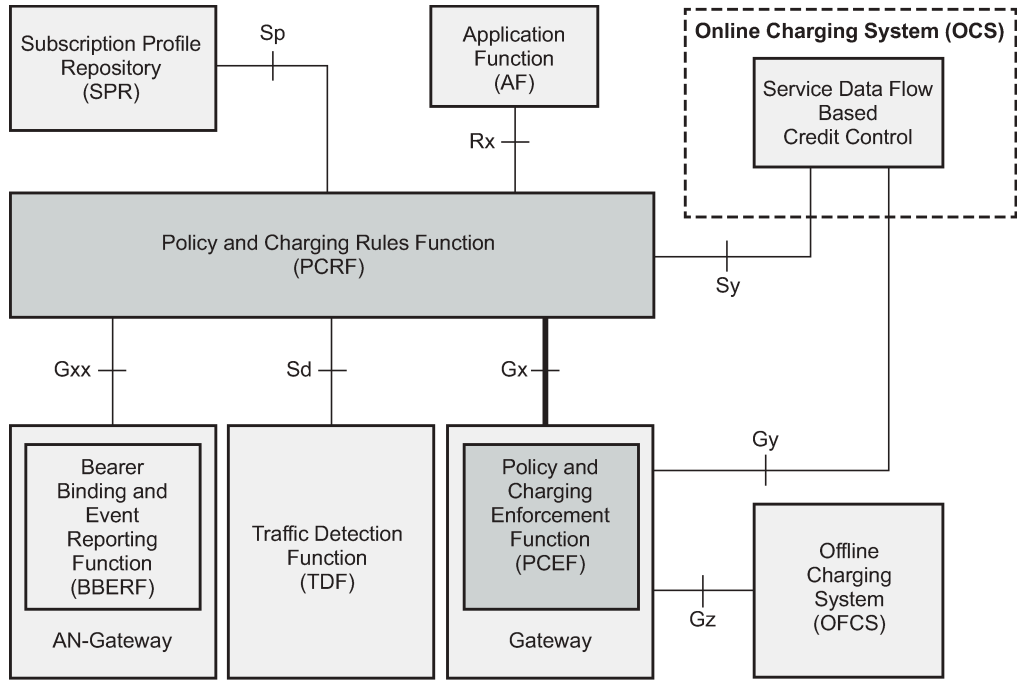
This chapter was initially based on SR OS Release 13.0.R1, but the Diameter Base configuration in the current edition is based on SR OS Release 19.10.R1.

Overview

The Gx reference point is defined in the Policy and Charging Control (PCC) architecture within the 3rd Generation Partnership Project (3GPP) standardization body. The Gx reference point is used for policy and charging control. The PCC architecture is defined in the 23.203 3GPP technical specification, while the Gx functionality is defined in the 29.212 3GPP technical specification. The SR OS implementation of Gx supports both Release 11 and Release 12 of the specification. Gx is an application of the Diameter Base Protocol (RFC 6733).

As shown in [Figure 37: Gx reference point](#), Gx is placed between a policy server Policy and Charging Rule Function (PCRF) and a traffic forwarding node Policy and Charging Enforcement Function (PCEF) that enforces rules set by the policy server.


Figure 37: Gx reference point



al_0651

Although the Gx reference point is defined within the 3GPP standardization body, its applicability has also spread to wire-line operations to achieve mobile–fixed convergence gains by streamlining policy management functions into a single Gx based infrastructure, see [Figure 38: Convergence](#).

Figure 38: Convergence

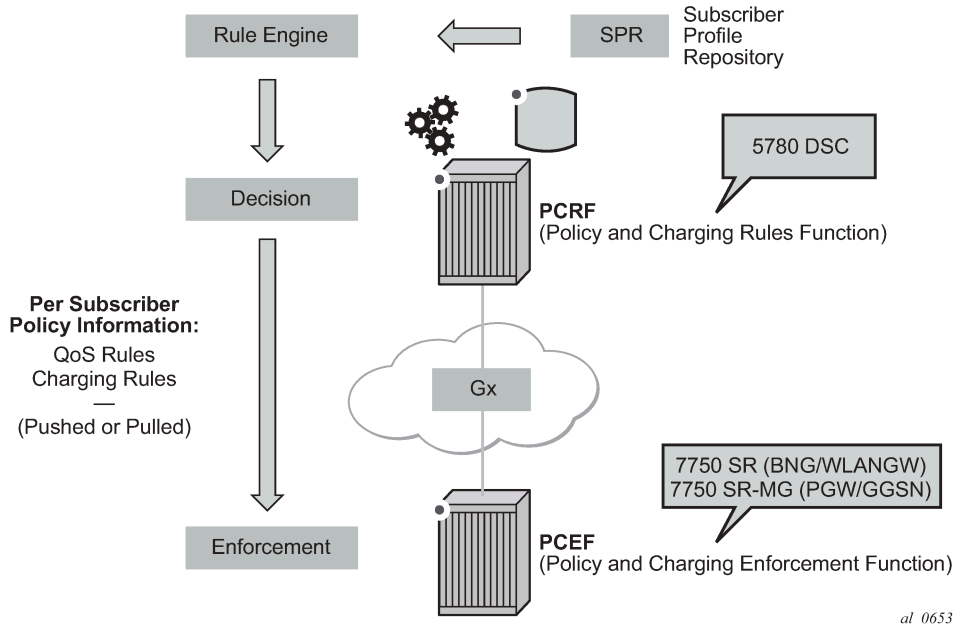


Gx support on SR OS is applicable to Enhanced Subscriber Management (ESM) functions, including the Application Assurance (AA) functions. The focus of this chapter is on the AA aspects of Gx.

The SR OS based Gx interface offers the following functionalities:

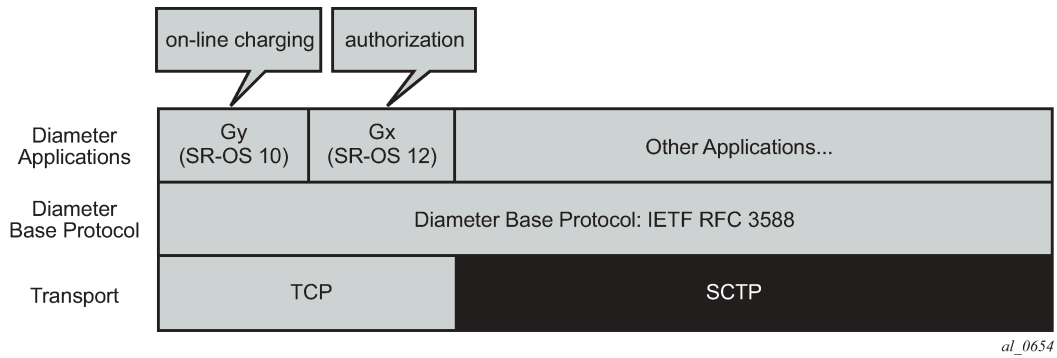
- ESM subscriber-based policy decision providing
 - QoS attributes
 - charging attributes
 - subscriber identification
- Usage management
 - usage reporting from PCEF to PCRF

Figure 39: Gx reference point



Note that Gx does not provide subscriber authentication or subscriber IP address assignment.

Figure 40: Diameter protocol stack



Policy assignment use case

The SR OS accepts the following policy information from PCRF using Gx:

- Subscriber profile strings and SLA profile strings.
- Subscriber-QoS-overrides.
- Application profile strings.
- Application subscriber options (ASOs) related to AA.

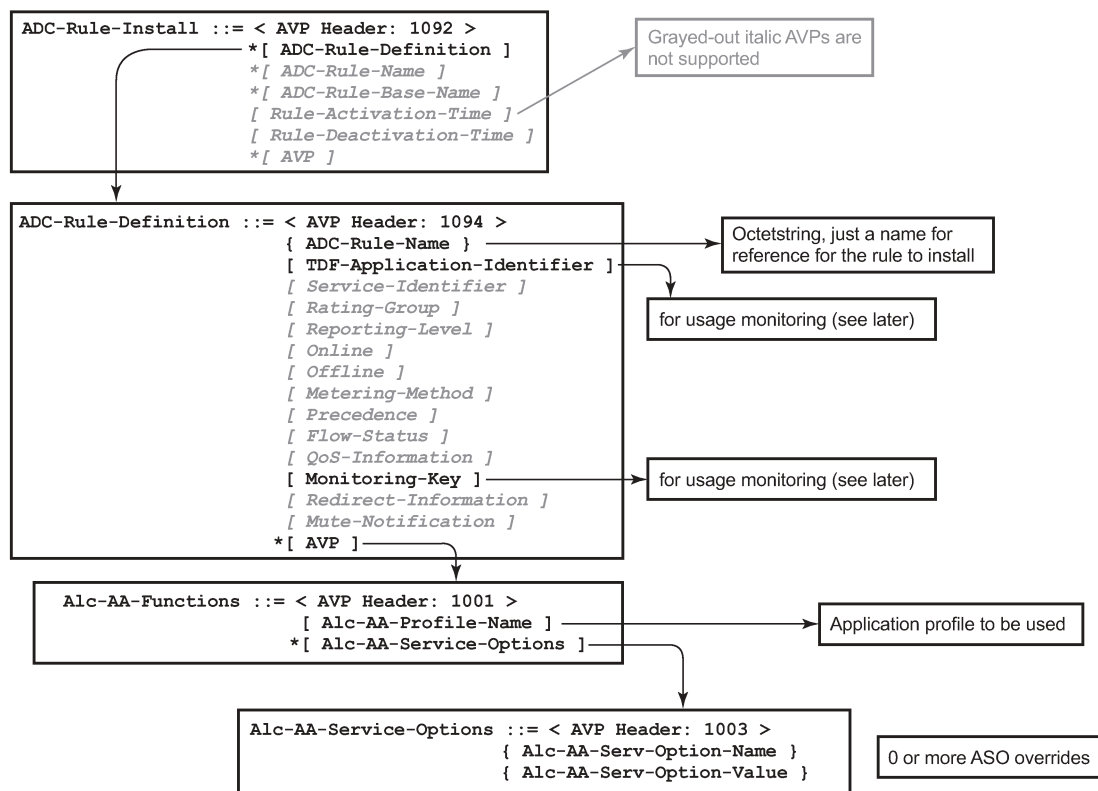
Gx operates at subscriber host level and creates an "IP-CAN Session" (IP Connectivity Access Network) for every subscriber host. However, as AA operates at the subscriber level, AA related policies apply to all the hosts belonging to that subscriber.

This chapter covers AA related functionalities, namely: application profile and ASO assignments and override. These functionalities are defined in either:

1. Application Detection and Control (ADC) rules—per 3GPP Release 11— **or**
2. Policy and Charging Control (PCC) rules—per 3GPP Release 12—
 - **Application Profile** Alc-AA-Profile-Name Attribute-Value-Pair (AVP)
 - RADIUS equivalent is Alc-App-Prof-Str Vendor-Specific-Attribute (VSA)
 - **ASO overrides** Alc-AA-Service-Options AVP
 - RADIUS equivalent is Alc-AA-App-Service-Options VSA

Details of the ADC rules and related Nokia defined AVPs defined for use by AA are shown in [Figure 41: ADC rules and related Nokia-defined AVPs defined for use by AA.](#)

Figure 41: ADC rules and related Nokia-defined AVPs defined for use by AA

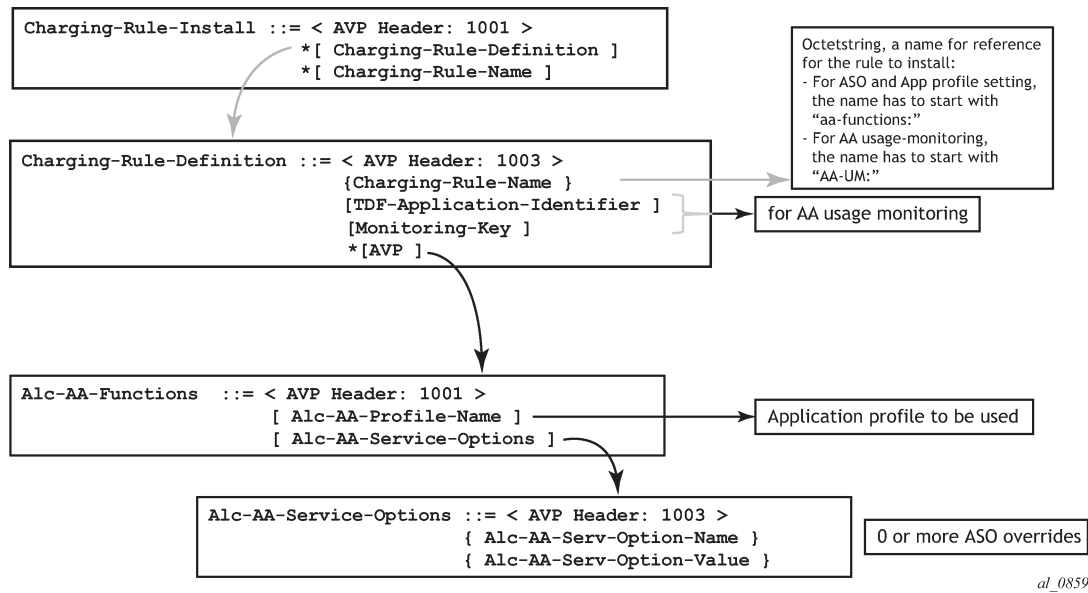


al_0655

The ADC-Rule-Install is at the root level of the GX message.

As for 3GPP Release 12, the details of the PCC rules and related Nokia-defined AVPs defined for use by AA are shown in [Figure 42: PCC rules and related Nokia-defined AVPs defined for use by AA.](#)


Figure 42: PCC rules and related Nokia-defined AVPs defined for use by AA



The PCC-Rule-Install, as in the case of ADC-Rule-Install, is at the root level of the GX message.


An example of the AVPs to install the application profile "gold_level" using 3GPP Release 11 (/ADC rules) is shown in Figure 42: PCC rules and related Nokia-defined AVPs defined for use by AA.

Figure 43: ADC rule example of AVPs to install the application profile "gold_level"



An example of the AVPs to install the "gold_level" application profile using 3GPP Release 12 (/PCC rules) is shown in Figure 45: Capture of the ADC rule assignment of the "gold_level" appProfile.

Figure 44: PCC rule example of AVPs to install the application profile "gold_level"



Note:

ADC-Rule-Names and PCC-Rule-Names have to start with *aa-functions* when they contain an Alc-AA-Functions AVP.

The assignment of the *gold_level* appProfile is shown in another format in [Figure 45: Capture of the ADC rule assignment of the "gold_level" appProfile.](#)

Figure 45: Capture of the ADC rule assignment of the "gold_level" appProfile

```

adc-rule-install (1092) V----- [184]
  vendor-id TGPP
  data [172] (Grouped)
    adc-rule-definition (1094) V----- [172]
      vendor-id TGPP
      data [160] (Grouped)
        adc-rule-name (1096) V----- [32]
          vendor-id TGPP
          data [20] (UTF8String) : aa-functions:appprof
        AA-Functions (1001) V----- [128]
          vendor-id ALU
          data [116] (Grouped)
            AA-Profile-Name (1002) V----- [17]
              vendor-id ALU
              data [5] (UTF8String) : gold level
            AA-App-Service-Options (1003) V----- [48]
              vendor-id ALU
              data [36] (Grouped)
                AA-App-Service-Options-Name (1004) V----- [17]
                  vendor-id ALU
                  data [5] (UTF8String) : level
                AA-App-Serv-Options-Value (1005) V----- [16]
                  vendor-id ALU
                  data [4] (UTF8String) : high
            AA-App-Service-Options (1003) V----- [48]
              vendor-id ALU
              data [36] (Grouped)
                AA-App-Serv-Options-Name (1004) V----- [18]
                  vendor-id ALU
                  data [6] (UTF8String) : p2p
                AA-App-Serv-Options-Value (1005) V----- [14]
                  vendor-id ALU
                  data [2] (UTF8String) : unlimited
          
```

al_0657

Application profiles and ASO overrides can be changed on-the-fly with a Re-Authentication-Request (RAR) message according to these rules:

- If an Application profile is present in the Gx message it is applied first. Then ASO AVPs are applied when present (in the Gx message). In other words:

- If a RAR message only contains the same application profile and no ASO overrides, then all previous ASO overrides are removed.
- When a RAR message contains the same application profile and new ASO overrides, then the new ASO overrides are applied, and the previous ASO overrides are removed.
- When a RAR message contains a new application profile, all previous ASO overrides are removed and replaced with the ASOs in the RAR if present.
- When a RAR message does not contain an application profile but only ASO overrides, then the new ASO overrides are added to the existing ASO overrides.

Note that a single Gx ADC (or PCC) rule cannot contain both AA subscriber policies (appProfile/ASO) and AA Usage monitoring (as outlined later). These have to be in separate ADC (or PCC) rules.

Usage management/monitoring use-case

The AA-ISA can monitor application usage at the subscriber level and report back to the PCRF whenever the usage exceeds the threshold(s) set by the PCRF when receiving requests from the PCRF over the Gx interface.

Usage monitoring can be used by operators to report to PCRF when:

- The AA-ISA detects the start of a subscriber application by setting the usage threshold to a very low value.
- A pre-set usage volume per subscriber application is exceeded.

AA can monitor subscriber's traffic for any defined:

- Application
- Application group, and/or
- Charging group

The AA-ISA reports the accumulated usage when:

- A usage threshold is reached.
- The PCRF explicitly disables the usage monitoring.
- The PCRF requests a report.
- The ADC (or PCC) rule associated with the monitoring instance is removed or deactivated.
- A session is terminated.

An AA defined application, application group and/or charging group is automatically allowed to be referenced by an ADC (or PCC) rule for the purpose of usage monitoring only if:

{It is already selected for either XML or RADIUS per subscriber accounting

OR

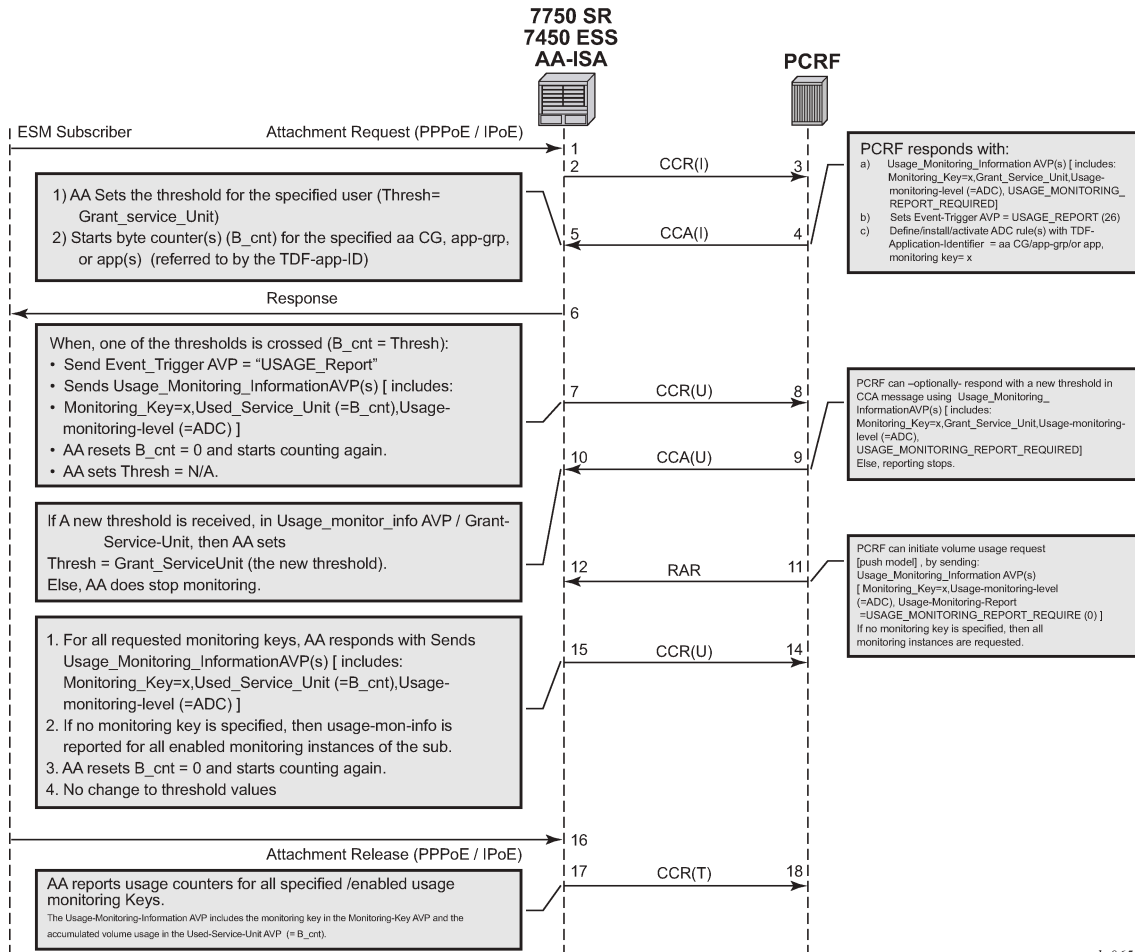
It is explicitly enabled by the operator for per subscriber statistics collection}

AND

Usage monitoring is enabled for the given AA group:partition

Figure 46: Call flow diagram illustrates the different messaging/call flows involved in application level usage monitoring. Details of the different supported AVPs used in these messages are illustrated later.

Figure 46: Call flow diagram



al_0658

The AA-ISA/PCEF supports Usage-Thresholds AVPs that refer to the thresholds (in bytes) at which point an event needs to be sent back to the PCRF, (see [Figure 44: PCC rule example of AVPs to install the application profile "gold_level"](#)).

Time based thresholds are not supported.

AA supports the "grant-service-unit" AVP using the following possible values (AVP):

- CC-Input-Octets AVP (code 412): from subscriber total byte count threshold.
- CC-Output-Octet AVP (code 414): to subscriber total byte count threshold.
- CC-Total-octets AVP (code 421): threshold of aggregate traffic (input and output byte counters).

As shown in [Figure 46: Call flow diagram](#), (T=7), AA sends a Credit Control Request (CCR_message) with a "USAGE_REPORT" Event-Trigger AVP to the PCRF when the usage counter reaches the configured usage monitoring threshold for a given subscriber (and given application group). AA counters are reset (to zero) when the monitoring threshold is reached (and an event is sent back to the PCRF). The counter(s) however does not stop counting newly arriving traffic. AA counters only include "admitted" packets. Any packets that were discarded by AA due to, for example, policing actions are not counted for usage-monitoring purposes.

The TDF-Application-Identifier AVP (within the ADC or PCC rule) refers to an AA Charging group, an AA application group or to an AA application. TDF-Application-Identifiers (for example, charging-groups) have to be manually entered at the PCRF to match the AA charging groups defined in the AA. If the TDF-Application-Identifier refers to a name that is used for both a charging group and an application (or an application group), AA monitors the charging group. In other words, the AA charging group has a higher precedence than the AA application group.

Gx usage monitoring AVP summary

For 3GPP Release 11 (using ADC rules), the following AVPs are used for AA-Usage monitoring:

```

ADC-Rule-Install ::= < AVP Header: 1092 >
    *[ ADC-Rule-Definition ]
    *[ ADC-Rule-Name ]

ADC-Rule-Definition ::= < AVP Header: 1094 >
    { ADC-Rule-Name }
    [ TDF-Application-Identifier ]; AA app/app-grp/chrp-grp
    [ Monitoring-Key ];

Usage-Monitoring-Information ::= < AVP Header: 1067 >
    [ Monitoring-Key ]
    0,2[ Granted-Service-Unit ]
        Granted-Service-Unit ::= < AVP Header: 431 >
            [ CC-Total-Octets ]
            [ CC-Input-Octets ]
            [ CC-Output-Octets ]

    0,2[ Used-Service-Unit ]
Used-Service-Unit ::= < AVP Header: 446 >
    [ CC-Total-Octets ] ;
    [ CC-Input-Octets ]
    [ CC-Output-Octets ]

    [ Usage-Monitoring-Level ]
; ADC_RULE_LEVEL (2)

    [ Usage-Monitoring-Report ]
; immediate report -- USAGE_MONITORING_REPORT_REQUIRED (0)

    [ Usage-Monitoring-Support ]
; to disable : USAGE_MONITORING_DISABLED (0)

```

For 3GPP Release 12 (using PCC rules), the following AVPs are used for AA-Usage monitoring:

```

Charging-Rule-Install ::= < AVP Header: 1001 >
    *[ Charging-Rule-Definition ]
    *[ Charging-Rule-Name ]

Charging-Rule-Definition ::= < AVP Header: 1003 >
    { Charging-Rule-Name } ;/ starts with "UM-AA:"
    [ TDF-Application-Identifier ]; AA app/app-grp/chrp-grp
    [ Monitoring-Key ];

Usage-Monitoring-Information ::= < AVP Header: 1067 >
    [ Monitoring-Key ]

```

```

0,2[ Granted-Service-Unit ]
  Granted-Service-Unit ::= < AVP Header: 431 >
    [ CC-Total-Octets ]
    [ CC-Input-Octets ]
    [ CC-Output-Octets ]

0,2[ Used-Service-Unit ]
Used-Service-Unit ::= < AVP Header: 446 >
  [ CC-Total-Octets ] ;
  [ CC-Input-Octets ]
  [ CC-Output-Octets ]

  [ Usage-Monitoring-Level ]
; PCC_RULE_LEVEL (1)

  [ Usage-Monitoring-Report ]
; immediate report -- USAGE_MONITORING_REPORT_REQUIRED (0)

  [ Usage-Monitoring-Support ]
; to disable : USAGE_MONITORING_DISABLED (0)

```

Configuration


This configuration example highlights the commands illustrating how Gx can be used to:

- Override AppProfile and ASO characteristics.
- Set and retrieve AA level usage monitoring metrics.

While the configuration associated with setting up the Gx interface toward the PCRF is shown for the sake of completeness, that aspect of the configuration is not explored in detail, and only a 3GPP Release 11 model is shown. Similarly, the Gx policies and usage monitoring associated with ESM host policies (non-AA aspects) are out of the scope of this chapter.

The configuration on the 7750 node is the same, independent of whether the PCRF supports 3GPP Release 11 (ADC) or Release 12(PCC) to provide AA policy control function.

Figure 47: Example configuration setup



al_0659

The BNG is set up with at least one IOM and one MS-ISA MDA configured as ISA-AA.

```
configure
  card 1
    card-type iom3-xp
    mda 1
      mda-type m20-1gb-xp-sfp
      no shutdown
    exit
    mda 2
      mda-type isa-aa
      no shutdown
    exit
  no shutdown
exit
card 3
  card-type iom3-xp
  mda 1
    mda-type isa-aa
    no shutdown
  exit
  mda 2
    mda-type isa-aa
    no shutdown
  exit
no shutdown
exit
```

The configurations in this example are broken down into four main steps:

1. Configuring the Gx interface (high-level)
2. Configuring AA application profiles and ASOs (high-level)
3. Configuring AA applications filters (high-level)
4. Configuring AA usage-monitoring

The focus of this configuration example is on Step 4, and the updated show routines related to AA ESM subscriber state are shown at the end of Step 2.

1. Configuring the Gx interface (high-level).

These commands bring up the Gx diameter control channel between the Gx Controller(/Server), also known as PCRF, and the PCEF(/BNG).

```
configure
  aaa
    diameter
      node "bng-gx.realm-1.com" create
      description "Authentication and Policy Management"
      source-address 192.0.2.2
      peer index 1 "dra-1.realm-1.com" create
      address 10.1.0.10
      no shutdown
    exit
  exit
exit
exit
exit
```

The diameter node "*bng-gx.realm-1.com*" is then referenced under subscriber management.

```
configure
 subscriber-mgmt
   diameter-application-policy "diamAppPlcy" create
     application gx
     diameter-node "bng-gx.realm-1.com" destination-realm "realm-1.com"
   exit
 exit
```

Then the created subscriber management policy "*diamAppPlcy*" is applied to the subscriber interface.

```
configure
 service
   customer 1 create
     description "Default customer"
   exit
   ies 1 customer 1 vpn 1 create
     description "Default Ies description for service id 1"
     subscriber-interface "ies-1-172.16.0.0" create
       address 172.16.0.0/12
       group-interface "grp-1-35782656-1" create
         dhcp
           server 172.16.200.200
           trusted
           lease-populate 2000
           gi-address 172.16.0.0
           no shutdown
         exit
         diameter-application-policy "diamAppPlcy"
         sap 1/1/4:1 create
           description "sap-grp-1"
           sub-sla-mgmt
             def-sub-profile "sub_prof"
             def-sla-profile "sla_prof"
             def-app-profile "app_prof_1"
             sub-ident-policy "sub_ident_A_1"
             multi-sub-sap 2
             no shutdown
           exit
         exit
       exit
     exit
   service-name "ACG Ies 1"
   no shutdown
 exit
```

Now verify the configuration and connectivity towards the PCRF by running the following command:

```
# show aaa diameter-node "bng-gx.realm-1.com" peers

=====
Peers
=====
Host identity                               Status           Default Preference Active
-----
dra-1.realm-1.com                           I-Open           No           50           Yes
-----
No. of peers: 1
```

The Peer-State-Machine State (PSM), as per RFC 6733, has the value I-OPEN indicating that the peer is operational. The "I-" stands for Initiator state, in this case the BNG is the initiator.

A detailed look into the traffic statistics between the PCEF and the PCRF (Gx controller) can be viewed using a show statistics command (see below). These statistics provide a breakdown of the messages exchanged:

```
# show aaa diameter-node "bng-gx.realm-1.com" peer "dra-1.realm-1.com" statistics
=====
Peer "dra-1.realm-1.com"
=====
Message                               Sent           Received
-----
Capabilities-Exchange-Request         7              0
Capabilities-Exchange-Answer          0              7
Disconnect-Peer-Request                0              0
Disconnect-Peer-Answer                 0              0
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
=====

# show subscriber-mgmt diameter-application-policy "diamAppPlcy" statistics
=====
Diameter node statistics for policy "diamAppPlcy"
=====
Message                               Requests       Answers
-----
Initial Credit-Control                 2              2
Update Credit-Control                  14             14
Termination Credit-Control              1              1
Re-Auth                                 2              2
Abort-Session                           0              0
-----

Request message transmission failure* 0
Request message retransmissions        0

Result code                            Sent           Received
-----
(1xxx) Informational                    0              0
(2xxx) Success                           0             14
(3xxx) Protocol Errors                   0              0
(4xxx) Transient Failures                 0              0
(5xxx) Permanent Failures                 0              0
=====
* indicates that the corresponding row element may have been truncated.
```

2. Configuring AA application profiles and ASOs (high-level)

To illustrate the use of application profiles and ASO overrides using Gx RAR messages, four ASOs and 2 appProfiles are defined, as follows.

"app_prof_1" is the default app-profile used when a subscriber is created on AA.

```
configure
  application-assurance
    group 129:34883 create
      policy
        begin
          app-service-options
            characteristic "permitDNS" persist-id 1 create
            value "no"
            value "yes"
            default-value "yes"
          exit
          characteristic "permitRDP" persist-id 2 create
            value "no"
            value "yes"
            default-value "yes"
          exit
          characteristic "permitHTTP" persist-id 3 create
            value "no"
            value "yes"
            default-value "yes"
          exit
        exit
      app-profile "app_prof_1" create
        description "Application Profile Id app_prof_1"
        divert
      exit
      app-profile "app_prof_2" create
        description "Application Profile Id app_prof_2"
        divert
      exit
    exit
```

3. Configuring AA applications filters (high-level)

First create the application group, as follows.

```
configure
  isa
    application-assurance-group 129 create
      primary 3/2
      backup 1/2
      partitions
      divert-fc be
      no shutdown
    exit
```

Then create the partition and associated charging groups, application groups, applications, etc.

```
configure
  application-assurance
    group 129:34883 create
      policy
        begin
          charging-group "0_rated" create
            export-id 1
          exit
          charging-group "default_charge_group" create
            export-id 255
          exit
        default-charging-group "default_charge_group"
        app-group "Other" create
      exit
    exit
```

```

        export-id 8
    exit
    app-group "Peer to Peer" create
        export-id 3
    exit
    app-group "Remote Connectivity" create
        export-id 4
    exit
    app-group "Unknown"
        charging-group "0_rated"
        export-id 1
    exit
    app-group "Web" create
        export-id 10
    exit
    application "DNS" create
        description "default-description for application DNS"
        app-group "Other"
        export-id 12
    exit
    application "BitTorrent" create
        app-group "Peer to Peer"
        export-id 3
    exit
    application "HTTP" create
        description "default-description for application HTTP"
        app-group "Web"
        export-id 26
    exit
    application "RDP" create
        description "default-description for application RDP"
        app-group "Remote Connectivity"
        export-id 61
    exit
    application "Unknown"
        charging-group "0_rated"
        export-id 1
    exit
    exit
    commit
    exit
    exit
    exit

```

Example app-filter definitions defining HTTP, DNS, Bittorrent and RDP applications are as follows.

```

configure
  application-assurance
    group 129:34883
      policy
        begin
          app-filter
            entry 6 create
              description "default-description for AppFilter entry 6"
              protocol eq "rdp"
              ip-protocol-num eq tcp
              application "RDP"
              no shutdown
            exit
            entry 9 create
              description "default-description for AppFilter entry 9"
              protocol eq "dns"
              ip-protocol-num eq udp

```



```
server-port eq range 53 55
application "DNS"
no shutdown
exit
entry 20 create
description "default-description for AppFilter entry 20"
protocol eq "bittorrent"
ip-protocol-num eq tcp
application "BitTorrent"
no shutdown
exit
entry 38 create
description "default-description for AppFilter entry 38"
protocol eq "http"
ip-protocol-num eq tcp
server-port gt 8738
application "HTTP"
no shutdown
exit
exit
commit
exit
exit
exit
```

**Note:**

The focus of this example is on the definition of app-filters and/or AQPs. These are listed above (and below) for illustration purposes. The "sample" AQP configurations and app-filters shown here should not be used in a real-life configuration. Their configuration should follow the information in [Application Assurance — Application Identification and User-Defined Applications](#).

Example AQP configurations for blocking DNS, RDP and HTTP traffic are as follows.

```
configure
application-assurance
group 129:34883
policy
begin
app-qos-policy
entry 2 create
match
application eq "DNS"
characteristic "permitDNS" eq "no"
exit
action
drop
exit
no shutdown
exit
entry 3 create
match
application eq "HTTP"
characteristic "permitHTTP" eq "no"
ip-protocol-num neq 0
exit
action
drop
exit
no shutdown
exit
entry 4 create
```

```

        match
        application eq "RDP"
        app-group eq "Remote Connectivity"
        characteristic "permitRDP" eq "no"
        ip-protocol-num neq udp
    exit
    action
        drop
    exit
    no shutdown
exit
    exit
    commit
exit
    exit
exit

```

When an ESM subscriber is created, it is associated with the default AA app-profile, as seen using the show command below.

```

*A:BNG-1>show>app-assure>group# aa-sub esm "sub_172.16.0.2" summary
=====
Application-Assurance Subscriber Summary (realtime)
=====
AA-Subscriber      : sub_172.16.0.2 (esm)
ISA assigned       : 3/2
App-Profile        : app_prof_1
App-Profile divert : Yes
Capacity cost      : 1
Aarp Instance Id   : N/A
HTTP URL Parameters : (Not Specified)
Last HTTP Notified Time : N/A
-----
Traffic                Octets                Packets                Flows
-----
From subscriber:
  Admitted              0                    0                    0
  Denied                0                    0                    0
  Active flows          0
To subscriber:
  Admitted              0                    0                    0
  Denied                0                    0                    0
  Active flows          0
Flow counts:
  Terminated          0
  Short duration        0
  Med duration          0
  Long duration         0
Total flow duration : 0 seconds
-----
Top App-Groups                Octets                Packets                Flows
-----
None
-----
Application Service Options (ASO)
-----
Characteristic            Value                Derived from
-----
permitDNS                  yes                  default
permitRDP                  yes                  default
permitHTTP                 yes                  default
=====
*A:BNG-1>show>app-assure>group#

```

After the PCRF sends out AppProfile and ASO override AVPs, using either PCC or ADC rules, in RAR messages (as shown below) it can be seen that the new parameters (new profile and new values for permitDNS and permitHTTP ASOs) are updated for that ESM subscriber.

Figure 48: PCRF AVPs override call flow diagram




Figure 49: RAR containing ASOs and AppProfile override AVPs example

```

adc-rule-install (1092) V----- [184]
  vendor-id TGPP
  data [172] (Grouped)
    adc-rule-definition (1094) V----- [172]
      vendor-id TGPP
      data [160] (Grouped)
        adc-rule-name (1096) V----- [32]
          vendor-id TGPP
          data [20] (UTF8String) : aa-functions:appprof
        AA-Functions (1001) V----- [128]
          vendor-id ALU
          data [116] (Grouped)
            AA-Profile-Name (1002) V----- [17]
              vendor-id ALU
              data [5] (UTF8String) : app_prof 2
            AA-App-Service-Options (1003) V----- [48]
              vendor-id ALU
              data [36] (Grouped)
                AA-App-Serv-Options-Name (1004) V----- [17]
                  vendor-id ALU
                  data [5] (UTF8String) : permitDNS
                AA-App-Serv-Options-Value (1005) V----- [16]
                  vendor-id ALU
                  data [4] (UTF8String) : no
            AA-App-Service-Options (1003) V----- [48]
              vendor-id ALU
              data [36] (Grouped)
                AA-App-Serv-Options-Name (1004) V----- [18]
                  vendor-id ALU
                  data [6] (UTF8String) : permitHTTP
                AA-App-Serv-Options-Value (1005) V----- [14]
                  vendor-id ALU
                  data [2] (UTF8String) : no
          
```

al_0661

```

*A:BNG-1>show>app-assure>group# aa-sub esm "sub_172.16.0.2" summary
=====
Application-Assurance Subscriber Summary (realtime)
=====
AA-Subscriber      : sub_172.16.0.2 (esm)
ISA assigned       : 3/2
App-Profile        : app_prof_2
App-Profile divert : Yes
Capacity cost      : 1
Aarp Instance Id   : N/A
HTTP URL Parameters : (Not Specified)
Last HTTP Notified Time : N/A
=====
    
```

| Traffic | Octets | Packets | Flows |
|-----------------------------------|--------|--------------|-------|
| ----- | | | |
| From subscriber: | | | |
| Admitted | 0 | 0 | 0 |
| Denied | 0 | 0 | 0 |
| Active flows | | | 0 |
| To subscriber: | | | |
| Admitted | 0 | 0 | 0 |
| Denied | 0 | 0 | 0 |
| Active flows | | | 0 |
| Flow counts: | | | |
| Terminated | | | 0 |
| Short duration | | | 0 |
| Med duration | | | 0 |
| Long duration | | | 0 |
| Total flow duration : 0 seconds | | | |
| ----- | | | |
| Top App-Groups | Octets | Packets | Flows |
| ----- | | | |
| None | | | |
| ----- | | | |
| Application Service Options (ASO) | | | |
| ----- | | | |
| Characteristic | Value | Derived from | |
| ----- | | | |
| permitDNS | no | dyn-override | |
| permitRDP | yes | default | |
| permitHTTP | no | dyn-override | |
| ===== | | | |

4. Configuring AA usage monitoring

Once the applications, application groups and/or charging groups are defined and configured (see previous steps), the operator needs:

- to enable the collection of per-subscriber statistics so they can be used for Gx based usage-monitoring. This step is not needed for any app/appgrp or charging group that is already enabled for per-subscriber statistics. In other words, if XML or RADIUS accounting is enabled for a given app/appgrp or charging group, then Gx usage-monitoring is also automatically enabled.
- to enable usage-monitoring for the given AA group:partition.

```
configure
  application-assurance
    group 129:34883
      statistics
        aa-sub
          usage-monitoring
          app-group "Unknown" export-using accounting-policy
                                radius-accounting-policy
          charging-group "0_rated" export-using accounting-policy
                                radius-accounting-policy
          charging-group "default_charge_group" export-using
                                accounting-policy
          radius-accounting-policy
          application "BitTorrent" no-export
        exit
```

In the preceding example:

- The usage-monitoring command is used to enable Gx usage monitoring for the specified AA partition.

- The aa-group and charging-group commands specify which charging groups and AA groups are selected for export. In this case *0-rated*, *Unknown*, and *default_charge_group* are selected for RADIUS accounting and they automatically qualify for Gx-usage monitoring.
- The BitTorrent application however needs to be explicitly configured as “no-export” as it needs to be enabled for Gx-usage monitoring.

The operator can display the number of usage monitoring rules for a given subscriber. This is shown below after the ESM subscriber is created, but before any ADC rules are installed for usage-monitoring by PCRF, so AA reports that no rules apply (“0”).

```
*A:BNG-1>show>app-assure>group# aa-sub esm "alcatel_A_1" usage-monitor status
=====
Application-Assurance Subscriber "alcatel_A_1" (esm)
Usage Monitor Status
=====
Type           Name           Type           Oper Status
-----
No. of rules: 0
=====
*A:BNG-1>show>app-assure>group#
```

The PCRF then sends a RAR message with a usage monitoring ADC or PCC rule for the BitTorrent application to set the usage thresholds for BitTorrent for the ESM subscriber “alcatel_A_1” to (in bytes):

Input (from sub) 1378168

Output (to sub) 1381148

Total traffic (up and down) 18446744073709551614

Figure 50: RAR containing usage monitoring ADC rules example

```

adc-rule-install (1092) V----- [96]
  vendor-id TGPP
  data [84] (Grouped)
    adc-rule-definition (1094) V----- [84]
      vendor-id TGPP
      data [72] (Grouped)
        adc-rule-name (1096) V----- [20]
          vendore-id TGPP
          data [8] (UTF8String) whatever
          tdf-application-id (1088) V-----[22]
            vendor-id ALU
            data [10] (UTF8String) : BitTorrent
            monitoring-key (1066) V----- [25]
              vendor-id TGPP
              data [13] (UTF8String) : torrentmonkey

usage-monitoring-information (1067) V----- [80]
  vendor-id TGPP
  data [68] (Grouped)
    monitoring-key (1066) V----- [25]
      vendor-id TGPP
      data [13] (UTF8String) : torrentmonkey
    granted-service-units (431) ----- [24]
      data [16] (Grouped)
        cc-input-octets (412) ----- [16]
          data [8] (Unsigned64) : 1378168
        cc-output-octets (414) ----- [16]
          data [8] (Unsigned64) : 1378168
        cc-total-octets (421) ----- [16]
          data [8] (Unsigned64) : 18446744073709551614
    monitoring-key (1068) V----- [16]
      vendor-id TGPP
      data [4] (Enumerated) : 2 : ADC RULE LEVEL
    
```

al_0662

This is then reflected on the AA-ISA:

```

*A:BNG-1>show>app-assure>group# aa-sub esm "alcatel_A_1" usage-monitor status
=====
Application-Assurance Subscriber "alcatel_A_1" (esm)
Usage Monitor Status
=====
Type                Name                Oper Status
-----
application         BitTorrent          active
=====
No. of rules: 1
=====
*A:BNG-1>show>app-assure>group#
    
```

Note the "active" oper status is set since there is at least one usage monitoring threshold associated with this application.

Given that there is no traffic flowing yet to or from the subscriber the counters currently are "0":

```

*A:BNG-1>show>app-assure>group# aa-sub esm "alcatel_A_1" usage-monitor count
=====
Application-Assurance Subscriber "alcatel_A_1" (esm)
Usage Monitor Credit Statistics
=====
Application: "BitTorrent"
Direction  Status                Granted                Used                % Used
-----
    
```

```

to sub      valid          1378168          0          0%
from sub    valid          1381148          0          0%
both        valid          18446744073709551614  0          0%
=====
*A:BNB-1>show>app-assure>group#

```

The status is set to "valid" since a threshold (or Grant) is received.

When, at a later stage, traffic starts flowing again usage-monitor subscriber statistics are updated as shown below.

```

*A:BNB-1>show>app-assure>group# aa-sub esm "alcatel_A_1" usage-monitor count
=====
Application-Assurance Subscriber "alcatel_A_1" (esm)
Usage Monitor Credit Statistics
=====
Application: "BitTorrent"
Direction      Status          Granted          Used          % Used
-----
to sub         valid          1378168          137816         10%
from sub       valid          1381148          13781          1%
both           valid          18446744073709551614  151597         5%
=====
*A:BNB-1>show>app-assure>group#

```

The PCRF can also at the same time set ADC or PCC rules for other applications (such as the *0_rated* and the *default_charging_group* charging groups).

In the following case, the PCRF installs an ADC usage monitoring rule for:

- Charging group: "0-rated", but without usage thresholds
- Charging group: "default_charge_group", and sets only a threshold for "to sub" traffic.

This results in having a usage policy for the "0-rated" charging group installed but this is not active since there are no grants associated with it:

```

*A:BNB-1>show>app-assure>group# aa-sub esm "alcatel_A_1" usage-monitor status
=====
Application-Assurance Subscriber "alcatel_A_1" (esm)
Usage Monitor Status
=====
Type          Name              Oper Status
-----
application   BitTorrent        active
charging-group 0_rated           inactive
charging-group default_charge_group active
-----
No. of rules: 3
=====
*A:BNB-1>show>app-assure>group#

```

Note that the "inactive" status for the "0-rated" charging group is due to no grants being received.

Moreover, detailed counters show:

```

*A:BNB-1>show>app-assure>group# aa-sub esm "alcatel_A_1" usage-monitor count
=====
Application-Assurance Subscriber "alcatel_A_1" (esm)
Usage Monitor Credit Statistics
=====
Application: "BitTorrent"

```

```

Direction      Status      Granted      Used      % Used
-----
to sub         valid      1378168     137816     10%
from sub       valid      1381148     13781      1%
both           valid      18446744073709551614  151597     5%
-----
Charging-Group: "0_rated"
Direction      Status      Granted      Used      % Used
-----
to sub         invalid     n/a         0          n/a
from sub       invalid     n/a         0          n/a
both           invalid     n/a         0          n/a
-----
Charging-Group: "default_charge_group"
Direction      Status      Granted      Used      % Used
-----
to sub         valid      1000000     1378084    100%
from sub       invalid     n/a         1574       n/a
both           invalid     n/a         1379658    n/a
=====
*A: BNG-1>show>app-assure>group#
    
```

Again, the "invalid" status above reflects the fact that no grants have been received.

Conclusion

The introduction of the diameter (/Gx) control feature on the 7x50 BNG enables operators to consolidate policy management systems used in wire-line and wireless environments into a single system. This provides an increase in operational efficiency as mobile and fixed networks convergence gains more traction.

This example illustrates how policy control and usage monitoring of the 7x50 BNG Application Assurance services can be achieved over standard 3GPP Diameter Gx protocol.

Deterministic Large Scale NAT44

This chapter provides information about deterministic large scale NAT44 configurations.

Topics in this chapter include:

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

Applicability

The information and configuration in this chapter was initially based on SR OS Release 11.0.R3, and is updated to SR OS Release 14.0.R4.

Overview

Deterministic Network Address Translation (NAT) is a mode of operation where mappings between the NAT subscriber and the outside IP address and port range are allocated at the time of configuration.

In deterministic NAT for Large Scale NAT IPv4-to-IPv4 (LSN44) subscribers, each LSN44 subscriber is permanently mapped to an outside IP address and a dedicated (deterministic) port-block based on a specific algorithm.

Logging is not needed in this case because the reverse mapping can be obtained using the reverse of the preceding algorithm.


A deterministic LSN44 subscriber can have only one deterministic port-block that can (optionally) be extended by one or multiple dynamic port-blocks in case all ports in deterministic port-block are exhausted.

In case an LSN44 subscriber has been assigned both deterministic and dynamic port blocks, logging for the dynamic port-block allocation/de-allocation is required.

A scalable logging solution for dynamic port-blocks is achievable using RADIUS or IPFIX.

Logging for dynamic port-blocks is out of the scope of this chapter.

Figure 51: Deterministic NAT Mapping



26145


Algorithm

The deterministic NAT algorithm makes a predictable mapping between the (inside IP, routing instance) and the (outside IP, routing instance, deterministic port block).

The algorithm is revertive, meaning that a given (outside IP, routing instance, deterministic port block) will derive a given (inside IP, routing Instance).

The algorithm is loosely based on draft-donley-behave-deterministic-cgn-00.txt, which allows for the dynamic expansion of the port-blocks once the ports in the original deterministic port-block are exhausted.

Figure 52: Deterministic NAT Algorithm



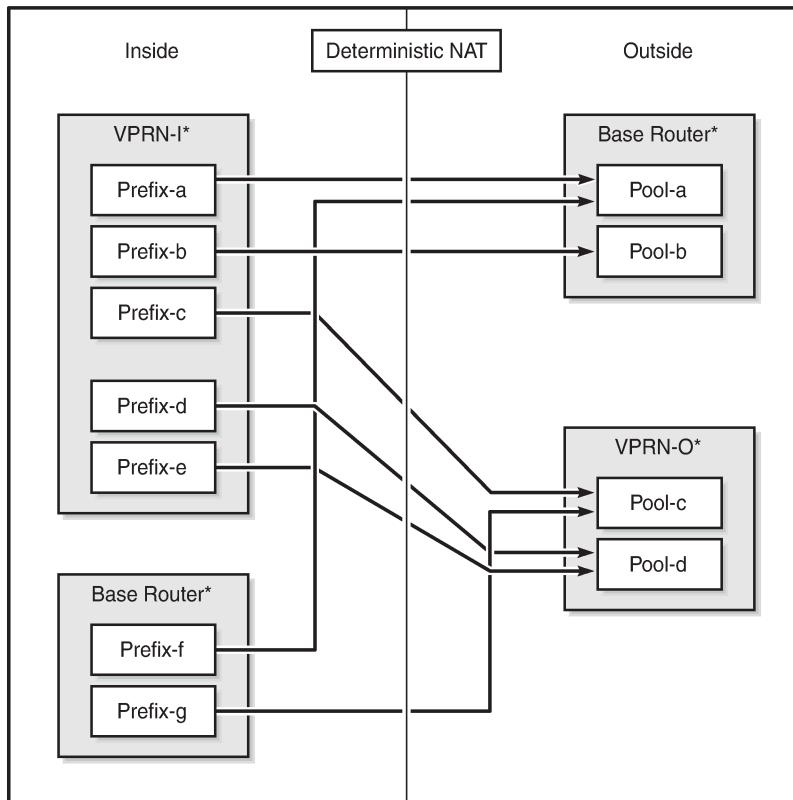
26146

Deterministic mapping

Any inside prefix in any routing instance can be mapped to any pool in any routing instance.

In deterministic NAT, prefixes from multiple routing instances can be mapped to the same outside pool, also prefixes from a single inside routing instance can be selectively mapped to different outside pools.

Figure 53: Deterministic Mapping: Inside -> Outside Routing Instances



*Routing-Based NAT cannot be used if inside/outside routing instances are the same

26147

Mapping rules

A deterministic LSN44 subscriber is mapped to only one deterministic block which can further be extended to multiple dynamic blocks if ports within the deterministic block are exhausted.

The subscriber-limit is the number of subscribers that can be deterministically mapped to one outside IP address (i.e. compression ratio) and **must** be a power of 2.

The total number of deterministic ports (DetP) per outside IP address is determined by this subscriber-limit and the number of deterministic ports per subscriber.

The remaining ports (DynP) beyond the deterministic port range up to 65535 will be dedicated for dynamic use when a deterministic block is exhausted.

Every host using an inside prefix is guaranteed one dedicate block in the deterministic port ranges.

If the inside prefix length is $m < 32-n$, where $2^n = \text{subscriber-limit}$, then the prefix must be broken into pieces so that all hosts (subscriber-limit) in each piece maps exactly to one outside IP address.


For example, if there is an inside prefix 192.168.0.0/23, with $m=23$ and a maximum number of 256 hosts; and the subscriber-limit set to 256, then $n=8$. This results in $23 < 24 (32-8)$ and so this inside prefix has to be broken into 2 pieces, in other words, this inside prefix will fit into 2 outside IP addresses, each of 256 port-blocks.

In case that the prefix length is $m \geq 32-n$, where $2^n = \text{subscriber-limit}$, then all hosts from the configured prefix are mapped to the same outside IP.

For example, if there is an inside prefix 192.168.1.0/25, with $m=25$ and a maximum number of 128 hosts, and the subscriber-limit set to 256, then $n=8$. This results in $25 > 24 (32-8)$, so definitely 128 hosts can fit in one outside IP because there are 256 available port-blocks, in other words, this inside prefix will fit into one outside IP where 128 blocks have been used out of the 256 port-blocks available, and the rest (256-128) are wasted.

Overbooking of the outside address pool is not supported in deterministic NAT.


Figure 54: Deterministic mapping: outside IP port-blocks/ranges



26148

Configuration

Figure 55: Example topology



26149

Configuration prerequisites

Card and MDA configuration.

```
configure
  card 2
    card-type iom3-xp
    mda 1
      mda-type isa-bb
      no shutdown
    exit
    mda 2
      mda-type isa-bb
      no shutdown
    exit
  no shutdown
exit
exit
```



Note:

Private address ranges are used in outside pools within this chapter but normally public address ranges would be used.

Create the NAT group, and add the MS-ISAs created above to the NAT group; up to 10 MS-ISAs of type isa-bb can be configured under the NAT group.

```
configure
  isa
    nat-group 1 create
      mda 2/1
      mda 2/2
      active-mda-limit 1
      no shutdown
    exit
  exit
exit
exit
```

Configuration commands

A NAT **outside pool** is configured using the following command:

```
configure {router | service vprn <service-id>}
  nat
    outside
      pool <nat-pool-name> [nat-group <nat-group-id> type <pool-type> create]
      port-reservation {blocks <num-blocks> | ports <num-ports>}
      port-forwarding-range <range-end>
      subscriber-limit <subscriber-limit>
      deterministic
      port-reservation <det-num-ports>
    exit
      address-range <start-ip-address> <end-ip-address> create
    exit
  exit
exit
exit
exit
```

where:

nat-pool-name — Specifies the name of the NAT pool up to 32 characters max.

nat-group-id — Specifies the NAT group ID. The values are 1 — 4.

pool-type — Species the pool type (**large-scale**).

num-blocks — Specifies the number of dynamic port-blocks per outside IP address. The values are 1 — 64512

num-ports — Specifies the number of ports per dynamic block. The values are 1 — 32256

range-end — Specifies the upper limit of the port range available for static port forwarding. The values are 1023 — 65535

subscriber-limit — Specifies the maximum number of subscribers per outside IP address.

A power of 2 (2^n) number for deterministic NAT

[1,2,4,8,16,32,64,128,256,512,1024,2048, 4096, 8192,16348, 32768]

1..65535 for non-deterministic NAT

default: 65535 for non-deterministic

det-num-ports — Specifies the number of ports in a deterministic port block that is allocated and dedicated to a single subscriber during the configuration phase. The values are 1..65535

start-ip-address — Specifies the first outside IP address in the a.b.c.d format.

end-ip-address — Specifies the last outside IP address in the a.b.c.d format.



Note:

- When the subscriber limit equals 1, each subscriber is mapped to a single outside IP address, though the NAPT (port translation) function is still performed.
- 1:1 NAT mode in combination with deterministic NAT is not supported.

A NAT **policy** is configured using the following command:

```
configure service nat
  nat-policy <nat-policy-name> [create]
    block-limit <[1..40]>
    pool <nat-pool-name> {router <router-instance> | service-name <service-name>}
  exit
```

where:

nat-policy-name — Specifies the NAT policy name up to 32 characters max.

block-limit —The maximum number of deterministic plus dynamic port blocks that can be assigned to a single inside IP address. In other words, the maximum number of dynamic port blocks that can be assigned to an inside IP address when the deterministic port block is exhausted equals (block-limit - 1).

nat-pool-name — Specifies the NAT pool name up to 32 characters max.

router-instance — Specifies the router instance the pool belongs to, either by router name or service ID.

<router-name> | *<service-id>*

The router name values are *Base* or *service-id* [1..2147483647]

service-name — Specifies the name of the service up to 64 characters max.

A NAT **inside prefix** is configured using the following command:

```
configure [router] service vprn <service-id>
  nat
    inside
      classic-lsn-max-subscriber-limit <max>
      deterministic
      prefix <ip-prefix/length> subscriber-type <nat-sub-type>
      nat-policy <nat-policy-name> create
        map start <lsn-sub-address> end <lsn-sub-address> to <outside-ip-address>
        no shutdown
      exit
    exit
  exit
exit
```

where:

max — The power of 2 (2^n) number that must match the largest subscriber limit number in a deterministic pool referenced from this inside routing instance. The range for this command is the same as the subscriber-limit command under the pool hierarchy. The values are 1,2,4,8 — 32768

ip-prefix/length — A prefix on the inside encompassing subscribers that will be deterministically mapped to an outside IP address and port block in the corresponding pool.

| | |
|----------------------|---|
| <ip-prefix/ip-pref*> | <ipv4-prefix>/<ipv4-prefix-length> <ipv6-prefix>/<ipv6-prefix-length> |
| <ipv4-prefix> | a.b.c.d (host bits must be 0) |
| <ipv4-prefix-length> | [0..32] |
| <ipv6-prefix> | x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x - [0..FFFF]H d - [0..255]D |
| <ipv6-prefix-length> | [0..128] |
| <nat-sub-type>: | classic-lsn-sub |
| <nat-policy-name> | Specifies a NAT policy name up to 32 characters in length. |

Following rules apply to the **classic-lsn-max-subscriber-limit**:

- Should be greater than or equal to the largest subscriber-limit of all pools referenced by the NAT policies within the corresponding inside routing instance.
- Must be configured before any inside prefix configuration.
- Must be 2^n and affects the ingress hashing of deterministic subscribers and also non-deterministic subscribers in case both are configured under the same inside router instance.

Three cases are now configured to demonstrate the use of deterministic and dynamic port-block usage:

- **Case 1:** Mapping multiple prefixes from the same VRF (VPRN 15001) into the same outside pool, routing instance "Base".
- **Case 2:** Mapping multiple prefixes from the same VRF (VPRN 15001) into different outside pools, routing instance VPRN 15002
- **Case 3:** Mapping overlapping prefixes from different VRFs (VPRN 15001 and VPRN 15002) into the same outside pool, routing instance "Base".

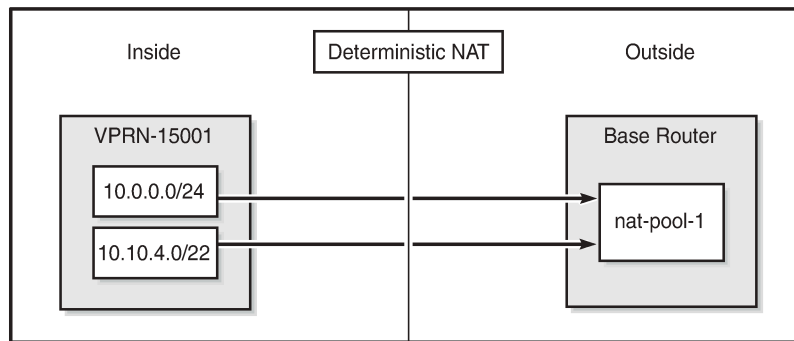
In each case all of the traffic is NATed.

Case 1

Configured with:

- Mapping multiple prefixes of the same VRF into the same outside pool.
- NAT all traffic.

Figure 56: Case 1



26150

The NAT **outside pool** is configured as follows:

```

configure
router
nat
  outside
    pool "nat-pool-1" nat-group 1 type large-scale create
    port-reservation ports 180
    port-forwarding-range 4023
    subscriber-limit 128
    deterministic
    port-reservation 300
    exit
    address-range 192.168.0.1 192.168.0.100 create
    exit
    no shutdown
  exit
exit
exit
exit
exit
exit

```


The NAT **policy** is configured as follows:

```
configure
  service
    nat
      nat-policy "nat-policy-1" create
      block-limit 4
      pool "nat-pool-1" router Base
    exit
  exit
exit
```

The NAT **inside prefixes** are configured as follows:

```
configure
  service
    vprn 15001 customer 1 create
    nat
      inside
        destination-prefix 0.0.0.0/0
        classic-lsn-max-subscriber-limit 256
        deterministic
        prefix 10.0.0.0/24 subscriber-type classic-lsn-sub
        nat-policy "nat-policy-1" create
        map start 10.0.0.0 end 10.0.0.255 to 192.168.0.1
        no shutdown
        exit
        prefix 10.10.4.0/22 subscriber-type classic-lsn-sub
        nat-policy "nat-policy-1" create
        map start 10.10.4.0 end 10.10.7.255 to 192.168.0.3
        no shutdown
        exit
      exit
    exit
  exit
  no shutdown
exit
exit
```

map statements are automatically created when the prefix is created and it is **no shutdown**.

Show commands

The subscriber-limit is set to 128 for the 10.0.0.0/24 prefix, so it is broken into two smaller /25 prefixes each. Each of these smaller prefixes are mapped into a specific outside IP address.

To show the first Large Scale NAT (LSN) subscriber of the first /25 prefix for inside routing instance 15001, the following command can be used:

```
*A:PE1# show service nat lsn-subscribers inside-router 15001 inside-ip 10.0.0.0

=====
NAT LSN subscribers
=====
Subscriber          : [LSN-Host@10.0.0.0]
NAT policy          : nat-policy-1
Subscriber ID       : 276824064
-----
```

```
Type : classic-lsn-sub
Inside router : 15001
Inside IP address prefix : 10.0.0.0/32
ISA NAT group : 1
ISA NAT group member : 1
Outside router : "Base"
Outside IP address : 192.168.0.1

-----
No. of LSN subscriber instances: 1
=====
*A:PE1#
```

The last subscriber mapping to the same 192.168.0.1 outside IP address has inside address 10.0.0.127. To show the first LSN subscriber of the second /25 prefix for inside routing instance 15001, the following command can be used:

```
*A:PE1# show service nat lsn-subscribers inside-router 15001 inside-ip 10.0.0.128

=====
NAT LSN subscribers
=====
Subscriber : [LSN-Host@10.0.0.128]
NAT policy : nat-policy-1
Subscriber ID : 276824192
-----
Type : classic-lsn-sub
Inside router : 15001
Inside IP address prefix : 10.0.0.128/32
ISA NAT group : 1
ISA NAT group member : 1
Outside router : "Base"
Outside IP address : 192.168.0.2

-----
No. of LSN subscriber instances: 1
=====
*A:PE1#
```

The last subscriber mapping to the same 192.168.0.2 outside IP address has inside address 10.0.0.255. To show the base router LSN blocks corresponding to the first inside IP address within the 10.0.0.0/24 prefix, the following command can be used:

```
*A:PE1# show router nat lsn-blocks inside-ip 10.0.0.0

=====
Large-Scale NAT blocks for Base
=====
192.168.0.1 [4024..4323]
Pool : nat-pool-1
Policy : nat-policy-1
Started : 2016/10/27 11:18:59
Inside router : vprn15001
Inside IP address : 10.0.0.0

-----
Number of blocks: 1
=====
*A:PE1#
```

To show the base router LSN blocks corresponding to the last inside IP address within the 10.0.0.0/24 prefix, the following command can be used:

```
*A:PE1# show router nat lsn-blocks inside-ip 10.0.0.255

=====
Large-Scale NAT blocks for Base
=====
192.168.0.2 [42124..42423]
Pool                : nat-pool-1
Policy              : nat-policy-1
Started             : 2016/10/27 11:18:59
Inside router       : vprn15001
Inside IP address   : 10.0.0.255

-----
Number of blocks: 1
=====
*A:PE1#
```

The subscriber-limit is 128 for the 10.10.4.0/22 prefix, so it is broken into eight /25 prefixes. Each of these smaller prefixes are mapped into a specific outside IP address.

To show the first LSN subscriber of the first /25 prefix for inside routing instance 15001, the following command can be used:

```
*A:PE1# show service nat lsn-subscribers inside-ip 10.10.4.0

=====
NAT LSN subscribers
=====
Subscriber          : [LSN-Host@10.10.4.0]
NAT policy          : nat-policy-1
Subscriber ID       : 276824320

-----
Type                : classic-lsn-sub
Inside router       : 15001
Inside IP address prefix : 10.10.4.0/32
ISA NAT group       : 1
ISA NAT group member : 1
Outside router      : "Base"
Outside IP address   : 192.168.0.3

-----
No. of LSN subscriber instances: 1
=====
*A:PE1#
```

The last subscriber mapping to the same 192.168.0.3 outside IP address has inside address 10.10.4.127.

To show the first LSN subscriber of the last /25 prefix for inside routing instance 15001, the following command can be used:

```
*A:PE1# show service nat lsn-subscribers inside-ip 10.10.7.128

=====
NAT LSN subscribers
=====
Subscriber          : [LSN-Host@10.10.7.128]
NAT policy          : nat-policy-1
Subscriber ID       : 276825216

-----
```

```
Type : classic-lsn-sub
Inside router : 15001
Inside IP address prefix : 10.10.7.128/32
ISA NAT group : 1
ISA NAT group member : 1
Outside router : "Base"
Outside IP address : 192.168.0.10
```

```
-----
No. of LSN subscriber instances: 1
=====
```

```
*A:PE1#
```

To show the base router LSN blocks corresponding to the first inside IP within 10.10.4.0/24 prefix, the following command can be used:

```
*A:PE1# show router nat lsn-blocks inside-ip 10.10.4.0
```

```
=====
Large-Scale NAT blocks for Base
=====
```

```
192.168.0.3 [4024..4323]
Pool : nat-pool-1
Policy : nat-policy-1
Started : 2016/10/27 11:18:59
Inside router : vprn15001
Inside IP address : 10.10.4.0
```

```
-----
Number of blocks: 1
=====
```

```
*A:PE1#
```

To show the base router LSN blocks corresponding to the last inside IP within 10.10.4.0/24 prefix, the following command can be used:

```
*A:PE1# show router nat lsn-blocks inside-ip 10.10.7.255
```

```
=====
Large-Scale NAT blocks for Base
=====
```

```
192.168.0.10 [42124..42423]
Pool : nat-pool-1
Policy : nat-policy-1
Started : 2016/10/27 11:18:59
Inside router : vprn15001
Inside IP address : 10.10.7.255
```


```
-----
Number of blocks: 1
=====
```

```
*A:PE1#
```

Mapping results

According to this configuration, each inside IP address has one deterministic block of 300 ports and can have up to three dynamic blocks (block-limit = 4) each of 180 ports, allowing a maximum of $300+3*180 = 840$ flows.


Figure 57: Case 1 results



Sending flows

For the inside IP 10.0.0.1, several UDP flows will be sent and both the deterministic and dynamic blocks mappings will be verified.

Figure 58: Case 1 flows



26152

When sending 300 UDP flows or less, all flows are mapped to a single deterministic block because the number of ports in a deterministic block is 300. There is no logging; because no dynamic blocks are used, only the deterministic block is used.

To show LSN blocks on the outside routing instance *Base* and the outside ports allocated for the inside IP 10.0.0.1, the following command can be used:

```
*A:PE1# show router nat lsn-blocks inside-ip 10.0.0.1
```

```
=====
Large-Scale NAT blocks for Base
=====
```

```
192.168.0.1 [4324..4623]
Pool                : nat-pool-1
Policy              : nat-policy-1
Started             : 2016/10/27 11:18:59
Inside router       : vprn15001
Inside IP address   : 10.0.0.1
```

```
-----
Number of blocks: 1
=====
```

```
*A:PE1#
```

When increasing the number of flows such that: $301 < \text{number of flows} < 480$

- In addition to the deterministic block (300 ports), there will be an extension by 1 dynamic block of 180 ports (port-reservation=180).
- Logging occurs for the dynamic port-block.

To show the base router LSN blocks and the outside ports allocated to the inside IP address 10.0.0.1, the following command can be used:

```
*A:PE1# show router nat lsn-blocks inside-ip 10.0.0.1

=====
Large-Scale NAT blocks for Base
=====
192.168.0.1 [4324..4623]
Pool                : nat-pool-1
Policy              : nat-policy-1
Started             : 2016/10/27 11:18:59
Inside router       : vprn15001
Inside IP address   : 10.0.0.1
```

```
192.168.0.1 [44044..44223]
Pool                : nat-pool-1
Policy              : nat-policy-1
Started             : 2016/10/28 12:40:41
Inside router       : vprn15001
Inside IP address   : 10.0.0.1
```

```
-----
Number of blocks: 2
=====
```

```
*A:PE1#
```

Logging is verified using Log 99 (in case event-control *nat* events are generated) which shows the mapping details to the new dynamic block as follows:

```
2 2016/10/28 12:40:41.51 UTC MINOR: NAT #2012 Base NAT
"{12} Map 192.168.0.1 [44044-44223] MDA 2/1 -- 276824065 classic-lsn-sub %1 vprn15001 10.0.0.1
at 2016/10/28 12:40:41"
```

When increasing the number of flows such that: $481 < \text{number of flows} < 660$

- In addition to the deterministic block (300 ports), there will be an extension by 2 dynamic blocks of 180 ports each.
- Logging occurs for the dynamic port-blocks.

To show LSN blocks on the outside routing instance *Base* and the outside ports allocated for the inside IP 10.0.0.1, the following command is used:

```
*A:PE1# show router nat lsn-blocks inside-ip 10.0.0.1

=====
Large-Scale NAT blocks for Base
=====
192.168.0.1 [4324..4623]
Pool                : nat-pool-1
Policy              : nat-policy-1
Started             : 2016/10/27 11:18:59
Inside router      : vprn15001
Inside IP address  : 10.0.0.1

192.168.0.1 [44044..44223]
Pool                : nat-pool-1
Policy              : nat-policy-1
Started             : 2016/10/28 12:40:41
Inside router      : vprn15001
Inside IP address  : 10.0.0.1

192.168.0.1 [44224..44403]
Pool                : nat-pool-1
Policy              : nat-policy-1
Started             : 2016/10/28 12:41:52
Inside router      : vprn15001
Inside IP address  : 10.0.0.1

-----
Number of blocks: 3
=====
*A:PE1#
```

Logging is verified using Log 99 (in case event-control *nat* events are generated) which shows the mapping details to the new dynamic block as follows:

```
3 2016/10/28 12:41:52.66 UTC MINOR: NAT #2012 Base NAT
"{13} Map 192.168.0.1 [44224-44403] MDA 2/1 -- 276824065 classic-lsn-sub %1 vprn15001 10.0.0.1
at 2016/10/28 12:41:52"
```

When increasing the number of flows such that $:661 < \text{number of flows} < 840$

- In addition to the deterministic block (300 ports), there will be an extension by 3 dynamic blocks of 180 ports each.
- Logging occurs for the dynamic port-blocks.

To show LSN blocks on the outside routing instance "Base" and the outside ports allocated for the inside IP 10.0.0.1, the following command can be used:

```
*A:PE1# show router nat lsn-blocks inside-ip 10.0.0.1

=====
Large-Scale NAT blocks for Base
=====
192.168.0.1 [4324..4623]
Pool                : nat-pool-1
Policy              : nat-policy-1
Started             : 2016/10/27 11:18:59
Inside router      : vprn15001
Inside IP address  : 10.0.0.1
```

```

192.168.0.1 [44044..44223]
Pool                : nat-pool-1
Policy              : nat-policy-1
Started             : 2016/10/28 12:40:41
Inside router      : vprn15001
Inside IP address  : 10.0.0.1

192.168.0.1 [44224..44403]
Pool                : nat-pool-1
Policy              : nat-policy-1
Started             : 2016/10/28 12:41:52
Inside router      : vprn15001
Inside IP address  : 10.0.0.1

192.168.0.1 [44404..44583]
Pool                : nat-pool-1
Policy              : nat-policy-1
Started             : 2016/10/28 12:43:46
Inside router      : vprn15001
Inside IP address  : 10.0.0.1
    
```

```
-----
Number of blocks: 4
=====
```

```
*A:PE1#
```

Logging is verified using Log 99 (in case event-control *nat* events are generated) which shows the mapping details to the new dynamic block as follows:

```

4 2016/10/28 12:43:46.71 UTC MINOR: NAT #2012 Base NAT
"{14} Map 192.168.0.1 [44404-44583] MDA 2/1 -- 276824065 classic-lsn-sub %1 vprn15001 10.0.0.1
at 2016/10/28 12:43:46"
    
```

When increasing number of flows such that the number of flows > 840

- No more extension by dynamic blocks (block-limit = 4) allowed.
- Any flows more than 840 will be dropped and the relevant NAT statistics incremented.

To verify NAT statistics, first check the NAT group/member and MS-ISA associated with the outside IP 192.168.0.1/32:

```
*A:PE1# show router route-table 192.168.0.1/32
```

```
=====
Route Table (Router: Base)
=====
```

| Dest Prefix[Flags] Next Hop[Interface Name] | Type | Proto | Age | Metric | Pref |
|--|--------|-------|-----------|--------|------|
| 192.168.0.1/32 NAT outside to mda 2/1 | Remote | NAT | 01d01h26m | 0 | 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:PE1#
```


To check which group/member does this MS-ISA belong to, the following command can be used:

```
*A:PE1# show isa nat-group 1 members

=====
ISA Group 1 members
=====
Group Member      State           Mda  Addresses  Blocks    Se-% Hi Se-Prio
-----
1      1      active        2/1  175        23648    < 1  N  0
-----
No. of members: 1
=====
*A:PE1#
```

To verify relevant statistics for this NAT group/member, the following command can be used:


```
*A:PE1# show isa nat-group 1 member 1 statistics | match flow
no matching flow           : 56818
max flow exceeded         : 0
TCP no flow for RST       : 0
TCP no flow for FIN       : 0
TCP no flow                : 128094
flow log failed           : 0
new flow                  : 1470768
found flow                : 39661850
flow create logged        : 0
flow delete logged        : 0
flow log pkt tx           : 0
flow create failed, key ambiguous : 0
flow create failed, conflicting policies : 0
*A:PE1#
```

Case 2

Configured with:

- Mapping multiple prefixes from the same VRF into different outside pools.
- NAT all traffic.

Figure 59: Case 2



26153

The NAT **outside pool** are configured as follows:

```
configure
  service
    vprn 15002 customer 1 create
    nat
      outside
        pool "nat-pool-2" nat-group 1 type large-scale create
          port-reservation ports 80
          subscriber-limit 256
          deterministic
            port-reservation 180
          exit
          address-range 192.168.2.1 192.168.2.200 create
          exit
          no shutdown
        exit
        pool "nat-pool-3" nat-group 1 type large-scale create
          port-reservation ports 120
          port-forwarding-range 4023
          subscriber-limit 64
          deterministic
            port-reservation 840
          exit
          address-range 192.168.3.1 192.168.3.200 create
          exit
          no shutdown
        exit
      exit
    exit
  exit
exit
```

The NAT **policies** are configured as follows:

```
configure
  service
    nat
      nat-policy "nat-policy-2" create
        block-limit 4
        pool "nat-pool-2" router 15002
      exit
      nat-policy "nat-policy-3" create
        block-limit 2
        pool "nat-pool-3" router 15002
      exit
    exit
  exit
exit
```

The NAT **inside prefix** is configured as follows:

```
configure
  service
    vprn 15001 customer 1 create
    nat
      inside
        destination-prefix 0.0.0.0/0
        classic-lsn-max-subscriber-limit 256
        deterministic
          prefix 10.1.0.0/23 subscriber-type classic-lsn-sub
```



```
Type : classic-lsn-sub
Inside router : 15001
Inside IP address prefix : 10.1.1.0/32
ISA NAT group : 1
ISA NAT group member : 1
Outside router : 15002
Outside IP address : 192.168.2.2

-----
No. of LSN subscriber instances: 1
=====
*A:PE1#
```

The last subscriber mapping to the same 192.168.2.2 outside IP address has inside address 10.1.1.255.

To show the VPRN-15002 LSN blocks corresponding to the first inside IP address within 10.1.0.0/23 prefix, the following command can be used:

```
*A:PE1# show router 15002 nat lsn-blocks inside-ip 10.1.0.0

=====
Large-Scale NAT blocks for vprn15002
=====
192.168.2.1 [1024..1203]
Pool : nat-pool-2
Policy : nat-policy-2
Started : 2016/10/28 12:53:23
Inside router : vprn15001
Inside IP address : 10.1.0.0

-----
Number of blocks: 1
=====
*A:PE1#
```

To show the VPRN-15002 LSN blocks corresponding to the last inside IP address within 10.1.0.0/23 prefix, the following command can be used:

```
*A:PE1# show router 15002 nat lsn-blocks inside-ip 10.1.1.255

=====
Large-Scale NAT blocks for vprn15002
=====
192.168.2.2 [46924..47103]
Pool : nat-pool-2
Policy : nat-policy-2
Started : 2016/10/28 12:53:23
Inside router : vprn15001
Inside IP address : 10.1.1.255

-----
Number of blocks: 1
=====
*A:PE1#
```

The subscriber-limit corresponding to the 10.2.0.0/22 prefix is 64,so the 10.2.0.0/22 prefix is broken into sixteen /26 prefixes. Each of these /26 prefixes is mapped to a specific outside IP address.

To show the first LSN subscriber for the inside routing instance 15001 for the first /26 prefix, the following command can be used:

```
*A:PE1# show service nat lsn-subscribers inside-ip 10.2.0.0

=====
NAT LSN subscribers
=====
Subscriber          : [LSN-Host@10.2.0.0]
NAT policy          : nat-policy-3
Subscriber ID       : 276829984
-----
Type                : classic-lsn-sub
Inside router       : 15001
Inside IP address prefix : 10.2.0.0/32
ISA NAT group       : 1
ISA NAT group member : 1
Outside router      : 15002
Outside IP address   : 192.168.3.1
-----
No. of LSN subscriber instances: 1
=====
*A:PE1#
```

The last inside address mapping to the 192.168.3.1 outside address is 10.2.0.63.

To show the first LSN subscriber for the inside routing instance 15001 for the last /26 prefix, the following command can be used:

```
*A:PE1# show service nat lsn-subscribers inside-ip 10.2.3.192

=====
NAT LSN subscribers
=====
Subscriber          : [LSN-Host@10.2.3.192]
NAT policy          : nat-policy-3
Subscriber ID       : 276830944
-----
Type                : classic-lsn-sub
Inside router       : 15001
Inside IP address prefix : 10.2.3.192/32
ISA NAT group       : 1
ISA NAT group member : 1
Outside router      : 15002
Outside IP address   : 192.168.3.16
-----
No. of LSN subscriber instances: 1
=====
*A:PE1#
```

The last inside address mapping to the 192.168.3.16 outside address is 10.2.3.255.

To show the VPRN-15002 LSN blocks corresponding to the first inside IP address within the 10.2.0.0/22 prefix, the following command can be used:

```
*A:PE1# show router 15002 nat lsn-blocks inside-ip 10.2.0.0

=====
Large-Scale NAT blocks for vprn15002
=====
```

```
192.168.3.1 [4024..4863]
Pool                : nat-pool-3
Policy              : nat-policy-3
Started             : 2016/10/28 12:53:23
Inside router       : vprn15001
Inside IP address   : 10.2.0.0
```

```
-----
Number of blocks: 1
=====
```

```
*A:PE1#
```

To show the VPRN-15002 LSN blocks corresponding to the last inside IP within 10.2.0.0/22 prefix, the following command can be used:

```
*A:PE1# show router 15002 nat lsn-blocks inside-ip 10.2.3.255
```

```
=====
Large-Scale NAT blocks for vprn15002
=====
```

```
192.168.3.16 [56944..57783]
Pool                : nat-pool-3
Policy              : nat-policy-3
Started             : 2016/10/28 12:53:23
Inside router       : vprn15001
Inside IP address   : 10.2.3.255
```

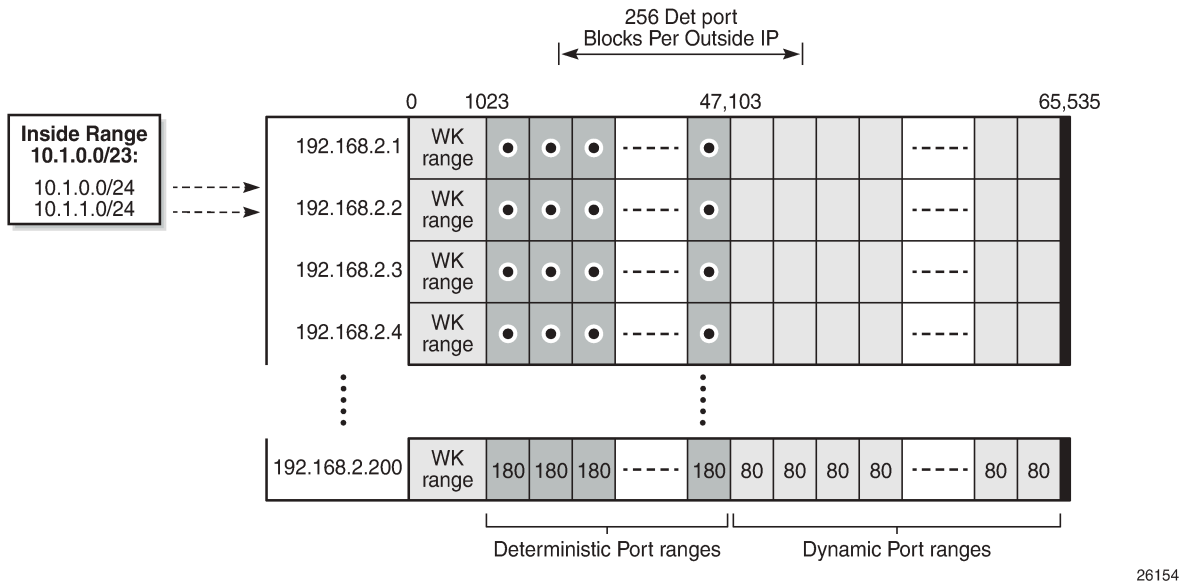
```
-----
Number of blocks: 1
=====
```

```
*A:PE1#
```

Mapping results


According to this configuration, for the 10.1.0.0/23 prefix, each inside IP address has one deterministic block of 180 ports and can have up to three dynamic blocks (block-limit =4) each of 80 ports, allowing for a maximum of $180+3*80 = 420$ flows.

Figure 60: Case 2: Prefix 10.1.0.0/23 results



According to this configuration, for the 10.2.0.0/22 prefix, each inside IP address has one deterministic block of 840 ports, and can have up to one dynamic block (block-limit =2) of 120 ports, allowing for a maximum of 840+120 = 960 flows.

Figure 61: Case 2: Prefix 10.2.0.0/22 results




Case 3

Configured with:

- Mapping overlapping prefixes from different VRFs into the same outside pool.
- NAT all traffic.

Figure 62: Case 3



26156

The NAT **outside pool** is configured as follows:

```
configure
  router "Base"
  nat
  outside
    pool "nat-pool-4" nat-group 1 type large-scale create
    port-reservation ports 461
    port-forwarding-range 4023
    subscriber-limit 64
    deterministic
    port-reservation 500
  exit
  address-range 192.168.4.1 192.168.4.100 create
  exit
  no shutdown
  exit
exit
exit
exit
exit
```

The NAT **policy** is configured as follows:

```
configure
  service
  nat
  nat-policy "nat-policy-4" create
  block-limit 4
  pool "nat-pool-4" router Base
  exit
exit
exit
exit
```


The NAT **inside prefix** is configured as follows:

```

configure
  service
    vprn 15001 customer 1 create
      nat
        inside
          destination-prefix 0.0.0.0/0
          classic-lsn-max-subscriber-limit 256
          deterministic
            prefix 10.5.0.0/20 subscriber-type classic-lsn-sub
              nat-policy "nat-policy-4" create
                map start 10.5.0.0 end 10.5.15.255 to 192.168.4.1
                no shutdown
            exit
          exit
        exit
      exit
    exit
  exit
configure
  service
    vprn 15002 customer 1 create
      nat
        inside
          destination-prefix 0.0.0.0/0
          classic-lsn-max-subscriber-limit 128
          deterministic
            prefix 10.5.0.0/27 subscriber-type classic-lsn-sub
              nat-policy "nat-policy-4" create
                map start 10.5.0.0 end 10.5.0.31 to 192.168.4.65
                no shutdown
            exit
          exit
        exit
      exit
    exit
  exit
exit

```

Show commands

For the 10.5.0.0/20 prefix on VPRN 15001, the subscriber-limit is 64. The 10.5.0.0/20 prefix will be broken into 64 smaller /26 prefixes, each will be mapped into a specific outside IP address.

To show the first LSN subscriber for the inside routing instance 15001 of the first /26 prefix, the following command can be used:

```

*A:PE1# show service nat lsn-subscribers inside-ip 10.5.0.0 inside-router 15001
=====
NAT LSN subscribers
=====
Subscriber          : [LSN-Host@10.5.0.0]
NAT policy          : nat-policy-4
Subscriber ID       : 276825344
-----
Type                : classic-lsn-sub

```

```

Inside router      : 15001
Inside IP address prefix : 10.5.0.0/32
ISA NAT group     : 1
ISA NAT group member : 1
Outside router    : "Base"
Outside IP address : 192.168.4.1
    
```

```

-----
No. of LSN subscriber instances: 1
=====
    
```

```
*A:PE1#
```

The last inside address mapping to the 192.168.4.1 outside address is 10.5.0.63.

To show the first Large Scale NAT (LSN) subscriber for the inside routing instance 15001 of the last /26 prefix, the following command can be used:

```
*A:PE1# show service nat lsn-subscribers inside-ip 10.5.15.192 inside-router 15001
```

```

=====
NAT LSN subscribers
=====
    
```

```

Subscriber      : [LSN-Host@10.5.15.192]
NAT policy      : nat-policy-4
Subscriber ID   : 276829376
    
```

```

-----
Type            : classic-lsn-sub
Inside router   : 15001
Inside IP address prefix : 10.5.15.192/32
ISA NAT group   : 1
ISA NAT group member : 1
Outside router  : "Base"
Outside IP address : 192.168.4.64
    
```

```

-----
No. of LSN subscriber instances: 1
=====
    
```

```
*A:PE1#
```

The last inside address mapping to the 192.168.4.64 outside address is 10.5.15.255.

To show the base router LSN blocks corresponding to the first inside IP address within the 10.5.0.0/20 prefix, the following command can be used:

```
*A:PE1# show router nat lsn-blocks inside-ip 10.5.0.0 inside-router 15001
```

```

=====
Large-Scale NAT blocks for Base
=====
    
```

```

192.168.4.1 [4024..4523]
Pool            : nat-pool-4
Policy         : nat-policy-4
Started        : 2016/10/27 13:11:38
Inside router  : vprn15001
Inside IP address : 10.5.0.0
    
```

```

-----
Number of blocks: 1
=====
    
```

```
*A:PE1#
```

To show the base router LSN blocks corresponding to the last inside IP address within the 10.5.0.0/20 prefix, the following command can be used:

```
*A:PE1# show router nat lsn-blocks inside-ip 10.5.15.255 inside-router 15001

=====
Large-Scale NAT blocks for Base
=====
192.168.4.64 [35524..36023]
Pool                : nat-pool-4
Policy              : nat-policy-4
Started             : 2016/10/27 13:11:38
Inside router       : vprn15001
Inside IP address   : 10.5.15.255

-----
Number of blocks: 1
=====
*A:PE1#
```

For the 10.5.0.0/27 prefix in VPRN 15002, the subscriber-limit is 64. The 10.5.0.0/27 prefix will be mapped into one outside IP address.

To show the first LSN subscriber for the inside routing instance 15002 of the 10.5.0.0/27 prefix, the following command can be used:

```
*A:PE1# show service nat lsn-subscribers inside-ip 10.5.0.0 inside-router 15002

=====
NAT LSN subscribers
=====
Subscriber          : [LSN-Host@10.5.0.0]
NAT policy          : nat-policy-4
Subscriber ID       : 276829440

-----
Type                : classic-lsn-sub
Inside router       : 15002
Inside IP address prefix : 10.5.0.0/32
ISA NAT group       : 1
ISA NAT group member : 1
Outside router      : "Base"
Outside IP address   : 192.168.4.65

-----
No. of LSN subscriber instances: 1
=====
*A:PE1#
```

To show the LSN blocks corresponding to the first inside IP address within the 10.5.0.0/27 prefix, the following command can be used:

```
*A:PE1# show router nat lsn-blocks inside-ip 10.5.0.0 inside-router 15002

=====
Large-Scale NAT blocks for Base
=====
192.168.4.65 [4024..4523]
Pool                : nat-pool-4
Policy              : nat-policy-4
Started             : 2016/10/27 13:12:02
Inside router       : vprn15002
```

```

Inside IP address           : 10.5.0.0
-----
Number of blocks: 1
=====
*A:PE1#
    
```

To show the LSN blocks for the last inside IP address within the 10.5.0.0/27 prefix, the following command can be used:

```

*A:PE1# show router nat lsn-blocks inside-ip 10.5.0.31 inside-router 15002

=====
Large-Scale NAT blocks for Base
=====
192.168.4.65 [19524..20023]
Pool           : nat-pool-4
Policy        : nat-policy-4
Started       : 2016/10/27 13:12:02
Inside router  : vprn15002
Inside IP address : 10.5.0.31
-----
Number of blocks: 1
=====
*A:PE1#
    
```


Mapping results

According to this configuration, each inside IP address within VPRN 15001 has one deterministic block of 500 ports and can have up to three dynamic blocks (block-limit=4) of 461 ports each, allowing a maximum of $500+3*461 = 1883$ flows.

According to this configuration each inside IP address within VPRN 15002 has one deterministic block of 500 ports and can have up to three dynamic blocks (block-limit=4) of 461 ports each, allowing a maximum of $500+3*461 = 1883$ flows.

For VPRN 15002, since the number of LSN subscribers (32) is less than the number of deterministic blocks (64), then 32 deterministic blocks will be wasted, specifically $32*500 = 16,000$ ports will be wasted which is not good in terms of capacity planning.

Figure 63: Case 3 results




Inverse mapping

In deterministic LSN44, the inside IP addresses are mapped to outside IP addresses and corresponding port blocks based on a deterministic algorithm. The inverse mapping that reveals the subscriber identity behind the NAT is based on the reversal of this algorithm.

Inverse mappings can be done either online or offline:

- Online — Locally on the 7x50 node, via CLI (MIB)
- Offline — Externally, via a Python script. The purpose of such an offline approach is to provide fast queries without accessing the 7x50.

Figure 64: Inverse mapping approach



Online approach

A **tools** command is available which shows the reverse mapping (outside to inside) for deterministic NAT instead of using logging.

```
tools dump nat deterministic-mapping outside-ip <ipv4-address> router <router-instance>
outside-port <[1..65535]>

<ipv4-address>      : a.b.c.d
<router-instance>  : <router-name>|<service-id>
                    router-name   - "Base"
                    service-id    - [1..2147483647]
```

Using Case 3 as an example, to obtain (inside IP, inside routing instance), the inverse mapping for a specific (outside IP, outside routing instance, outside port) is done as follows:

```
*A:PE1# tools dump nat deterministic-mapping outside-ip 192.168.4.1 router "Base" outside-port
4024
classic-lsn-sub inside router 15001 ip 10.5.0.0 -- outside router Base ip 192.168.4.1 port 4024
at Fri Oct 28 13:04:22 UTC 2016
*A:PE1#
```

```
*A:PE1# tools dump nat deterministic-mapping outside-ip 192.168.4.65 router "Base" outside-port
4024
classic-lsn-sub inside router 15002 ip 10.5.0.0 -- outside router Base ip 192.168.4.65 port
4024 at Fri Oct 28 13:04:45 UTC 2016
*A:PE1#
```

Offline approach

The purpose of such an offline approach is to provide fast queries without the need to directly query the 7x50.

This is achieved by generating and exporting a Python script for reverse querying, which is a manual operation that needs to be repeated every time there is configuration change in deterministic NAT.

The script is exported (manually) to the external system.

To configure remotely the location for the Python script, the following command is used:

```
configure service nat deterministic-script location <remote-url>
```

remote-url — A remote location where the script is stored:

```
[[ftp://|tftp://]<login>:<pswd>@ <remote-locn>][<file-path>]
```

Maximum length is 180 characters.

Once the script location is specified, the script can be exported to that location using the following command:

```
admin nat save-deterministic-script
```

Using the following command the status of the script can be checked, and whether it is necessary to re-save (export) the script or not:

```
*A:PE1# show service nat deterministic-script

=====
Deterministic NAT script data
=====
Location           : ftp://*:*@123.123.123.123/pub/python/detnat.py
Save needed        : no
Last save result    : success
Last save time     : 2016/10/28 13:05:41
=====
*A:PE1#
```

The external system must have the Python scripting language installed with the following modules: getopt, math, os, socket, and sys.

The Python script can then be run on the external server; the parameters are as follows:

```
[user@123.123.123.123 ~]$ ./detnat.py
Error: need exactly one of --forward or --backward arguments

Usage: detnat.py DIRECTION PARAMETERS
Perform forward or backward NAT according to the configured deterministic rules.

DIRECTION:
  -f, --forward          Translate from inside to outside address/port
  -b, --backward        Translate from outside to inside address/port

PARAMETERS:
  -a, --address=IP-ADDRESS The address to translate. IPv6 addresses can be
                           specified in shorthand or full notation.
  -p, --port=PORT          The outside port in case of backward translation.
  -s, --service=SERVICE-ID This is the inside service in case of forward
                           translation and the outside service in case of
                           backward translation.
                           To specify the base router, this option must be
                           omitted.

  -h, --help              Show this help message
[user@123.123.123.123 ~]$
```

where deterministic-nat.py is the name of the python script previously exported.

As an example of a forward query:

```
[user@123.123.123.123 ~]$ ./detnat.py -f -s 15001 -a 10.0.0.1
classic-lsn-sub has public ip address 192.168.0.1 from base router and is using ports [4324 -
4623]
[user@123.123.123.123 ~]$
```

As an example of a reverse query:

```
[user@123.123.123.123 ~]$ ./detnat.y -b -s 0 -a 192.168.0.1 -p 4325
classic-lsn-sub has private ip address 10.0.0.1 from service 15001
[user@123.123.123.123 ~]$
```

Simultaneous support of deterministic and non-deterministic NAT

Deterministic NAT can be used simultaneously with non-deterministic NAT within the same inside routing instance. However, they cannot share the same pool.

An outside pool can be only deterministic (although expandable by dynamic ports blocks) or non-deterministic at any given time (a non-deterministic pool is a pool that contains dynamic port-blocks only).

The following show a configuration using deterministic NAT simultaneously with non-deterministic NAT.

The NAT **outside pool** are configured as follows:

```
configure
  router
    nat
      outside
        pool "nat-pool-1" nat-group 1 type large-scale create
          port-reservation ports 180
          port-forwarding-range 4023
          subscriber-limit 128
          deterministic
            port-reservation 300
          exit
          address-range 192.168.0.1 192.168.0.100 create
          exit
          no shutdown
        exit
        pool "nat-pool-Non-Deterministic" nat-group 1 type large-scale create
          address-range 192.168.7.1 192.168.7.100 create
          exit
          no shutdown
        exit
      exit
    exit
  exit
exit
```

The NAT **policies** are configured as follows:

```
configure
  service
    nat
      nat-policy "nat-policy-1" create
        block-limit 4
        pool "nat-pool-1" router Base
      exit
      nat-policy "nat-policy-Non-Deterministic" create
        pool "nat-pool-Non-Deterministic" router Base
      exit
    exit
  exit
exit
```

The NAT **inside prefixes** are configured as follows:

```
configure
  service
    vprn 15001 customer 1 create
      nat
        inside
          destination-prefix 0.0.0.0/0
          classic-lsn-max-subscriber-limit 256
```




```

deterministic
  prefix 10.0.0.0/24 subscriber-type classic-lsn-sub
  nat-policy "nat-policy-1" create
  map start 10.0.0.0 end 10.0.0.255 to 192.168.0.1
  no shutdown
  exit
exit
nat-policy "nat-policy-Non-Deterministic"
exit
exit
no shutdown
exit
exit
exit
exit

```

In this example, the inside IP prefixes that do not match any of the deterministic prefixes will be NATed using a non-deterministic pool.

Figure 65: Sending flows: deterministic + non-deterministic NAT



26159

To check which NAT pool/NAT policy is used for NATing the inside IP 10.7.0.1, the following command can be used:

```

*A:PE1# show router nat lsn-blocks inside-ip 10.7.0.1

=====
Large-Scale NAT blocks for Base
=====
192.168.7.100 [1024..1527]
Pool                : nat-pool-Non-Deterministic
Policy              : nat-policy-Non-Deterministic
Started             : 2016/10/28 13:24:56
Inside router       : vprn15001
Inside IP address   : 10.7.0.1

-----
Number of blocks: 1
=====
*A:PE1#

```

Conclusion

This example provides the commands required for configuring deterministic LSN44 NAT. Both deterministic as well as non-deterministic NAT are supported, with simultaneous operation being possible.

Inverse query can be done online or offline to retrieve the NAT mappings. Logging is not needed as long as there are no dynamic blocks assigned to LSN44 subscribers.

IP/GRE Termination

This chapter provides configuration and troubleshooting commands for IP/GRE termination.

Topics in this chapter include:

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

Applicability

The chapter was initially written for SR OS Release 9.0.R8. The CLI in the current edition corresponds to SR OS Release 22.2.R2.

Initially, the 7750 SR only supported GRE SDP tunnels which use pseudowire encapsulation. In SR OS Release 8.0.R5 and later, the 7750 SR supports tunneling IPv4 packets in an IPv4 Generic Routing Encapsulation (GRE) tunnel.


The IP GRE tunnel termination configuration described in this chapter requires an MS-ISA. IP GRE tunnels without ISA are beyond the scope of this chapter.

Overview

A common use case for IP/GRE tunneling is remote access to a VPRN over a public IP network because IP/GRE tunneling allows encapsulated packets to follow a path based on the outer IP header which is useful when the inner IP packet cannot or should not be forwarded natively over this path.

GRE allows packets of one protocol, the payload protocol, to be encapsulated by packets of another protocol, called the delivery protocol. [Figure 66: GRE packet format](#) shows the GRE packet format with an outer delivery header, GRE header, and payload packet:

Figure 66: GRE packet format



al_0132

The outer delivery and GRE header for outgoing traffic is as follows.

- Outer delivery header
 - The source address in the IPv4 delivery header is the configured source address.

- The destination address in the IPv4 delivery header is the configured remote IP (or the backup remote IP) address.
- The IP protocol value in the IPv4 delivery header is 0x2F or 47 (GRE).
- The DSCP in the IPv4 outer delivery header is:
 - set to the value configured for the tunnel;
 - otherwise, the DSCP value from the payload packet is copied into the outer delivery header.
- The TTL in the IPv4 outer delivery header is set to 255.
- GRE header
 - The checksum (C) bit in the GRE header is set to 0 (no checksum present).
 - The version in the GRE header is 0.
 - The protocol type in the GRE header is 0x0800 for IPv4.

The outer delivery and GRE header for incoming traffic is as follows:

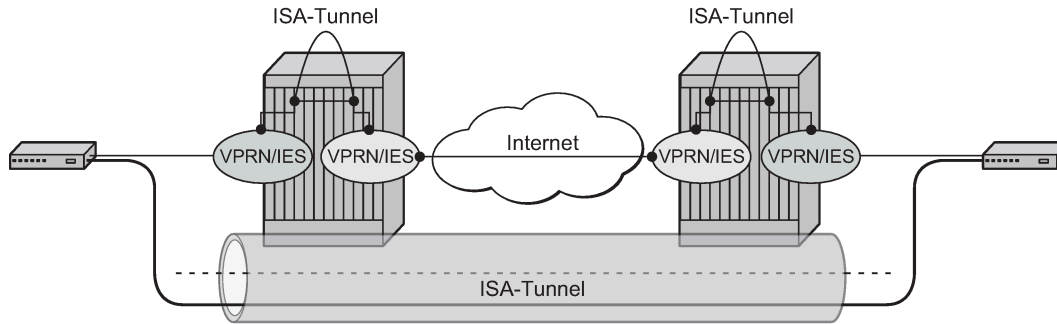
- Outer delivery header
 - If the packet is a fragment (more fragments=1, non-zero fragment offset), it is dropped.
 - If the checksum (C) bit in the GRE header is set, then the included checksum is validated; if the checksum is incorrect, the packet is discarded.
 - If the version in the GRE header is not 0, the packet is discarded.
 - If the source/destination address pair in the IPv4 delivery header is any other combination, the packet is dropped.
- GRE header
 - If the checksum (C) bit in the GRE header is set, then the included checksum is validated; if the checksum is incorrect, the packet is discarded.
 - If the version in the GRE header is not 0, the packet is discarded.

Implementation

Encapsulation, de-encapsulation and other datapath operations related to IP/GRE are handled by the ISA-tunnel MDA.

For GRE tunnels configured as SDPs (which are not covered by this section), no ISA-tunnel MDA is required.


Figure 67: Implementation



al_0133

SR OS nodes initially supported the IP/GRE tunnels with static routes and BGP. IP/GRE tunnels have been enhanced by adding support for OSPF and BFD on private tunnel interfaces (used with static routes, OSPF, or BGP) and GRE protection by tunneling into an IPsec tunnel.

Figure 68: IP/GRE over IPsec tunnel



al_0134

Configuration

ISA-tunnel MDA

The ISA-tunnel MDA supports IP/GRE and IPsec tunnels and is configured as follows:

```
# on PE-1:
configure
  card 1
    mda 2
      mda-type isa2-tunnel
      no shutdown
    exit
  exit
  card 2
    mda 2
      mda-type isa2-tunnel
      no shutdown
    exit
  exit
```

The following command checks the MDA configuration:

```
*A:PE-1# show mda

=====
MDA Summary
=====
```

| Slot | Mda | Provisioned Type Equipped Type (if different) | Admin State | Operational State |
|------|-----|--|----------------|----------------------|
| 1 | 1 | p10-10g-sfp | up | up |
| | 2 | isa2-tunnel p-isa2-ms | up | up |
| 2 | 1 | me40-1gb-csfp | up | up |
| | 2 | isa2-tunnel me-isa2-ms | up | up |

```
=====
```

Tunnel groups and tunnel group restrictions

The first step of the GRE tunnel configuration is to configure a tunnel group.

A tunnel group can have one tunnel ISA designated primary and optionally one tunnel-ISA designated backup. When a GRE tunnel is created, it is assigned to the primary tunnel-ISA in its tunnel group. If the primary tunnel-ISA fails, the backup tunnel-ISA (if not already claimed by another tunnel group) takes over for the failed card.

```
*A:PE-1>config>isa# tunnel-group 1 ?
- tunnel-group <tunnel-group-id> [create]
- tunnel-group <tunnel-group-id> isa-scale-mode <isa-scale-mode> [create]
- no tunnel-group <tunnel-group-id>

<tunnel-group-id>      : [1..16]
<isa-scale-mode>      : tunnel-limit-2k
                       k=1024
<create>              : keyword - mandatory while creating an entry.

[no] active-mda-num* - Configure number of active MDAs
[no] backup          - Configure ISA-Tunnel-Group backup ISA
[no] description    - Configure the ISA group description
[no] esa-vm         - Configure the esa-vm
[no] ipsec-responde* - Enable/Disable responder-only for IPsec Ikev2 tunnels only
[no] mda            - Configure MDA to associate with
[no] multi-active   - Configure multi-active status of tunnel-group
[no] primary        - Configure ISA-Tunnel-Group primary ISA
[no] reassembly     - Configure reassembly wait time
[no] shutdown       - Administratively enable/disable an ISA-Tunnel-Group
stats-collecti* +  Configure ISA statistics collection parameters

# on PE-1:
configure
  isa
    tunnel-group 1 create
    primary 1/2
    backup 2/2
    no shutdown
  exit
```

The failed tunnels are re-established using a cold-standby on the backup tunnel-ISA. Cold-standby means the backup tunnel-ISA has no state or configuration information about the tunnels prior to the failure.

A tunnel ISA cannot be primary for more than one tunnel group:

```
*A:PE-1>config>isa# tunnel-group 2 create
*A:PE-1>config>isa>tunnel-grp$ primary 1/2
MINOR: IPSECGRPMGR #1003 The specified MDA is primary in another Tunnel Group
```

A tunnel ISA cannot be primary in one tunnel group and backup in another tunnel group:

```
*A:PE-1>config>isa# tunnel-group 2 create
*A:PE-1>config>isa>tunnel-grp# backup 1/2
MINOR: IPSECGRPMGR #1003 The specified MDA is primary in another Tunnel Group
```

The following commands shows the ISA tunnel group (after tunnel group 2 has been removed):

```
*A:PE-1# show isa tunnel-group

=====
ISA Tunnel Groups
=====
Tunnel      PrimaryIsa      BackupIsa      ActiveIsa      Admin      Oper
GroupId
-----
1           1/2              2/2            1/2            Up         Up
-----
No. of ISA Tunnel Groups: 1
=====
```

The following command shows the number of the IP (GRE) tunnels, after configuring IES and VPRN services with tunnel interfaces:

```
*A:PE-1# show ip tunnel count
-----
IP Tunnels: 2
-----
```

The following command shows all IP tunnels:

```
*A:PE-1# show ip tunnel

=====
IP Tunnels
=====
TunnelName      SapId      SvcId      Admn
Local Address   DlvrySvcId Oper
OperRemoteAddress
-----
gre-tunnel-1    tunnel-1.private:1    1          Up
192.168.1.1     2            Up
192.168.2.1
protected-gre-tunnel    tunnel-1.private:5    3          Up
192.168.11.1    3            Up
192.168.22.1
-----
IP Tunnels: 2
=====
```

The detailed tunnel information is as follows:

```
*A:PE-1# show ip tunnel "gre-tunnel-1"

=====
IP Tunnel Configuration Detail
=====
Service Id       : 1                Sap Id          : tunnel-1.private:1
Tunnel Name      : gre-tunnel-1
Description      : None
GRE Header       : Yes
Delivery Service : 2
GRE Keys Set     : False
GRE Send Key     : N/A              GRE Receive Key : N/A
Admin State      : Up                Oper State      : Up
Source Address   : 192.168.1.1
Remote Address   : 192.168.2.1
Backup Address   : (Not Specified)
Oper Remote Addr : 192.168.2.1
DSCP             : None
Reassembly       : inherit
Clear DF Bit     : false            IP MTU          : max
Encap IP MTU     : max
Pkt Too Big     : true
Pkt Too Big Num : 100              Pkt Too Big Intvl: 10 secs
Frag Required    : true
Frag Req Count  : 100              Frag Req Interval: 10 secs
Propagate IPv6 P*: true
Propagate IPv4 P*: true
Oper Flags       : None
Transport Profile: (Not Specified)
Last Oper Changed: 05/12/2022 08:40:02
Host ISA        : 1/2
TCP MSS Adjust
  Public         : Disabled
  Private        : Disabled

-----
Target Address Table
-----
Destination IP          IP Resolved Status
-----
10.0.0.2                Yes
-----

=====
IP Tunnel Statistics: gre-tunnel-1
=====
Errors Rx       : 0                Errors Tx       : 0
Pkts Rx        : 51                Pkts Tx        : 49
Bytes Rx       : 3575              Bytes Tx       : 3483
Key Ignored Rx : 0                Too Big Tx     : 0
Seq Ignored Rx : 0
Vers Unsup. Rx : 0
Invalid Chksum Rx: 0
Key Mismatch Rx : 0

=====

=====
Fragmentation Statistics
=====
Encapsulation Overhead : 24
Temporary Private MTU  : max
Pre-Encapsulation
```

```

Fragmentation Count      : 0
Last Fragmented Packet Size : 0
Post-Encapsulation
Fragmentation Count      : 0
Last Fragmented Packet Size : 0
=====
* indicates that the corresponding row element may have been truncated.

```

Interfaces

The interface toward the Internet (or WAN):

- can be a network interface or VPRN/IES interface.
- provides IP reachability.


The tunnel public interface:

- can be an IES or VPRN interface.
- represents the public side of the tunnel-ISA.

The delivery VPRN/IES service (the service connected to the Internet) must have at least one IP interface associated with a public tunnel SAP in order to receive and process GRE encapsulated packets.

The public tunnel SAP type has the format **tunnel-id.private|public:tag** (where the *id* corresponds to the tunnel group). [Figure 70: GRE for remote access to a VPRN service](#) shows the example topology, where CE-2 in customer site A is connected to PE-1.

Figure 69: GRE for remote access to a VPRN service



The IES service with public tunnel SAP is configured on PE-1 as follows:

```

*A:PE-1>config>service>ies>if# sap ?
- no sap <sap-id>
- sap <sap-id> [create]

<sap-id>
---snip---
      tunnel-id      - tunnel-<id>.<private|public>:<tag>
      tunnel         - keyword
      id              - [1..16]
      tag             - [0..4094]
---snip---

# on PE-1:
configure
service
  ies 2 name "IES 2" customer 1 create
  interface "int-tunnel-public" create
  address 192.168.1.2/30

```



```

        tos-marking-state untrusted
        sap tunnel-1.public:1 create
        exit
    exit
    interface "int-PE-1-CE-2" create
        address 192.168.12.1/24
        sap 1/1/2:2 create
        exit
    exit
    no shutdown
exit

```

PE-1 has address 192.168.1.2/30 assigned to the interface "int-tunnel-public" in IES 2. In a similar way, CE-2 has address 192.168.2.2/30 assigned to the interface "int-tunnel-public" in IES 2.

In order to reach 192.168.2.0/30 on CE-2, a static route is configured on PE-1, as follows:

```

# on PE-1:
configure
    router Base
        static-route-entry 192.168.2.0/30
            next-hop 192.168.12.2
            no shutdown

```

In a similar way, a static route is configured on CE-2 to reach 192.168.1.0/30 on PE-1.

Mask /32 is not supported on the public tunnel. When address 192.168.1.2/32 is configured on the interface "int-tunnel-public", the public tunnel cannot be created, as follows:

```

*A:PE-1>config>service>ies>if# address 192.168.1.2/32
*A:PE-1>config>service>ies>if# sap tunnel-1.public:1 create
INFO: PIP #1288 Cannot bind when there are /32 or /128 addresses configured

```

Therefore, the address configured on the interface will have mask /30 instead of /32, as shown earlier.

The tunnel private interface:

- can be an IES or VPRN interface.
- represents the private side of the tunnel ISA.

The private tunnel SAP has the format **tunnel-*id*.private|public:tag** (where the *id* corresponds to the tunnel group) as shown in the following example where an unprotected GRE tunnel is configured in the SAP context.

```

*A:PE-1>config>service>vprn>if# sap ?
- no sap <sap-id>
- sap <sap-id> [create]

<sap-id>
    ---snip---
        tunnel-id      - tunnel-<id>.<private|public>:<tag>
        tunnel         - keyword
        id             - [1..16]
        tag            - [0..4094]
    ---snip---

# on PE-1:
configure
    service
        vprn 1 name "VPRN 1" customer 1 create
        interface "int-gre-tunnel" tunnel create

```

```

address 10.0.0.1/30
sap tunnel-1.private:1 create
ip-tunnel "gre-tunnel-1" create
  dest-ip 10.0.0.2
  gre-header
---snip---

```

It is not mandatory to have the same tag (internal dot1q) in private and public GRE tunnels.

```
sap tunnel-1.private:1 <=> sap tunnel-1.public:2
```

Unprotected GRE tunnel configuration

To associate an unprotected GRE tunnel with a private tunnel SAP, the **ip-tunnel** command is configured in the SAP context.

```

# on PE-1:
configure
  service
    vprn 1 name "VPRN 1" customer 1 create
    ---snip---
    interface "int-gre-tunnel" tunnel create
      address 10.0.0.1/30
      sap tunnel-1.private:1 create
      ip-tunnel "gre-tunnel-1" create
        dest-ip 10.0.0.2
        gre-header
      ---snip---

```

The **dest-ip** keyword followed by the private IP address of the remote tunnel endpoint is mandatory.

If this remote IP address is not within the subnet of the local private endpoint, then the tunnel will not come up.

The following parameters are configured in the **ip-tunnel** context:

- The source address of the GRE tunnel. This is the source IPv4 address of GRE encapsulated packets sent by the delivery service. It must be an address in the subnet of the associated public tunnel SAP interface.
- The remote IP address. If this address is reachable in the delivery service (there is a route), then this is the destination IPv4 address of GRE encapsulated packets sent by the delivery service.
- The backup remote IP address. If the remote IP address of the tunnel is not reachable, then the backup remote IP address is the destination IPv4 address of GRE encapsulated packets sent by the delivery service.
- The delivery service. This is the identifier or name of the IES or VPRN service where GRE encapsulated packets are injected and terminated. The delivery service can be the same service where the private tunnel SAP interface resides.
- The DSCP marking in the outer IP header of GRE encapsulated packets. If this is not configured, then the default copies the DSCP from the inner IP header to the outer IP header.

```

# on PE-1:
configure
  service
    vprn 1 name "VPRN 1" customer 1 create
    interface "int-gre-tunnel" tunnel create

```

```

address 10.0.0.1/30
sap tunnel-1.private:1 create
  ip-tunnel "gre-tunnel-1" create
    dest-ip 10.0.0.2
    gre-header
    source 192.168.1.1
    remote-ip 192.168.2.1
    delivery-service 2
    dscp af22
    no shutdown
  exit
  ---snip---
  
```

- A private tunnel SAP can have only one IP/GRE tunnel (per SAP).


```

*A:PE-1>config>service>vprn>if# sap tunnel-1.private:1 ip-tunnel "gre-tunnel-2" create
MINOR: SVCMGR #5120 Only one IP tunnel allowed per SAP
  
```

IP/GRE tunneling via static route


A static route can reference the GRE tunnel directly (by next-hop IP address) or the GRE tunnel can be the resolved next-hop for an indirect static route (Figure 70: GRE for remote access to a VPRN service).

Figure 70: GRE for remote access to a VPRN service



The details of both ends on the GRE tunnel, at site A and PE-1, are shown in Figure 71: IP/GRE tunneling via static route. The node at left hand side is CE-2 at site A.

Figure 71: IP/GRE tunneling via static route



The following shows the configuration of VPRN 1 on PE-1.

```

# on PE-1:
configure
  
```

```

service
  vprn 1 name "VPRN 1" customer 1 create
  interface "int-gre-tunnel" tunnel create
  address 10.0.0.1/30
  sap tunnel-1.private:1 create
  ip-tunnel "gre-tunnel-1" create
  dest-ip 10.0.0.2
  gre-header
  source 192.168.1.1
  remote-ip 192.168.2.1
  delivery-service 2
  no shutdown
  exit
  exit
  interface "loopback1" create
  address 172.16.1.1/32
  loopback
  exit
  static-route-entry 172.16.2.1/32
  next-hop 10.0.0.2
  no shutdown
  exit
  exit
  bgp-ipvpn
  mpls
  route-distinguisher 64496:1
  vrf-target target:64496:1
  no shutdown
  exit
  exit
  no shutdown
  exit

```

The configuration of the VPRN on CE-2 is similar.

To check the static route status:

```
*A:PE-1# show router 1 static-route
```

```

=====
Static Route Table (Service: 1) Family: IPv4
=====
Prefix          Tag      Met   Pref Type Act
  Next Hop      Interface
-----
172.16.2.1/32   0         1     5   NH   Y
  10.0.0.2      int-gre-tunnel
-----
No. of Static Routes: 1
=====

```

IP/GRE tunneling via BGP peering

In this section, the configuration has BGP running inside the GRE tunnel.

```

# on PE-1:
configure
  service
    vprn 1 name "VPRN 1" customer 1 create

```

```

autonomous-system 64496
interface "int-gre-tunnel" tunnel create
  address 10.0.0.1/30
  sap tunnel-1.private:1 create
  ip-tunnel "gre-tunnel-1" create
    dest-ip 10.0.0.2
    gre-header
    source 192.168.1.1
    remote-ip 192.168.2.1
    delivery-service 2
    no shutdown
  exit
exit
interface "loopback1" create
  address 172.16.1.1/32
  loopback
exit
interface "loopback2" create
  address 172.31.1.1/24
  loopback
exit
static-route-entry 172.16.2.1/32
  next-hop 10.0.0.2
  no shutdown
exit
exit
bgp-ipvpn
  mpls
    route-distinguisher 64496:1
    vrf-target target:64496:1
    no shutdown
  exit
exit
bgp
  group "group-1"
    type internal
    export "export-bgp-172.31"
    local-address 172.16.1.1
    neighbor 172.16.2.1
  exit
  exit
  no shutdown
exit
no shutdown

```

It is mandatory to configure the autonomous system in the **vprn** context, otherwise the BGP session will not be established.

The configuration of the VPRN on CE-2 is similar.

The following command on PE-1 shows the summary of the BGP sessions. The BGP session between peers 172.16.1.1 in VPRN 1 on PE-1 and 172.16.2.1 in VPRN 1 on CE-2 is established for address family IPv4.

```
*A:PE-1# show router 1 bgp summary all
```

```
=====
BGP Summary
=====
```

```
Legend : D - Dynamic Neighbor
=====
```

```
Neighbor
```

| Description | ServiceId | AS | PktRcvd | InQ | Up/Down | State | Rcv/Act/Sent (Addr Family) |
|-------------|-----------|-------|---------|-----|-----------|-------|----------------------------|
| ----- | | | | | | | |
| 172.16.2.1 | 1 | 64496 | 7 | 0 | 00h01m12s | 1/1/1 | (IPv4) |
| | | | 7 | 0 | | | |
| ----- | | | | | | | |

In this example, PE-1 exports BGP route 172.31.1.0/24 and CE-2 exports BGP route 172.31.2.0/24. The route table for VPRN 1 on PE-1 includes the following BGP route:

```
*A:PE-1# show router 1 route-table protocol bgp

=====
Route Table (Service: 1)
=====
Dest Prefix[Flags]
Next Hop[Interface Name]
Type Proto Age Pref
Metric
-----
172.31.2.0/24 Remote BGP 00h00m45s 170
10.0.0.2 1
-----
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
=====
```

IP/GRE tunneling via OSPFv2 peering

OSPF can be run on IES and VPRN IP interfaces associated with private IP/GRE tunnel SAPs.

All OSPF features are supported, including area 0 and non-area 0 support, virtual links, authentication, BFD, configurable protocol timers.

```
# on PE-1:
configure
service
  vprn 1 name "VPRN 1" customer 1 create
  interface "int-gre-tunnel" tunnel create
  address 10.0.0.1/30
  sap tunnel-1.private:1 create
  ip-tunnel "gre-tunnel-1" create
  dest-ip 10.0.0.2
  gre-header
  source 192.168.1.1
  remote-ip 192.168.2.1
  delivery-service 2
  no shutdown
  exit
  exit
  interface "loopback1" create
  address 172.16.1.1/32
  loopback
  exit
  bgp-ipvpn
  mpls
```

```

        route-distinguisher 64496:1
        vrf-target target:64496:1
        no shutdown
    exit
exit
ospf
    area 0.0.0.0
        interface "int-gre-tunnel"
        exit
        interface "loopback1"
        exit
    exit
    no shutdown
exit
no shutdown
exit

```

The configuration on CE-2 is similar.

The following command shows the OSPF neighbors for VPRN 1:

```

*A:PE-1# show router 1 ospf neighbor
=====
Rtr vprn1 OSPFv2 Instance 0 Neighbors
=====
Interface-Name          Rtr Id      State      Pri  RetxQ  TTL
  Area-Id
-----
int-gre-tunnel          192.0.2.2   Two Way    1    0      32
  0.0.0.0
-----
No. of Neighbors: 1
=====

```

The OSPF routes in the routing table of VPRN 1 are as follows:

```

*A:PE-1# show router 1 route-table protocol ospf
=====
Route Table (Service: 1)
=====
Dest Prefix[Flags]      Type  Proto  Age      Pref
  Next Hop[Interface Name]      Metric
-----
172.16.2.1/32          Remote OSPF  00h00m22s  10
  10.0.0.2                2
-----
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
=====

```

IP/GRE tunneling protection using IPsec tunnel mode

To provide protection against potential threats such as spoofing, the GRE packets can be encrypted and authenticated using IPsec.

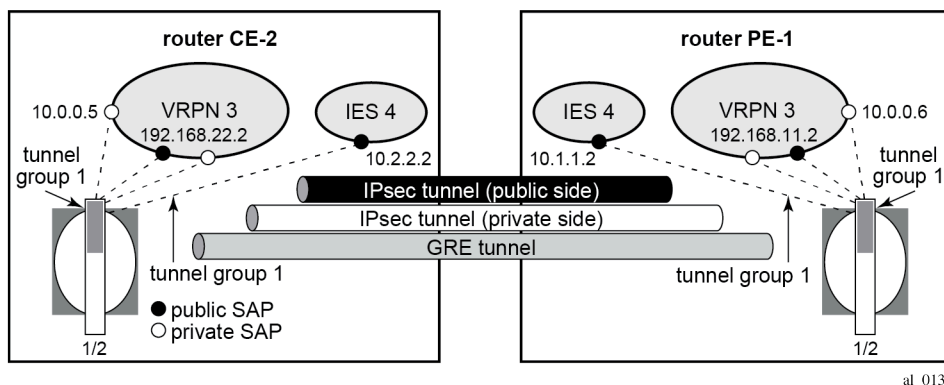
GRE packets receive IPsec protection by forwarding them, after encapsulation by a tunnel-ISA, into an IPsec tunnel supported by another (or the same) tunnel ISA.

Note that when configuring GRE protection by an IPsec tunnel:

- A GRE tunnel and its protecting IPsec tunnel may belong to the same or different tunnel groups (the same tunnel group is assumed in the following example).
- A GRE tunnel and its protecting IPsec tunnel may be assigned to the same tunnel ISA (if they belong to the same tunnel group) or different tunnel ISAs.
- A single IPsec tunnel can protect one or more GRE tunnels in addition to other IP traffic that meets the IPsec security policy.
- The private IPsec tunnel SAP interface and public GRE tunnel SAP interface are always part of the same VPRN. The private GRE tunnel SAP interface can be part of this same VPRN or a different VPRN.

In the following example, the GRE tunnel and its protecting IPsec tunnel belong to the same tunnel group.

Figure 72: Example GRE over IPsec tunnel



IPsec configuration

An **ike-policy** and **ipsec-transform** must be configured on PE-1 and CE-2, as follows:

```
# on PE-1, CE-2:
configure
  ipsec
    ike-transform 1 create
      dh-group 5
    exit
    ike-policy 1 create
      ike-transform 1
    exit
    ipsec-transform 1 create
      esp-encryption-algorithm aes256
    exit
```

The public/private side of the GRE tunnel and the private side of the IPsec tunnel are in the same VPRN, as shown in the following configuration example:

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



```

service
  vprn 3 name "VPRN 3" customer 1 create
  ipsec
    security-policy 1 create
      entry 1 create
        local-ip 192.168.11.0/24
        remote-ip 192.168.22.0/24
      exit
    exit
  exit
interface "int-private-ipsec-1" tunnel create
  sap tunnel-1.private:3 create
  ipsec-tunnel "ipsec-tunnel-for-gre-tunnel" create
    security-policy 1
    local-gateway-address 10.1.1.1 peer 10.2.2.1 delivery-service 4
    dynamic-keying
      ike-policy 1
      pre-shared-key "pass"
      transform 1
    exit
  no shutdown
  exit
exit
interface "int-public-gre-1" create
  address 192.168.11.2/24
  sap tunnel-1.public:4 create
  exit
exit
interface "int-private-gre-1" tunnel create
  address 10.0.0.6/30
  sap tunnel-1.private:5 create
  ip-tunnel "protected-gre-tunnel" create
    dest-ip 10.0.0.5
    gre-header
    source 192.168.11.1
    remote-ip 192.168.22.1
    delivery-service 3
    no shutdown
  exit
  exit
exit
static-route-entry 192.168.22.0/24
  ipsec-tunnel "ipsec-tunnel-for-gre-tunnel"
  no shutdown
  exit
exit
bgp-ipvpn
  mpls
    route-distinguisher 64496:3
    vrf-target target:64496:3
    no shutdown
  exit
  exit
no shutdown
exit

```

The following displays a configuration example of the public side of the IPsec tunnel:

```

# on PE-1:
configure
  service
    ies 4 name "IES 4" customer 1 create

```

```

interface "public-ipsec-1" create
  address 10.1.1.2/24
  tos-marking-state untrusted
  sap tunnel-1.public:3 create
  exit
exit
interface "int2-PE-1-CE-2" create
  address 192.168.112.1/30
  sap 1/1/2:4 create
  exit
exit
no shutdown
exit

```

The following static route is configured in the base router on PE-1:

```

# on PE-1:
configure
  router Base
    static-route-entry 10.2.2.0/24
      next-hop 192.168.112.2
      no shutdown

```

The configuration is similar on CE-2.

The following command shows that the tunnel "protected-gre-tunnel" with SAP tunnel-1.private:5 is up:

```

*A:PE-1# show ip tunnel
=====
IP Tunnels
=====
TunnelName                SapId                SvcId    Admn
Local Address             OperRemoteAddress    DlvrySvcId Oper
-----
gre-tunnel-1              tunnel-1.private:1   1         Up
192.168.1.1                192.168.2.1         2         Up
protected-gre-tunnel    tunnel-1.private:5  3       Up
192.168.11.1           192.168.22.1      3       Up
-----
IP Tunnels: 2
=====

```

The following command shows the IP/GRE tunnel information for this IPsec tunnel:

```

*A:PE-1# show ip tunnel "protected-gre-tunnel"
=====
IP Tunnel Configuration Detail
=====
Service Id      : 3                Sap Id          : tunnel-1.private:5
Tunnel Name     : protected-gre-tunnel
Description     : None
GRE Header      : Yes
Delivery Service : 3
GRE Keys Set    : False
GRE Send Key    : N/A                GRE Receive Key : N/A
Admin State     : Up                  Oper State      : Up
Source Address  : 192.168.11.1

```

```

Remote Address   : 192.168.22.1
Backup Address  : (Not Specified)
Oper Remote Addr : 192.168.22.1
DSCP            : None
Reassembly      : inherit
Clear DF Bit    : false           IP MTU           : max
Encap IP MTU    : max
Pkt Too Big     : true
Pkt Too Big Num : 100             Pkt Too Big Intvl: 10 secs
Frag Required   : true
Frag Req Count  : 100             Frag Req Interval: 10 secs
Propagate IPv6 P* : true
Propagate IPv4 P* : true
Oper Flags      : None
Transport Profile: (Not Specified)
Last Oper Changed: 05/12/2022 08:46:26
Host ISA        : 1/2
TCP MSS Adjust
  Public        : Disabled
  Private       : Disabled
    
```

Target Address Table

| Destination IP | IP Resolved Status |
|----------------|--------------------|
| 10.0.0.5 | Yes |

=====
IP Tunnel Statistics: protected-gre-tunnel
=====

| | | | |
|-------------------|-----|------------|-----|
| Errors Rx | : 0 | Errors Tx | : 0 |
| Pkts Rx | : 0 | Pkts Tx | : 0 |
| Bytes Rx | : 0 | Bytes Tx | : 0 |
| Key Ignored Rx | : 0 | Too Big Tx | : 0 |
| Seq Ignored Rx | : 0 | | |
| Vers Unsup. Rx | : 0 | | |
| Invalid Chksum Rx | : 0 | | |
| Key Mismatch Rx | : 0 | | |

=====
Fragmentation Statistics
=====

| | |
|-----------------------------|-------|
| Encapsulation Overhead | : 24 |
| Temporary Private MTU | : max |
| Pre-Encapsulation | |
| Fragmentation Count | : 0 |
| Last Fragmented Packet Size | : 0 |
| Post-Encapsulation | |
| Fragmentation Count | : 0 |
| Last Fragmented Packet Size | : 0 |

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

By default, the IPsec tunnel is down if it is not used by any traffic, as follows:

```
*A:PE-1# show ipsec tunnel
```

```
=====  
IPsec Tunnels
```

```
=====
```

| TunnelName SapId | LocalAddress RemoteAddress | SvcId DlvrySvcId | Admn Oper | Keying Sec Plcy |
|-----------------------------|-------------------------------|---------------------|--------------|-----------------------|
| ipsec-tunnel-for-gre-tunnel | 10.1.1.1 | 3 | Up | Dynamic |
| tunnel-1.private:3 | 10.2.2.1 | 4 | Down | 1 |

```
-----
```

IPsec Tunnels: 1

```
=====
```

After it is used by traffic, the status will be changed to be up.

```
*A:PE-1# ping router 3 10.0.0.5
PING 10.0.0.5 56 data bytes
64 bytes from 10.0.0.5: icmp_seq=1 ttl=64 time=1.42ms.
64 bytes from 10.0.0.5: icmp_seq=2 ttl=64 time=1.35ms.
64 bytes from 10.0.0.5: icmp_seq=3 ttl=64 time=1.26ms.
64 bytes from 10.0.0.5: icmp_seq=4 ttl=64 time=1.34ms.
64 bytes from 10.0.0.5: icmp_seq=5 ttl=64 time=1.28ms.

---- 10.0.0.5 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 1.26ms, avg = 1.33ms, max = 1.42ms, stddev = 0.058ms
```

The IPsec tunnel is now up, as follows:

```
*A:PE-1# show ipsec tunnel

=====
IPsec Tunnels
=====
```

| TunnelName SapId | LocalAddress RemoteAddress | SvcId DlvrySvcId | Admn Oper | Keying Sec Plcy |
|-----------------------------|-------------------------------|---------------------|--------------|-----------------------|
| ipsec-tunnel-for-gre-tunnel | 10.1.1.1 | 3 | Up | Dynamic |
| tunnel-1.private:3 | 10.2.2.1 | 4 | Up | 1 |

```
-----
```

IPsec Tunnels: 1

```
=====
```

BFD support on private tunnel interfaces

BFD is supported on IP interfaces associated with private IP/GRE tunnel SAPs. The BFD state of the interface can be used by static routes, OSPFv2, or BGP configured on the interface. It is not used to trigger a switchover to the backup remote IP address of the GRE tunnel.

The following displays a static route example:

```
# on PE-1:
configure
  service
    vprn 1 name "VPRN 1" customer 1 create
      interface "int-gre-tunnel" tunnel create
        address 10.0.0.1/30
        bfd 100 receive 100 multiplier 3
        sap tunnel-1.private:1 create
        ip-tunnel "gre-tunnel-1" create
```

```

        dest-ip 10.0.0.2
        gre-header
        source 192.168.1.1
        remote-ip 192.168.2.1
        delivery-service 2
        no dscp
        no shutdown
    exit
exit
interface "loopback1" create
    address 172.16.1.1/32
    loopback
exit
static-route-entry 172.16.2.1/32
    next-hop 10.0.0.2
    bfd-enable
    no shutdown
exit
exit
bgp-ipvpn
    mpls
        route-distinguisher 64496:1
        vrf-target target:64496:1
        no shutdown
    exit
exit
no shutdown

```

The following command shows that the BFD session on interface "int-gre-tunnel" is up for protocol static:

```

*A:PE-1# show router 1 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
Loc Addr                                  LAG name
-----
int-gre-tunnel                          Up       N/A       N/A
10.0.0.2                                  3         1000     1000
static                                    cpm-np  N/A       N/A
10.0.0.1
-----
No. of BFD sessions: 1
=====

```

When no static routes are configured and OSPF is configured instead, the configuration of VPRN 1 on PE-1 is as follows:

```

# on PE-1:
configure
    service
        vprn 1 name "VPRN 1" customer 1 create
        interface "int-gre-tunnel" tunnel create
        address 10.0.0.1/30

```

```

bfd 100 receive 100 multiplier 3
sap tunnel-1.private:1 create
  ip-tunnel "gre-tunnel-1" create
    dest-ip 10.0.0.2
    gre-header
    source 192.168.1.1
    remote-ip 192.168.2.1
    delivery-service 2
    no shutdown
  exit
exit
interface "loopback1" create
  address 172.16.1.1/32
  loopback
exit
bgp-ipvpn
  mpls
    route-distinguisher 64496:1
    vrf-target target:64496:1
    no shutdown
  exit
exit
ospf
  area 0.0.0.0
    interface "int-gre-tunnel"
      bfd-enable
      no shutdown
    exit
    interface "loopback1"
      no shutdown
    exit
  exit
  no shutdown
exit
no shutdown

```

The following shows that the BFD session is up for protocol OSPF on interface "int-gre-tunnel":

```
*A:PE-1# show router 1 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 |
| Loc Addr | | | LAG name |
| int-gre-tunnel | Up | N/A | N/A |
| 10.0.0.2 | 3 | 1000 | 1000 |
| ospf2 | cpm-np | N/A | N/A |
| 10.0.0.1 | | | |

```
-----
No. of BFD sessions: 1
=====
```

When BGP is configured instead of OSPF, the configuration of VPRN 1 on PE-1 is as follows:

```
# on PE-1:
configure
service
  vprn 1 name "VPRN 1" customer 1 create
  autonomous-system 64496
  interface "int-gre-tunnel" tunnel create
  address 10.0.0.1/30
  bfd 100 receive 100 multiplier 3
  sap tunnel-1.private:1 create
  ip-tunnel "gre-tunnel-1" create
  dest-ip 10.0.0.2
  gre-header
  source 192.168.1.1
  remote-ip 192.168.2.1
  delivery-service 2
  no shutdown
  exit
  exit
  interface "loopback1" create
  address 172.16.1.1/32
  bfd 100 receive 100 multiplier 3
  loopback
  exit
  static-route-entry 172.16.2.1/32
  next-hop 10.0.0.2
  no shutdown
  exit
  exit
  bgp-ipvpn
  mpls
  route-distinguisher 64496:1
  vrf-target target:64496:1
  no shutdown
  exit
  exit
  bgp
  group "group-1"
  type internal
  local-address 172.16.1.1
  neighbor 172.16.2.1
  bfd-enable
  exit
  exit
  no shutdown
  exit
  no shutdown
```

The following command shows that the BFD session is up for protocol BGP on interface "loopback1":

```
*A:PE-1# show router 1 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 Multipl   Tx Intvl  Rx Intvl
```

| Protocols Loc Addr | Type | LAG Port | LAG ID LAG name |
|------------------------|---------------|----------|--------------------|
| Loopback1 | Up | N/A | N/A |
| 172.16.2.1 | 3 | 1000 | 1000 |
| bgp | cpm-np | N/A | N/A |
| 172.16.1.1 | | | |
| ----- | | | |
| No. of BFD sessions: 1 | | | |
| ===== | | | |

IP/GRE termination – Advanced topics

DSCP value of outer delivery header

- Default behavior — The DSCP value from the payload header is copied into the outer GRE header. This is a one to one copy and no QoS classifications are required. It is performed when no DSCP value is configured under the private GRE tunnel.
- Non default behavior — DSCP is configured under the private SAP (following example using DSCP af41).

```
# on PE-1:
configure
  service
    vprn 1 name "VPRN 1" customer 1 create
      interface "int-gre-tunnel" tunnel create
        address 10.0.0.1/30
        sap tunnel-1.private:1 create
          ip-tunnel "gre-tunnel-1" create
            dest-ip 10.0.0.2
            gre-header
            source 192.168.1.1
            remote-ip 192.168.2.1
            delivery-service 2
            dscp af41
            no shutdown
          exit
        exit
      exit
    exit
  exit
---snip---
```

The log filter output shows TOS=88 (DSCP=af41) in the public network.

```
*A:PE-1# show filter log 102

=====
Filter Log
=====
Admin state : Enabled
Description : (Not Specified)
Destination : Memory
Wrap       : Enabled
-----
Maximum entries configured : 1000
Number of entries logged   : 5
-----
2022/05/12 08:54:19 Ip Filter: 2:10 Desc:
SAP: tunnel-1.private:1 Direction: Egress
```



```
Src MAC: 02-0f-ff-00-02-c9 Dst MAC: 00-00-00-07-a0-bd EtherType: 0800
Src IP: 10.0.0.1 Dst IP: 10.0.0.2 Flags: 0 TOS: 88 TTL: 64 Len: 84
Protocol: ICMP Type: Echo Request Code: 0
---snip---
```

IP MTU

It is possible to configure the IP MTU of a private tunnel SAP interface. This sets the maximum IP packet size payload (including IP header) that can be sent into the tunnel (it applies to the packet size before the tunnel encapsulation is added).

```
# on PE-1:
configure
  service
    vprn 1 name "VPRN 1" customer 1 create
      interface "int-gre-tunnel"
        ip-mtu 1476
      ---snip---
```

When an IPv4 packet needs to be forwarded to the tunnel and is larger than IP MTU bytes:

- If the DF bit is clear, the payload packet is IP fragmented to the MTU size prior to tunnel encapsulation.
- If the DF bit is set, the payload packet is discarded.

The IP MTU range supported is from 512 to 9000 bytes.

The following command shows the configured IP MTU and the operational IP MTU for the GRE tunnel:

```
*A:PE-1# show router 1 interface "int-gre-tunnel" detail | match MTU
IP MTU      : 1476
IP Oper MTU : 1476
```

Statistics and accounting

Collect-stats can be configured under public and private SAPs.

For public SAPs:

```
# on PE-1:
configure
  service
    ies "IES 2"
      interface "int-tunnel-public"
        sap tunnel-1.public:1
          collect-stats
        exit
      exit
```

For private SAPs:

```
# on PE-1:
configure
  service
    vprn "VPRN 1"
      interface "int-gre-tunnel"
        sap tunnel-1.private:1
```

```

collect-stats
exit
exit

```

Filtering, policing, and QoS

An IP filter and QoS policy can be applied to the ingress and egress traffic of the private and public SAPs.

Public SAPs:

```

# on PE-1:
configure
service
  ies "IES 2"
  interface "int-tunnel-public"
    sap tunnel-1.public:1
    ingress
      qos 10
      filter ip 1
    exit
    egress
      qos 20
      filter ip 2
    exit
  exit
exit

```

Private SAPs:

```

# on PE-1:
configure
service
  vprn "VPRN 1"
  interface "int-gre-tunnel"
    sap tunnel-1.private:1
    ingress
      qos 10
      filter ip 1
    exit
    egress
      qos 20
      filter ip 2
    exit
  exit
  ---snip---

```

Mirroring

The public and private SAPs can be mirrored.

```

# on PE-1:
debug
  mirror-source 99
  sap tunnel-1.private:3 egress ingress
  sap tunnel-1.public:1 egress ingress
  no shutdown
exit

```

```
exit
```

Conclusion

This chapter provides configuration and show commands for IP/GRE termination.

L2-aware NAT (with dNAT and MNPs)

This chapter provides information about Network Address Translation (NAT) in combination with Enhanced Subscriber Management (ESM).

Topics in this chapter include:

- [Applicability](#)
- [Overview](#)
- [Configuration](#)
- [Conclusion](#)
- [Appendix A – Generic ESM Configuration](#)
- [Appendix B – Logging](#)

Applicability

The information and configuration in this chapter are applicable to all SR OS nodes that support Broadband (BB) MS-ISAs, including the virtual version of MS-ISA in VSR, and was tested on SR OS 15.0.R4.

Overview

L2-aware network address translation (NAT) is an enhanced NAT functionality that offers better IPv4 address conservation ratios than traditional NAT44. L2-aware NAT relies on tight integration between enhanced subscriber management (ESM) and NAT, and in the context of ESM, it maps (or binds) an ESM subscriber to an outside IP address and a single port block (PB). In this fashion, the private (inside) IPv4 address in L2-aware NAT is abstracted from the binding.



Note:

Binding is distinct from the flow/session concept. Flow or session state is maintained for each specific conversation between the two end nodes (this includes IP addresses, ports, and protocol), whereas binding is mapping between an inside entity (the subscriber ID in L2-aware NAT or the IPv4 source IP address in large scale NAT44 (LSN44)) and the outside IP address and the PB. Flows/session can then consume ports within the PB.

On the public side (outside), L2-aware NAT allows multiple hosts (inside IPv4 addresses) under the same subscriber to share the same outside IPv4 address and a PB. This contrasts with traditional NAT44 where each inside IPv4 address (host) is mapped to a unique IPv4 address and PB combination.

On the private side (inside), the abstraction of the host IPv4 address from the binding allows assignment of the same inside IPv4 address to multiple hosts that belong to different subscribers. This sharing of an inside IPv4 address between hosts that belong to different subscribers means that each host still has its own instance of the shared IPv4 address. For example, host 1 of subscriber 1 and host 2 of subscriber 2 can be both assigned the same inside IP address (for instance 10.10.1.1).

A binding in L2-aware NAT is defined as the following:

```
<subId, nat-policy> → <outside routing context, outside srcIP address, outside PB>
```

Where:

Inside

Sub-id Subscriber ID in the ESM context.

Nat-policy NAT policy associated with the ESM subscriber. This association is performed within the sub-profile. A subscriber can have multiple NAT policies (MNPs), in which case each subscriber will have one binding per NAT policy.

Outside


Outside routing context Outside routing context that contains an L2-aware pool. NAT traffic is sent out of the node via this routing context.

Outside SrcIP address A subscriber in conjunction with the NAT policy is assigned this outside IP address. The source IP address of the subscriber traffic will be replaced by this IP address.

Outside PB A subscriber in conjunction with the NAT policy is assigned this outside PB. A source (TCP/UDP) port of the subscriber traffic will be replaced by one of the source ports from this PB. In the ICMP case, the query ID is selected from this PB.

A simplified example of an L2-aware NAT binding is shown in [Figure 73: NAT Binding](#). Subscriber 1 is mapped to the outside IP address 172.16.1.1 and PB 1. This means that the source IP addresses (10.10.1.1 and 10.10.1.2) for all hosts under Subscriber 1 will be translated to 172.16.1.1 and the source port will be translated to one of the ports from PB 1. Similar logic can be applied to Subscriber 2 mappings, where Subscriber 2 hosts will use PB 2 from the same outside IP address.

Figure 73: NAT Binding



26868

NAT binding in MS-ISA will be used to translate traffic arriving from the subscriber. When the first packet of the flow from the subscriber is received, the translation is performed and a flow state is created. This flow

state is maintained for the duration of the flow. All consecutive packets of the flow will rely on the flow state lookup for translation.

The flow mapping in L2-aware NAT is characterized as:

```
<subId, inside srcIP address, inside srcPort, foreign ip address, foreign port, protocol> →  
<outside routing context, outside IP address, outside port, destination ip address, foreign  
port, protocol>
```

Flow mapping is more specific than a binding (which is a partial mapping), and the fields in flow mapping are self-explanatory. The flow mapping contains the original source IP address, which is needed to determine where to send the return (downstream) traffic.

The subId field needs to be unique in the mappings at the subscriber level, allowing all the other fields (including the source IP address) to overlap between the subscribers. This is the basis for sharing of the source IP addresses between the subscribers.

In flow mapping, there is a naming distinction between the foreign IP address and the destination IP address. A foreign IP address is the original destination IP address in the packet sourced by the host. This foreign IP address can be translated by destination-based NAT (dNAT), just like source IP is translated by NAT. In this case, the translated foreign IP address on the outside in the SR OS CLI syntax becomes a destination IP address. This naming distinction helps to differentiate two IP addresses of the same field in the IP header, one before dNAT is performed and the other after dNAT is performed.

The reason why there is only a foreign port (and no destination port) is that in SR OS the original destination port is not translated by dNAT, while the original destination IP is.

L2-aware NAT requires that ESM and NAT are collocated in the same SR OS node. Because of the tight integration between ESM and NAT, logging of L2-aware NAT resources can be integrated into ESM accounting via AAA.

This chapter does not discuss L2-aware NAT intra-chassis redundancy. Intra-chassis redundancy for L2-aware NAT offers protection against MS-ISA failure. The supported modes are:

- Active/Standby MS-ISA
- Active/Active MS-ISA
- L2-aware Bypass

Each L2-aware NAT intra-chassis redundancy mode is described in the *7450 ESS, 7750 SR, and VSR Multiservice ISA and ESA Guide*.

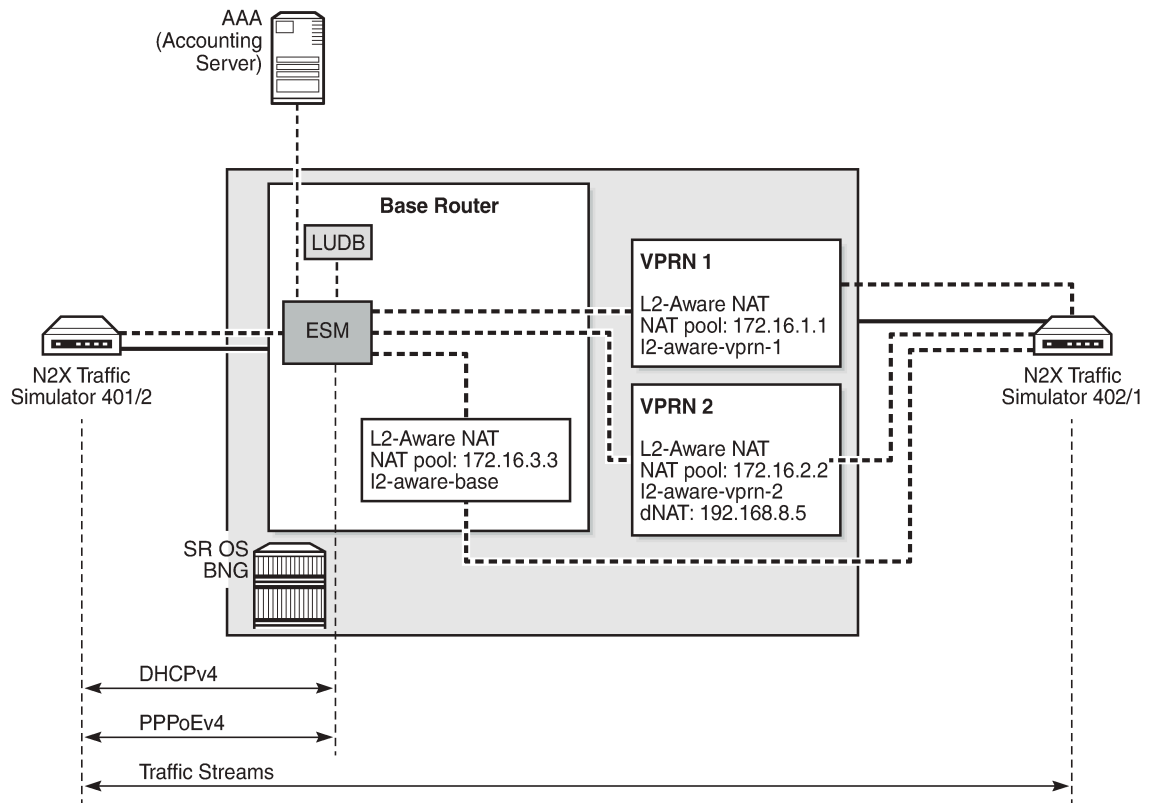
Configuration

The test setup in this configuration example will accommodate:

- Two L2-aware subscribers.
- Two hosts per L2-aware subscriber – one DHCPv4 host and one PPPoEv4 host per L2-aware subscriber.
- IP addresses overlap between the subscribers.
- Traffic split from a single host to three different NAT pools and outside routing contexts. This is achieved by selecting a NAT policy based on the destination IP address of the traffic.
- Destination-based NAT (dNAT) – some of the traffic is subjected to dNAT, where the destination IP address in the outgoing packet is translated (in addition to the source IP address and source port).

The test setup is shown in [Figure 74: Test Setup](#).

Figure 74: Test Setup



26866

The N2X traffic simulator on the left is used to simulate two homes (subscribers) with DHCPv4 and PPPoE hosts. Those hosts are terminated on an IES interface in a Base router in the BNG and they are further associated with the L2-aware NAT pools in VPRN1, VPRN2, and the Base router. ESM host authentication and IP address assignment is performed via local user database (LUDB). Accounting for the subscribers is sent to an AAA server and the accounting records contain NAT logging information.

Once the subscriber hosts are instantiated, traffic is sent through L2-aware NAT. This traffic is then used to further examine operation of L2-aware NAT in SR OS BNG.

The L2-aware NAT configuration in SR OS is split into two major parts:

- NAT-specific configuration – this is the subject of this document.
- ESM-specific configuration focuses on ESM subscribers and is described in other chapters in this volume. There are some key parts of ESM configuration that are essential for understanding L2-aware NAT operation and those parts will be described in this section.

For the sake of completeness, the remaining, more generic part of ESM configuration is described in [Appendix A – Generic ESM Configuration](#).

NAT-specific Configuration

L2-aware NAT requires the system IP address to be configured in the node. Without the system IP address, L2-aware NAT will not be operational. Lack of the system IP address within the L2-aware node is a common problem in most troubleshooting scenarios. The system IP is used within the system to pass traffic between the carrier IOM (where MS-ISA is attached) and the MS-ISA itself. The following is the command to configure the system IP address:

```
configure
router
  interface "system"
    address 192.0.2.2/32
    no shutdown
  exit
```

The **inside** NAT context in the following code identifies the private (inside) side of NAT, which in this example is configured in the Base routing context.

The address in the **l2-aware** context (in this example, address 10.10.1.254/24) represents the default gateway and an L2-aware subnet. All ESM hosts with IP addresses within the configured L2-aware subnet (in this example, 10.10.2.0/24) are subjected to L2-aware NAT.

ESM hosts with IP addresses out of the L2-aware subnet can be instantiated as regular (not L2-aware) hosts, effectively bypassing L2-aware NAT within the L2-aware enabled subscriber.

The **outside** NAT context contains the NAT pool configuration that is used for address translation. In this example, traffic is mapped to three outside contexts: Base, VPRN 1, and VPRN 2. IP address 172.16.3.3 is used for translation in the Base outside routing context, IP address 172.16.1.1 in the VPRN 1 outside routing context, and IP address 172.16.2.2 in the VPRN 2 outside routing context.

VPRN 1 and VPRN 2 also contain NAT pools with their own IP addresses.

```
configure
router
  nat
    inside
      l2-aware
        address 10.10.1.254/24
      exit
    exit
    outside
      pool "l2-aware-base" nat-group 1 type l2-aware create
        port-reservation ports 2000
        address-range 172.16.3.3 172.16.3.3 create
      exit
      no shutdown
    exit
  exit
exit
```

```
configure
service
  vprn 1 name "1" customer 1 create
  nat
    outside
      pool "l2-aware-vprn-1" nat-group 1 type l2-aware create
        port-reservation ports 2000
        address-range 172.16.1.1 172.16.1.1 create
      exit
      no shutdown
```



```

        exit
    exit
    exit
    vprn 2 name "2" customer 1 create
    nat
        outside
            port-reservation ports 2000
            pool "l2-aware-vprn-2" nat-group 1 type l2-aware create
            address-range 172.16.2.2 172.16.2.2 create
            exit
            no shutdown
        exit
    exit
    exit
    exit

```

NAT pool selection (and with this, the outside routing context) is performed based on the foreign IP address (traffic destination) in the packet. This is configured in the **nat-prefix-list**, which is then applied to the L2-aware subscriber via **sub-profile**.

Packets arriving from the inside and destined to networks 192.168.7.0/24 and 192.168.8.0/24 will be respectively routed to the VPRN 1 and VPRN 2 routing context, which are implied through NAT policies "l2-aware-vrf1" and "l2-aware-vrf2" referenced in the **nat-prefix-list**.

All other traffic will use the default NAT policy "l2-aware-base", which is directly referenced in the **sub-profile**. This NAT policy is pointing to the Base outside routing context. Also, traffic identified by NAT classifier "vrf2" will be subject to dNAT.

In the example, NAT classifier "vrf2" identifies UDP traffic destined to port 5001. This implies that destination IP address for traffic that is mapped to VPRN 2 with UDP destination port 5001 will be translated. The new destination IP address (192.168.8.5) is configured in the NAT classifier "vrf2", which is then applied via NAT policy "l2-aware-vrf2".

All other traffic (destined to a destination UDP port other than 5001) will pass transparently through the classifier without the destination IP address being modified. Traffic that is subject to dNAT can be identified (or classified) based on destination-port only.

The **block-limit** command in L2-aware pool is automatically set to "1". That is, each L2-aware subscriber is assigned exactly one PB. Additional PB allocation is not supported. That is why it is important that this single PB allocated to each L2-aware subscriber is initially configured with a large enough number of ports – this is performed via pool configuration (in this example, port-block size is set to 2000 ports). Once all the ports in this PB are exhausted, no new ports can be allocated.

```

configure
  subscriber-mgmt
    sub-profile "sub-prof-1" create
    nat-policy "l2-aware-base"
    nat-prefix-list "l2aware-mnp"
  exit

```

```

configure
  service
    nat
      nat-classifier "vrf2" create
      entry 1 create
        action dnat ip-address 192.168.8.5
        match protocol udp
          dst-port-range start 5001 end 5001
      exit

```

```
        exit
    exit
    nat-policy "l2-aware-base" create
        block-limit 1
        pool "l2-aware-base" router Base
    exit
    nat-policy "l2-aware-vrf1" create
        block-limit 1
        pool "l2-aware-vprn-1" router 1
    exit
    nat-policy "l2-aware-vrf2" create
        dnat
            nat-classifier "vrf2"
        exit
    exit
    nat-prefix-list "l2aware-mnp" application l2-aware-dest-to-policy create
        prefix 192.168.7.0/24 nat-policy "l2-aware-vrf1"
        prefix 192.168.8.0/24 nat-policy "l2-aware-vrf2"
    exit
exit
```

On the ESM side, **anti-spoof** configuration is of particular importance in L2-aware NAT.


- L2-aware deployments with bridged homes (multiple ESM hosts within a subscriber) must have the anti-spoofing option set to **ip-mac**. This will allow each host within the home to be distinctly identified in ESM and consequently receive proper ESM treatment (filters, QoS, accounting, and so on).
- L2-aware deployments with routed CPEs must have the anti-spoofing option set to **nh-mac**. In this case, only one host per ESM subscriber is allowed and this single host is the routed CPE. There can be multiple hosts behind this routed CPE and they will be all seen and properly treated (using NAT) by the MS-ISA. However, ESM will not be aware of them and all downstream traffic will be sent to the routed CPE (ESM host), which will then further route traffic to the proper destination.

From an ESM perspective, this approach is different from that for non-L2-aware ESM deployments. Non-L2-aware ESM deployments allow multiple ESM hosts per subscriber with the **nh-mac** anti-spoofing option enabled, where the ESM hosts have different MACs.

```
configure
    subscriber-mgmt
        msap-policy "msaps" create
            sub-sla-mgmt
                sub-ident-policy "sub_ident_pol"
                multi-sub-sap limit 10
            exit
        ies-vprn-only-sap-parameters
            anti-spoof ip-mac
        exit
    exit
```

This preceding described configuration is shown in [Figure 75: Logical Mapping of Subscribers to L2-aware Pool](#).

Figure 75: Logical Mapping of Subscribers to L2-aware Pool



26867

Logging

In this example, NAT logging is integrated with ESM accounting. This is configuration dependent and L2-aware NAT also supports syslog to convey information, about a subscriber that is hidden behind NAT, to the operator.

With RADIUS logging, NAT-specific information is carried in the **Aic-Nat-Port-Range** VSA, which must be explicitly enabled in **accounting-policy**. The ESM accounting model used in this example is per host-accounting with interim updates. For brevity, only the accounting records for host 1 of the subscriber 1 are shown following:

- **Acct Start** is generated when the ESM host is created. The first host for the subscriber will allocate NAT resources in the MS-ISA.
- **Acct Interim-Update** message is generated periodically (configuration dependent) and additionally carries information about consumed bytes flowing through queues/policers associated with the subscriber host.
- **Acct Stop** is generated when the ESM host is terminated. NAT resources are then released. The significance of this is that freeing the NAT resources coincides with the termination of the subscriber, without having to wait for NAT flows in the MS-ISA to time out, long after the user has ceased.

[Appendix B – Logging](#) lists a more comprehensive stream of logs for all four hosts that are established in this exercise.

```
root@linux:/var/log/freeradius/radacct/192.168.114.2# more detail-20170814
Mon Aug 14 09:14:45 2017
  Acct-Status-Type = Start
  NAS-IP-Address = 192.0.2.2
```

```
Framed-IP-Address = 10.10.1.1  
NAS-Identifrier = "BNG1"  
Acct-Session-Id = "D896FF0000002E59916D5E"  
Acct-Multi-Session-Id = "D896FF0000002F59916D5E"  
Event-Timestamp = "Aug 14 2017 02:29:02 PDT"  
NAS-Port-Type = Ethernet  
NAS-Port-Id = "1/1/5:1.1"  
ADSL-Agent-Circuit-Id = "sub-1"  
ADSL-Agent-Remote-Id = "host-1"  
Alc-Subsc-ID-Str = "sub-1"  
Alc-Subsc-Prof-Str = "sub-prof-1"  
Alc-SLA-Prof-Str = "sla-1"  
Alc-Nat-Port-Range = "172.16.3.3 1024-3023 router base l2-aware-base"  
Alc-Nat-Port-Range = "172.16.1.1 1024-3023 router 1 l2-aware-vrf1"  
Alc-Nat-Port-Range = "172.16.2.2 1024-3023 router 2 l2-aware-vrf2"  
Alc-Client-Hardware-Addr = "00:00:65:01:01:01"  
Acct-Unique-Session-Id = "fe5e55d7102d3f81"  
Timestamp = 1502727285  
Request-Authenticator = Verified
```

Mon Aug 14 09:20:11 2017

```
Acct-Status-Type = Interim-Update  
NAS-IP-Address = 192.0.2.2  
Framed-IP-Address = 10.10.1.1  
NAS-Identifrier = "BNG1"  
Acct-Session-Id = "D896FF0000002E59916D5E"  
Acct-Session-Time = 326  
Acct-Multi-Session-Id = "D896FF0000002F59916D5E"  
Event-Timestamp = "Aug 14 2017 02:34:28 PDT"  
NAS-Port-Type = Ethernet  
NAS-Port-Id = "1/1/5:1.1"  
ADSL-Agent-Circuit-Id = "sub-1"  
ADSL-Agent-Remote-Id = "host-1"  
Alc-Subsc-ID-Str = "sub-1"  
Alc-Subsc-Prof-Str = "sub-prof-1"  
Alc-SLA-Prof-Str = "sla-1"  
Alc-Nat-Port-Range = "172.16.3.3 1024-3023 router base l2-aware-base"  
Alc-Nat-Port-Range = "172.16.1.1 1024-3023 router 1 l2-aware-vrf1"  
Alc-Nat-Port-Range = "172.16.2.2 1024-3023 router 2 l2-aware-vrf2"  
Alc-Client-Hardware-Addr = "00:00:65:01:01:01"  
Alc-Acct-Triggered-Reason = regular  
Alc-Acct-I-Inprof-Octets-64 = 0x00010000000000000000  
Alc-Acct-I-Outprof-Octets-64 = 0x00010000000000000b220  
Alc-Acct-I-Inprof-Pkts-64 = 0x00010000000000000000  
Alc-Acct-I-Outprof-Pkts-64 = 0x0001000000000000001c8  
Alc-Acct-0-Inprof-Octets-64 = 0x00010000000000000000  
Alc-Acct-0-Outprof-Octets-64 = 0x00010000000000000000  
Alc-Acct-0-Inprof-Pkts-64 = 0x00010000000000000000  
Alc-Acct-0-Outprof-Pkts-64 = 0x00010000000000000000  
Acct-Unique-Session-Id = "fe5e55d7102d3f81"  
Timestamp = 1502727611  
Request-Authenticator = Verified
```

Mon Aug 14 09:30:34 2017

```
Acct-Status-Type = Stop  
NAS-IP-Address = 192.0.2.2  
Framed-IP-Address = 10.10.1.1  
NAS-Identifrier = "BNG1"  
Acct-Session-Id = "D896FF0000002E59916D5E"  
Acct-Session-Time = 949  
Acct-Terminate-Cause = User-Request  
Acct-Multi-Session-Id = "D896FF0000002F59916D5E"  
Event-Timestamp = "Aug 14 2017 02:44:51 PDT"  
NAS-Port-Type = Ethernet
```

```
NAS-Port-Id = "1/1/5:1.1"
ADSL-Agent-Circuit-Id = "sub-1"
ADSL-Agent-Remote-Id = "host-1"
Alc-Subsc-ID-Str = "sub-1"
Alc-Subsc-Prof-Str = "sub-prof-1"
Alc-SLA-Prof-Str = "sla-1"
Alc-Nat-Port-Range = "172.16.3.3 1024-3023 router base l2-aware-base"
Alc-Nat-Port-Range = "172.16.1.1 1024-3023 router 1 l2-aware-vrf1"
Alc-Nat-Port-Range = "172.16.2.2 1024-3023 router 2 l2-aware-vrf2"
Alc-Client-Hardware-Addr = "00:00:65:01:01:01"
Alc-Acct-I-Inprof-Octets-64 = 0x00010000000000000000
Alc-Acct-I-Outprof-Octets-64 = 0x0001000000000000037974
Alc-Acct-I-Inprof-Pkts-64 = 0x00010000000000000000
Alc-Acct-I-Outprof-Pkts-64 = 0x00010000000000000008e5
Alc-Acct-O-Inprof-Octets-64 = 0x00010000000000000000
Alc-Acct-O-Outprof-Octets-64 = 0x00010000000000000000
Alc-Acct-O-Inprof-Pkts-64 = 0x00010000000000000000
Alc-Acct-O-Outprof-Pkts-64 = 0x00010000000000000000
Acct-Unique-Session-Id = "fe5e55d7102d3f81"
Timestamp = 1502728234
Request-Authenticator = Verified

root@linux:/var/log/freeradius/radacct/192.168.114.2#
```

Operational Commands

In this example, three streams of traffic are run from host 1 of each subscriber (six streams in total). Traffic flows in the upstream direction (from the inside routing context to the outside routing context) and it will cause L2-aware NAT to allocate resources and create mappings that can be observed with the show and debug commands.

Host 1 of each subscriber is initiating three streams to three different destinations. Accordingly, traffic from the same host will be (based on the destination) routed to three different L2-aware NAT pools and outside routing contexts. One of the streams (stream 3) will, in addition to its source IP address/port, have its destination IP address translated as well.

The six traffic stream definitions are shown in [Table 11: Traffic Streams](#):

- All six streams have the same source IP address, 10.10.1.1. The significance of this is that both subscribers have hosts that are assigned the same inside (private) IP address.
- All streams carry UDP protocol.
- The source UDP port is 5000 for all streams.
- The UDP foreign port is 5000 except for the stream 3, which is assigned UDP foreign port 5001. Traffic with the foreign port 5001 is identified through the NAT classifier and is selected for dNAT (foreign IP address will be translated).
- Each host sends traffic to three different destinations and this is signified by three different foreign IP addresses:
 - The stream with the foreign IP address 192.168.5.2 is sent to the Base routing context (pool l2-aware-base).
 - The stream with the foreign IP address 192.168.7.2 is sent to the VPRN 1 routing context (pool l2-aware-vprn-1).
 - The stream with the foreign IP address 192.168.8.2 is sent to the VPRN 2 routing context (pool l2-aware-vprn-2).

- The column labelled Out IP (Outside IP address) represents a newly translated source (private) IP address by L2-aware NAT. Each of the three pools (one per routing context in Base, VPRN 1, and VPRN 2) has its own single outside IP address configured. In this example, each pool has a single outside IP address configured but, usually, multiple address ranges per NAT pool are supported.
- Each subscriber is allocated a single PB per pool with 2000 ports (configuration dependent) in each pool – this is reflected in the PB column. Subscriber 1 is allocated the same PB range in all three pools (one per routing context). A similar setup is shown for subscriber 2.

In this example, the PBs allocated per subscriber in three pools coincidentally match due to a low number of subscribers in the system (only two subscribers in our example). Usually, with an increased number of subscribers in the system, the probability for matching PB allocations per subscriber across pools would decrease. That is, each pool allocates PB independently of any other pool.

- The Out Port column indicates a new (translated) source UDP source port, after L2-aware NAT is performed.
- The Dest IP column indicates that the foreign IP address is modified only for stream 3, which is subject to dNAT (identified by UDP foreign port 5001 via a NAT classifier). The remaining streams do not have foreign IP addresses translated.
- The Destination ports on the outside in all three streams remain unchanged. This applies even for stream 3, which is subject to dNAT (only the foreign IP address is translated by dNAT functionality).

Dynamically allocated parameters by L2-aware NAT are in the following traffic streams table (fields in bold typeface - **Out IP, PB, Out Port, Dest IP**) populated based on the observation in the system once the traffic is run. The values are collected based on the output provided by debug commands and a **tools dump nat sessions** command (both of them are shown further in the text).

Table 11: Traffic Streams

| Host | Strm | Packet Fields Before NAT (inside routing context – Base) | | | | | Packet Fields After NAT (outside routing context) | | | | | |
|------------|------|--|------|----------|------------|-------------|---|-------------------|-----------|-------------|--------------------|-----------|
| | | Src IP | Prot | Src Port | Forgn Port | Forgn IP | Outside Rtr | Out IP | PB | Out Port | Dest IP | Dest Port |
| Sub1 Host1 | 1 | 10.10.1.1 | UDP | 5000 | 5000 | 192.168.5.2 | Base | 172.16.3.3 | 1024-3023 | 2186 | 192.168.5.2 | 5000 |
| | 2 | | | | | 192.168.7.2 | VRF1 | 172.16.1.1 | 1024-3023 | 1124 | 192.168.7.2 | |
| | 3 | | | | 5001 | 192.168.8.2 | VRF2 | 172.16.2.2 | 1024-3023 | 3690 | 192.168.8.5 | 5001 |

| Host | Strm | Packet Fields Before NAT (inside routing context – Base) | | | | | Packet Fields After NAT (outside routing context) | | | | | |
|------------|------|--|------|----------|------------|-------------|---|------------|-----------|----------|-------------|-----------|
| | | Src IP | Prot | Src Port | Forgn Port | Forgn IP | Outside Rtr | Out IP | PB | Out Port | Dest IP | Dest Port |
| Sub2 Host1 | 4 | 10.10.1.1 | UDP | 5000 | 5000 | 192.168.5.2 | Base | 172.16.3.3 | 3024-5023 | 4470 | 192.168.5.2 | 5000 |
| | 5 | | | | | 192.168.7.2 | VRF1 | 172.16.1.1 | 3024-5023 | 4666 | 192.168.7.2 | |
| | 6 | | | | | 192.168.8.2 | VRF2 | 172.16.2.2 | 3024-5023 | 4698 | 192.168.8.2 | |

Show Commands

A good starting point for troubleshooting is a generic NAT command, **show service nat overview**, with the output that would normally point to misconfiguration:

```
*A:BNG1# show service nat overview
=====
NAT overview
=====
Inside/      Policy/
Outside     Pool
Type
-----
*           l2-aware-base           l2aware
Base       l2-aware-base
*           l2-aware-vrf1           l2aware
vprn1     l2-aware-vprn-1
*           l2-aware-vrf2           l2aware
vprn2     l2-aware-vprn-2
=====

No firewall policies found.
*A:BNG1#
```

The output of the show port statistics command provides information relevant to packet flow between the carrier IOM (to which the MS-ISA is attached) and the MS-ISA itself.

In the following output, "1/2/nat-in-l2" is the port on the carrier IOM that represents the private (inside) side of NAT and "1/2/nat-out-ip" is the port on the carrier IOM that represents the public (outside) side of NAT. In this example, traffic is sent only in one direction, from private side to public side, and therefore:

- Traffic is egressing port "1/2/nat-in-l2" on the IOM and entering MS-ISA on the private side (11789 packets).
- Traffic is leaving MS-ISA on the public side and ingressing port "1/2/nat-out-ip" on the IOM (11789 packets).

```
*A:BNG1# show port statistics

=====
Port Statistics on Slot 1
=====
Port Id           Ingress Packets    Ingress Octets    Egress Packets    Egress Octets
-----
1/1/2             223162            25147835         0                 0
1/1/3             815114            78685140         614895            60630323
1/1/4             15484             1531339          230226            25728265
1/1/5             84116             8375228          1054              69854
1/1/10            5115              835288           345               110442
1/2/nat-out-ip    11789            1131744         0                0
1/2/nat-in-l2     0                0                11789            1532570
=====

Port Statistics on Slot A
=====
Port Id           Ingress Packets    Ingress Octets    Egress Packets    Egress Octets
-----
A/1               81952168          0                 28663             2721059
=====
```

NAT resources are allocated during the subscriber instantiation phase (when the first host for an L2-aware subscriber is created), before any data traffic from the subscriber side is even initiated.

The following command is run to verify that the two subscribers, each with two hosts (one DHCP and one PPPoE) are online. In addition, this command lists limited input related to L2-aware NAT:

- NAT policies associated with the subscriber
- Outside IP address allocated to the subscriber
- Port range in a PB

As shown in the output of this command, both L2-aware subscribers have hosts with overlapping IP addresses, which is a unique characteristic in L2-aware NAT functionality.

```
*A:BNG1# show service active-subscribers

=====
Active Subscribers
=====
Subscriber sub-1 (sub-prof-1)
-----
NAT Policy: l2-aware-base
Outside IP: 172.16.3.3
Ports      : 1024-3023

NAT Policy: l2-aware-vrf1
Outside IP: 172.16.1.1 (vprn1)
Ports      : 1024-3023

NAT Policy: l2-aware-vrf2
```



```

Outside IP: 172.16.2.2 (vprn2)
Ports      : 1024-3023
-----
(1) SLA Profile Instance sap:[1/1/5:1.1] - sla:sla-1
-----
IP Address          MAC Address          Session          Origin          Svc          Fwd
-----
10.10.1.2          00:00:65:01:01:02    PPP 1           IPCP            3            Y
10.10.1.1          00:00:65:01:01:01    N/A            DHCP            3            Y
-----

Subscriber sub-2 (sub-prof-1)
-----
NAT Policy: l2-aware-base
Outside IP: 172.16.3.3
Ports      : 3024-5023

NAT Policy: l2-aware-vrf1
Outside IP: 172.16.1.1 (vprn1)
Ports      : 3024-5023

NAT Policy: l2-aware-vrf2
Outside IP: 172.16.2.2 (vprn2)
Ports      : 3024-5023
-----
(1) SLA Profile Instance sap:[1/1/5:1.2] - sla:sla-1
-----
IP Address          MAC Address          Session          Origin          Svc          Fwd
-----
10.10.1.2          00:00:65:01:02:02    PPP 1           IPCP            3            Y
10.10.1.1          00:00:65:01:02:01    N/A            DHCP            3            Y
-----

Number of active subscribers : 2
=====
*A:BNGL#

```

NAT-specific **show** commands can provide more information about the L2-aware subscribers. One such command is **show service nat l2-aware-subscribers**. In this example, this command is run in a simplest form, but it can be expanded with additional filtering capabilities:

```

show service nat l2-aware-subscribers [nat-policy <policy-name>] [nat-group <nat-group-id>]
[member <[1..255]>] [outside-router <router-instance>] [outside-ip <outside-ip-address>]
[firewall-policy <policy-name>]

```

```

*A:BNGL# show service nat l2-aware-subscribers

```

```

=====
Layer-2-Aware NAT subscribers
=====

```

| | |
|----------------------------------|-------------------|
| Subscriber | : sub-1 |
| ----- | |
| ISA NAT group | : 1 |
| ISA NAT group member | : 1 |
| UPnP policy | : (None) |
| Default NAT policy | : l2-aware-base |
| Per-host port block size | : N/A |
| Firewall policy | : (None) |
| Policy | : l2-aware-base |
| Purpose | : nat |
| Outside router | : Base |
| Outside IP | : 172.16.3.3 |
| DNAT default IP address override | : (Not Specified) |
| DNAT disabled by override | : false |
| Ports | : 1024-3023 |
| Policy | : l2-aware-vrf1 |
| Purpose | : nat |
| Outside router | : vprn1 |
| Outside IP | : 172.16.1.1 |
| DNAT default IP address override | : (Not Specified) |
| DNAT disabled by override | : false |
| Ports | : 1024-3023 |
| Policy | : l2-aware-vrf2 |
| Purpose | : nat |
| Outside router | : vprn2 |
| Outside IP | : 172.16.2.2 |
| DNAT default IP address override | : (Not Specified) |
| DNAT disabled by override | : false |
| Ports | : 1024-3023 |
| Subscriber | : sub-2 |
| ----- | |
| ISA NAT group | : 1 |
| ISA NAT group member | : 1 |
| UPnP policy | : (None) |
| Default NAT policy | : l2-aware-base |
| Per-host port block size | : N/A |
| Firewall policy | : (None) |
| Policy | : l2-aware-base |
| Purpose | : nat |
| Outside router | : Base |
| Outside IP | : 172.16.3.3 |
| DNAT default IP address override | : (Not Specified) |
| DNAT disabled by override | : false |
| Ports | : 3024-5023 |
| Policy | : l2-aware-vrf1 |
| Purpose | : nat |
| Outside router | : vprn1 |
| Outside IP | : 172.16.1.1 |
| DNAT default IP address override | : (Not Specified) |
| DNAT disabled by override | : false |
| Ports | : 3024-5023 |
| Policy | : l2-aware-vrf2 |
| Purpose | : nat |
| Outside router | : vprn2 |
| Outside IP | : 172.16.2.2 |

```
DNAT default IP address override : (Not Specified)
DNAT disabled by override       : false
Ports                           : 3024-5023
```

```
-----
No. of subscribers: 2
=====
```

Host-level L2-aware NAT information can be obtained via the **show service nat l2-aware-hosts** command. This command also provides additional filters for more targeted results:

```
show service nat l2-aware-hosts [outside-router <router-instance>] [outside-ip <outside-ip-
address>] [inside-ip-prefix <ip-prefix/mask>]
```

In this case, the most generic version of this command is run where all four hosts and their L2-aware NAT specific information is shown:

```
*A:BNGL1# show service nat l2-aware-hosts
```

```
=====
Layer-2-Aware NAT hosts
=====
```

```
Subscriber           : sub-1
Inside IP address    : 10.10.1.1
-----
Policy               : l2-aware-base
Bypassing            : false
Outside router       : "Base"
Outside IP address   : 172.16.3.3
Port block           : N/A

Policy               : l2-aware-vrf1
Bypassing            : false
Outside router       : 1
Outside IP address   : 172.16.1.1
Port block           : N/A

Policy               : l2-aware-vrf2
Bypassing            : false
Outside router       : 2
Outside IP address   : 172.16.2.2
Port block           : N/A

Subscriber           : sub-1
Inside IP address    : 10.10.1.2
-----
Policy               : l2-aware-base
Bypassing            : false
Outside router       : "Base"
Outside IP address   : 172.16.3.3
Port block           : N/A

Policy               : l2-aware-vrf1
Bypassing            : false
Outside router       : 1
Outside IP address   : 172.16.1.1
Port block           : N/A

Policy               : l2-aware-vrf2
Bypassing            : false
```

```

Outside router      : 2
Outside IP address  : 172.16.2.2
Port block         : N/A

Subscriber       : sub-2
Inside IP address : 10.10.1.1
-----
Policy             : l2-aware-base
Bypassing         : false
Outside router    : "Base"
Outside IP address : 172.16.3.3
Port block        : N/A

Policy             : l2-aware-vrf1
Bypassing         : false
Outside router    : 1
Outside IP address : 172.16.1.1
Port block        : N/A

Policy             : l2-aware-vrf2
Bypassing         : false
Outside router    : 2
Outside IP address : 172.16.2.2
Port block        : N/A

Subscriber       : sub-2
Inside IP address : 10.10.1.2
-----
Policy             : l2-aware-base
Bypassing         : false
Outside router    : "Base"
Outside IP address : 172.16.3.3
Port block        : N/A

Policy             : l2-aware-vrf1
Bypassing         : false
Outside router    : 1
Outside IP address : 172.16.1.1
Port block        : N/A

Policy             : l2-aware-vrf2
Bypassing         : false
Outside router    : 2
Outside IP address : 172.16.2.2
Port block        : N/A

-----
No. of hosts: 4
=====

```

The following two commands are more generic and their output displays packet and event statistics related to NAT operation in two contexts:

- Per MS-ISA
- Per member within the ISA

A member within an MS-ISA is a concept related to intra-chassis redundancy. CPM maintains a copy of the member configuration so that it can download it to a rescuing MS-ISA during the switchover (when an MS-ISA fails). An MS-ISA can contain one or more (in A/A redundancy model) members.

```

show isa nat-group 1 mda 1/2 statistics
show isa nat-group 1 member 1 statistics

```

The output is the following:

```
*A:BNGL1# show isa nat-group 1 mda 1/2 statistics
=====
ISA NAT Group 1 MDA 1/1
=====
no resource : 0
pkt rx on wrong port : 0
unsupported protocol : 0
no host or host group : 0
no ip or port : 0
no matching flow : 0
max flow exceeded : 0
TCP no flow for RST : 0
TCP no flow for FIN : 0
TCP no flow : 0
addr. dep. filtering : 0
ICMP type unsupported : 0
ICMP local unsupported : 0
ICMP/ICMPv6 checksum error : 0
ICMP embedded checksum error : 0
ICMP/ICMPv6 unsupported embedded L4 : 0
ICMP/ICMPv6 too short : 0
ICMP/ICMPv6 length error : 0
Pkt not IPv4 or IPv6 : 0
Pkt rcv error : 0
Pkt error : 0
IPv4 header checksum violation : 0
IP header malformed : 0
IP malformed packet : 0
IP ttl zero : 0
IPv4 opt /IPv6 ext headers : 0
IP undefined error : 0
IPv6 fragments unsupported : 0
TCP/UDP malformed : 0
TCP/UDP checksum failure : 0
Pkt send error : 0
no buf to copy pkt : 0
no policy : 0
policy not in use by subscriber : 0
locked by mgmt core : 0
port range log failed : 0
MTU exceeded : 0
DS Lite unrecognized next hdr : 0
DS Lite unknown AFTR : 0
too many fragments for IP packet : 0
too many fragmented packets : 0
too many fragment holes : 0
too many frags buffered : 0
fragment list expired : 0
fragment rate too high : 0
flow log failed : 0
no multiple host or subscr. IPs allowed : 0
NAT64 disabled : 0
NAT64 invalid src addr : 0
NAT64 frag has zero checksum : 0
NAT64 v4 has zero checksum : 0
NAT64 ICMP frag unsupported : 0
Reassembly Failures : 0
subscriber-id unknown : 0
packet hashed to wrong MDA : 0
Nptv6 map failed : 0
Nptv6 IID all 1's : 0
UPnP rate too high : 0
```

```
Dest. IP is unroutable : 0
no TCP/UDP checksum : 0
to local : 0
to local ignored : 0
new flow : 0
TCP closed : 0
TCP expired : 0
UDP expired : 0
ICMP/ICMPv6 expired : 0
GRE expired : 0
ICMP local : 0
found flow : 0
ARPs ignored : 10
Fragments RX L2A : 0
Fragments RX LSN : 0
Fragments RX DSL : 0
Fragments RX DORMANT : 0
Fragments RX OUT : 0
Fragments TX L2A : 0
Fragments TX LSN : 0
Fragments TX DSL : 0
Fragments TX DORMANT : 0
Fragments TX NAT64 : 0
Fragments TX OUT : 0
flow create logged : 0
flow delete logged : 0
flow log pkt tx : 0
Reassembled Pkts : 0
subscriber-id cached : 0
subscriber-id delayed : 0
subscriber-id timeout : 0
forwarded DS-Lite V6 pkts : 0
V6 pkts fragmented : 0
dropped multicast : 0
SSDP rx : 0
SSDP tx : 0
SSDP dropped : 0
UPnP rx : 0
UPnP tx : 0
UPnP dropped : 0
no radius resources : 0
no radius connection resources : 0
Dest. NAT dest. IP mismatch : 0
Dest. NAT foreign IP mismatch : 0
temp. no policy : 0
no SNat enabled : 0
no default policy : 0
flow create failed, key ambiguous : 0
flow create failed, conflicting policies : 0
NAT64 unrecognized next hdr : 0
ICMP/ICMPv6 unsupported embedded L3 : 0
TCP reset waiting : 0
IPv6 downstream prohibited : 0
IPv6 ext. hdr parse error : 0
any V6 L4 expired : 0
ICMP embedded IPv6 ext. hdr parse error : 0
ICMPv6 type unsupported : 0
ICMP/ICMPv6 fragmented error : 0
Service-chaining rx : 0
Service-chaining tx : 0
Service-chaining encapsulation error : 0
Service-chaining decapsulation error : 0
Service-chaining filter drop : 0
PPPoE uplink down : 0
```

```
ICMP6 PTB dropped (mtu < 1280)           : 0
firewall addr. dep. filtering             : 0
unresolved L2-aware V6 host              : 0
=====
*A:BNG1#
```

Resources monitoring is performed via the following commands:

```
tools dump nat isa resources nat-group <id> member <id>
tools dump nat isa resources mda <card-id/mds-id>
show isa nat-system-resources nat-group <id> member <id>
```

The NAT scale depends on the hardware (vSIMs, MS-ISA1, MS-ISA2, VSR-NAT).

Tools Commands

To display the complete session state, the following command can be used:

```
tools dump nat sessions [nat-group <nat-group-id>] [mda <mda-id>] [protocol {gre|icmp|
icmp6|tcp|udp|unknown}] [inside-ip <ip-prefix[/ip-prefix-length]>] [inside-router <router-
instance>] [inside-port <port-number>] [outside-ip <ipv4-address>] [outside-port <port-
number>] [foreign-ip
  <ip-address>] [foreign-port <port-number>] [dslite-address
  <ipv6-address>] [wlan-gw-ue <ieee-address>] [next index <index>] [upnp] [member
  <member-id>] [nat-policy <policy-name>] [dest-ip <ip-address>] [firewall-policy <policy-
  name>] [address-type <addr-type>] [vas-filter
  <filter-name>] [vas-filter-entry <vas-filter-entry-id>]
  [l2-aware-subscriber <sub-ident>]
```

The output of this command is used to populate dynamically allocated fields (in bold) in [Table 11: Traffic Streams](#). The command, in its basic form, is run as:

```
*A:BNG1# tools dump nat sessions

=====
Matched 6 sessions on Slot #1 MDA #2
=====
Owner           : L2-aware Subscr sub-1
Policy          : l2-aware-vrfl
FlowType       : UDP                Timeout (sec)      : 299
Inside IP Addr  : 10.10.1.1
Inside Port     : 5000
Outside IP Addr : 172.16.1.1
Outside Port    : 1124
Foreign IP Addr : 192.168.7.2
Foreign Port    : 5000
Dest IP Addr    : 192.168.7.2
Nat Group       : 1
Nat Group Member : 1
-----
Owner           : L2-aware Subscr sub-2
Policy          : l2-aware-vrfl
FlowType       : UDP                Timeout (sec)      : 300
Inside IP Addr  : 10.10.1.1
Inside Port     : 5000
Outside IP Addr : 172.16.1.1
Outside Port    : 4666
Foreign IP Addr : 192.168.7.2
Foreign Port    : 5000
```

```
Dest IP Addr      : 192.168.7.2
Nat Group         : 1
Nat Group Member  : 1
-----
Owner             : L2-aware Subscr sub-2
Policy            : l2-aware-base
FlowType         : UDP                Timeout (sec)      : 300
Inside IP Addr   : 10.10.1.1
Inside Port      : 5000
Outside IP Addr  : 172.16.3.3
Outside Port     : 4470
Foreign IP Addr : 192.168.5.2
Foreign Port   : 5000
Dest IP Addr   : 192.168.5.2
Nat Group        : 1
Nat Group Member : 1
-----
Owner             : L2-aware Subscr sub-1
Policy            : l2-aware-vrf2
FlowType         : UDP                Timeout (sec)      : 300
Inside IP Addr   : 10.10.1.1
Inside Port      : 5000
Outside IP Addr  : 172.16.2.2
Outside Port     : 2068
Foreign IP Addr  : 192.168.8.2
Foreign Port     : 5001
Dest IP Addr     : 192.168.8.5
Nat Group        : 1
Nat Group Member : 1
-----
Owner             : L2-aware Subscr sub-1
Policy            : l2-aware-base
FlowType         : UDP                Timeout (sec)      : 299
Inside IP Addr   : 10.10.1.1
Inside Port      : 5000
Outside IP Addr  : 172.16.3.3
Outside Port     : 2186
Foreign IP Addr  : 192.168.5.2
Foreign Port     : 5000
Dest IP Addr     : 192.168.5.2
Nat Group        : 1
Nat Group Member : 1
-----
Owner             : L2-aware Subscr sub-2
Policy            : l2-aware-vrf2
FlowType         : UDP                Timeout (sec)      : 299
Inside IP Addr   : 10.10.1.1
Inside Port      : 5000
Outside IP Addr  : 172.16.2.2
Outside Port     : 3690
Foreign IP Addr  : 192.168.8.2
Foreign Port     : 5000
Dest IP Addr     : 192.168.8.2
Nat Group        : 1
Nat Group Member : 1
-----
```


Clear Commands

The following command clears L2-aware subscribers:

```
clear nat isa nat-group 1 member 1 l2-aware-
```

**Note:**

ESM subscribers are also deleted with this command. Therefore, this command should be used with caution because DHCP hosts in ESM do not send termination messages back to the client. This leaves the DHCP client in a state where it is not aware that its DHCP state in the L2-aware node has been deleted.

Debug Commands

Troubleshooting ESM subscribers is described in other configuration guides focusing on ESM. This chapter focuses only on the NAT aspect of ESM in L2-aware NAT.

L2-aware NAT debug output shows session (or flow) initiation information. As traffic is run, translations of the source IP address/port are performed and this information is displayed for debugging purposes. Traffic from both subscribers in this example is debugged with these two commands:

```
*A:BNG1# show debug
debug
  nat
    l2-aware-sub id 1 subscriber "sub-1"
    l2-aware-sub id 2 subscriber "sub-2"
  exit
exit
```

The following configuration sends debug output to the Telnet/SSH sessions screen:

```
A:BNG1>config>log# info
-----
  log-id 50
    from debug-trace
    to session
    no shutdown
  exit
```

Debug output shows translations that are being performed on the incoming traffic, which consists of six streams in this example:

```
1 2017/08/31 08:37:39.223 WEST MINOR: DEBUG #2001 Base NAT "NAT: L2-Aware@sub-1
Initiated flow from 172.16.3.3(port 2186) to 192.168.5.2(port 5000) protocol UDP"

2 2017/08/31 08:37:39.223 WEST MINOR: DEBUG #2001 Base NAT "NAT: L2-Aware@sub-1
Initiated flow from 172.16.1.1(port 1124) to 192.168.7.2(port 5000) protocol UDP"

3 2017/08/31 08:37:39.483 WEST MINOR: DEBUG #2001 Base NAT "NAT: L2-Aware@sub-1
Initiated flow from 172.16.2.2(port 3690) to 192.168.8.2(port 5001) protocol UDP"

4 2017/08/31 08:37:39.743 WEST MINOR: DEBUG #2001 Base NAT "NAT: L2-Aware@sub-2
Initiated flow from 172.16.3.3(port 4470) to 192.168.5.2(port 5000) protocol UDP"

5 2017/08/31 08:37:39.743 WEST MINOR: DEBUG #2001 Base NAT "NAT: L2-Aware@sub-2
Initiated flow from 172.16.1.1(port 4666) to 192.168.7.2(port 5000) protocol UDP"
```

```
6 2017/08/31 08:37:40.003 WEST MINOR: DEBUG #2001 Base NAT "NAT: L2-Aware@sub-2  
Initiated flow from 172.16.2.2(port 4698) to 192.168.8.2(port 5000) protocol UDP"
```

Note that dNAT information is not displayed in the debug output.

Conclusion

L2-aware NAT integrates NAT functionality into ESM. NAT resources (outside IP address and a PB) are allocated per ESM subscriber and, consequently, all the hosts within the subscriber share those NAT resources. Each subscriber is initially assigned a single PB that will be used during the lifetime of the subscriber. The number of ports in this PB must be large enough to accommodate the needs of the subscriber during its lifetime.

The integration between ESM and NAT in L2-aware NAT introduces the following:

- Using the subscriber-id field (instead of source IPv4 address) to identify the NAT binding allows the source IPv4 address to be abstracted from the binding. This allows hosts of different subscribers to have overlapping IP addresses .
IPv4 addresses cannot overlap between the hosts of the same subscriber, but they can overlap between the hosts of different subscribers.
- NAT logging can be integrated into ESM RADIUS accounting, simplifying operations and removing the need for a separate network resource dedicated to the collection of NAT logs.
- NAT resources are released when the ESM subscriber is disconnected. This prevents lingering of NAT resources in the system waiting for the timer of the last session to expire long after the entity that initially requested those resources ceased to exist. This improves resource utilization in a system with a high number of subscribers.

L2-aware NAT should be considered for deployment because it allows coupling between ESM and NAT.

Appendix A – Generic ESM Configuration

Appendix A describes the generic part of the ESM configuration.

ESM configuration starts with a **subscriber-interface** configured in an inside routing context. In this example, subscriber hosts are instantiated in the 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 local user database (LUDB) user-db "user-db-1". The IP addresses that are assigned to the hosts are statically configured in LUDB (no DHCP server is used in this setup).

```
configure
  service
    ies 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
    pppoe
        anti-spoof mac-sid-ip
        policy "pppoe_pol"
        session-limit 100
        sap-session-limit 100
        user-db "user-db-1"
        no shutdown
    exit
    no shutdown
exit
exit
exit
service-name "ies-3"
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/PPPoEv4 host instantiation, as follows:

```

configure
  service
    vpls 10 name "10" customer 1 create
    sap 1/1/5:1.* capture-sap create
      trigger-packet dhcp pppoe
      dhcp-user-db "user-db-1"
      pppoe-policy "pppoe_pol"
      pppoe-user-db "user-db-1"
    exit
  exit

```

The first interaction between ESM and L2-aware NAT is performed through a **sub-profile**. Specifically, a **nat-policy** within the **sub-profile** is used to associate the ESM subscriber with NAT.

A **nat-prefix-list** within the sub-profile provides the place to configure multiple nat-policies per subscriber. Selection of **nat-policy** (and with this, the NAT pool and outside routing context) is based on the destination IP address in the packet sent by the subscriber.

RADIUS accounting records passed to the accounting server will contain subscriber-host specific information, including NAT logs. This is configured via a RADIUS accounting policy, which is referenced in the **sub-profile**.

```

configure
  subscriber-mgmt
    sub-profile "sub-prof-1" create
      nat-policy "l2-aware-base"
      nat-prefix-list "l2aware-mnp"
      radius-accounting
        policy "acct"
    exit
  exit

```

Although an **sla-profile** is a mandatory configuration for a subscriber, the **sla-profile** content is not relevant for understanding L2-aware NAT concepts. A basic **sla-profile** is chosen with default settings: a default **qos-policy 1** and no **ip-filters**.

The **sub-profile** and **sla-profile** are associated with the subscriber(-host) during the authentication phase and, in this case, this is achieved through LUDB.

```
configure
  subscriber-mgmt
    sla-profile "sla-1" create
  exit
```

The MSAP policy is a mandatory configuration for dynamically created SAPs (Managed SAP or MSAPs). It is used to determine SAP parameters during the MSAP creation process based on control traffic (DHCP or PPPoE) of the first host (MSAP trigger traffic). Parameters defined in **msap-policy** that are of relevance to our example are:

- **Sub-ident-policy** → this is a mandatory parameter in ESM that governs mapping of subscriber-related profiles (sub-profile, SLA-profile) to the subscriber host during the authentication phase.
- **Anti-spoof filter** → this is of particular importance in L2-aware NAT and is already described in more detail in the [Configuration](#) section.

```
configure
  subscriber-mgmt
    msap-policy "msaps" create
    sub-sla-mgmt
      sub-ident-policy "sub_ident_pol"
      multi-sub-sap limit 10
    exit
  ies-vprn-only-sap-parameters
    anti-spoof ip-mac
  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. Subscriber 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
```

In this example, authentication of the subscriber hosts and IP address assignment is performed through LUDB. The hosts are identified based on **circuit-id** and **remote-id** fields in DHCP and PPPoE control packets. **Sla-sub-profile** strings in the LUDB are directly mapped to the configured sub/sla-profiles in SR OS node. This direct mapping is implied by the preceding **use-direct-map-as-default** command

within **sub-ident-policy**. The LUDB carries only ESM specific configuration. There is no NAT relevant configuration present in the 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 10.10.1.1
        identification-strings 254 create
          subscriber-id "sub-2"
          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
    exit
  ppp
    match-list circuit-id remote-id
    host "sub-1-host-2" create
      host-identification
        circuit-id string "sub-1"
        remote-id string "host-2"
      exit
      address 10.10.1.2/24
      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
        no shutdown
    exit
    host "sub-2-host-2" create
        host-identification
            circuit-id string "sub-2"
            remote-id string "host-2"
        exit
        address 10.10.1.2/24
        identification-strings 254 create
            subscriber-id "sub-2"
            sla-profile-string "sla-1"
            sub-profile-string "sub-prof-1"
        exit
        msap-defaults
            group-interface "group-int-1"
            policy "msaps"
            service 3
        exit
        no shutdown
    exit
    exit
    no shutdown
exit
```

The accounting policy identifies the type of accounting and the attributes that will be carried in an accounting message. Among the ESM-specific attributes, **nat-port-range** is the only NAT-related attribute that will carry NAT logging information for the subscriber.

```
configure
  subscriber-mgmt
    radius-accounting-policy "acct" create
    no queue-instance-accounting
    host-accounting interim-update
    update-interval 5
    include-radius-attribute
        circuit-id
        framed-ip-addr
        mac-address
        nas-identifier
        nas-port-id
        nat-port-range
        remote-id
        sla-profile
        sub-profile
        subscriber-id
        user-name
        alc-acct-triggered-reason
    exit
    session-id-format number
    radius-server-policy "aaa"
  exit
```

Appendix B – Logging

Appendix B shows the RADIUS logging stream for all four hosts used in this chapter:

```

root@ linux:/var/log/freeradius/radacct/192.168.114.2# more detail-20170814
Mon Aug 14 09:14:45 2017
  Acct-Status-Type = Start
  NAS-IP-Address = 192.0.2.2
  Framed-IP-Address = 10.10.1.1
  NAS-Identifier = "BNG1"
  Acct-Session-Id = "D896FF0000002E59916D5E"
  Acct-Multi-Session-Id = "D896FF0000002F59916D5E"
  Event-Timestamp = "Aug 14 2017 02:29:02 PDT"
  NAS-Port-Type = Ethernet
  NAS-Port-Id = "1/1/5:1.1"
  ADSL-Agent-Circuit-Id = "sub-1"
  ADSL-Agent-Remote-Id = "host-1"
  Alc-Subsc-ID-Str = "sub-1"
  Alc-Subsc-Prof-Str = "sub-prof-1"
  Alc-SLA-Prof-Str = "sla-1"
  Alc-Nat-Port-Range = "172.16.3.3 1024-3023 router base l2-aware-base"
  Alc-Nat-Port-Range = "172.16.1.1 1024-3023 router 1 l2-aware-vrf1"
  Alc-Nat-Port-Range = "172.16.2.2 1024-3023 router 2 l2-aware-vrf2"
  Alc-Client-Hardware-Addr = "00:00:65:01:01:01"
  Acct-Unique-Session-Id = "fe5e55d7102d3f81"
  Timestamp = 1502727285
  Request-Authenticator = Verified

Mon Aug 14 09:14:53 2017
  Acct-Status-Type = Start
  NAS-IP-Address = 192.0.2.2
  Service-Type = Framed-User
  Framed-Protocol = PPP
  Framed-IP-Address = 10.10.1.2
  NAS-Identifier = "BNG1"
  Acct-Session-Id = "D896FF0000003159916D66"
  Acct-Multi-Session-Id = "D896FF0000002F59916D5E"
  Event-Timestamp = "Aug 14 2017 02:29:10 PDT"
  NAS-Port-Type = PPPoEoQinQ
  NAS-Port-Id = "1/1/5:1.1"
  ADSL-Agent-Circuit-Id = "sub-1"
  ADSL-Agent-Remote-Id = "host-2"
  Alc-Subsc-ID-Str = "sub-1"
  Alc-Subsc-Prof-Str = "sub-prof-1"
  Alc-SLA-Prof-Str = "sla-1"
  Alc-Nat-Port-Range = "172.16.3.3 1024-3023 router base l2-aware-base"
  Alc-Nat-Port-Range = "172.16.1.1 1024-3023 router 1 l2-aware-vrf1"
  Alc-Nat-Port-Range = "172.16.2.2 1024-3023 router 2 l2-aware-vrf2"
  Alc-Client-Hardware-Addr = "00:00:65:01:01:02"
  Acct-Unique-Session-Id = "cc1701cf565fcd5"
  Timestamp = 1502727293
  Request-Authenticator = Verified

Mon Aug 14 09:15:02 2017
  Acct-Status-Type = Start
  NAS-IP-Address = 192.0.2.2
  Framed-IP-Address = 10.10.1.1
  NAS-Identifier = "BNG1"
  Acct-Session-Id = "D896FF0000003259916D70"
  Acct-Multi-Session-Id = "D896FF0000003359916D70"
  Event-Timestamp = "Aug 14 2017 02:29:20 PDT"
  NAS-Port-Type = Ethernet
  NAS-Port-Id = "1/1/5:1.2"

```

```
ADSL-Agent-Circuit-Id = "sub-2"
ADSL-Agent-Remote-Id = "host-1"
Alc-Subsc-ID-Str = "sub-2"
Alc-Subsc-Prof-Str = "sub-prof-1"
Alc-SLA-Prof-Str = "sla-1"
Alc-Nat-Port-Range = "172.16.3.3 3024-5023 router base l2-aware-base"
Alc-Nat-Port-Range = "172.16.1.1 3024-5023 router 1 l2-aware-vrf1"
Alc-Nat-Port-Range = "172.16.2.2 3024-5023 router 2 l2-aware-vrf2"
Alc-Client-Hardware-Addr = "00:00:65:01:02:01"
Acct-Unique-Session-Id = "57659d99038f7c0d"
Timestamp = 1502727302
Request-Authenticator = Verified

Mon Aug 14 09:15:09 2017
Acct-Status-Type = Start
NAS-IP-Address = 192.0.2.2
Service-Type = Framed-User
Framed-Protocol = PPP
Framed-IP-Address = 10.10.1.2
NAS-Identifier = "BNG1"
Acct-Session-Id = "D896FF0000003559916D76"
Acct-Multi-Session-Id = "D896FF0000003359916D70"
Event-Timestamp = "Aug 14 2017 02:29:26 PDT"
NAS-Port-Type = PPPoEoQinQ
NAS-Port-Id = "1/1/5:1.2"
ADSL-Agent-Circuit-Id = "sub-2"
ADSL-Agent-Remote-Id = "host-2"
Alc-Subsc-ID-Str = "sub-2"
Alc-Subsc-Prof-Str = "sub-prof-1"
Alc-SLA-Prof-Str = "sla-1"
Alc-Nat-Port-Range = "172.16.3.3 3024-5023 router base l2-aware-base"
Alc-Nat-Port-Range = "172.16.1.1 3024-5023 router 1 l2-aware-vrf1"
Alc-Nat-Port-Range = "172.16.2.2 3024-5023 router 2 l2-aware-vrf2"
Alc-Client-Hardware-Addr = "00:00:65:01:02:02"
Acct-Unique-Session-Id = "433924b3304e2d9e"
Timestamp = 1502727309
Request-Authenticator = Verified

Mon Aug 14 09:19:42 2017
Acct-Status-Type = Interim-Update
NAS-IP-Address = 192.0.2.2
Service-Type = Framed-User
Framed-Protocol = PPP
Framed-IP-Address = 10.10.1.2
NAS-Identifier = "BNG1"
Acct-Session-Id = "D896FF0000003159916D66"
Acct-Session-Time = 289
Acct-Multi-Session-Id = "D896FF0000002F59916D5E"
Event-Timestamp = "Aug 14 2017 02:33:59 PDT"
NAS-Port-Type = PPPoEoQinQ
NAS-Port-Id = "1/1/5:1.1"
ADSL-Agent-Circuit-Id = "sub-1"
ADSL-Agent-Remote-Id = "host-2"
Alc-Subsc-ID-Str = "sub-1"
Alc-Subsc-Prof-Str = "sub-prof-1"
Alc-SLA-Prof-Str = "sla-1"
Alc-Nat-Port-Range = "172.16.3.3 1024-3023 router base l2-aware-base"
Alc-Nat-Port-Range = "172.16.1.1 1024-3023 router 1 l2-aware-vrf1"
Alc-Nat-Port-Range = "172.16.2.2 1024-3023 router 2 l2-aware-vrf2"
Alc-Client-Hardware-Addr = "00:00:65:01:01:02"
Alc-Acct-Triggered-Reason = regular
Alc-Acct-I-Inprof-Octets-64 = 0x00010000000000000000
Alc-Acct-I-Outprof-Octets-64 = 0x000100000000000009024
Alc-Acct-I-Inprof-Pkts-64 = 0x00010000000000000000
```



```
Alc-Acct-I-Outprof-Pkts-64 = 0x000100000000000000171
Alc-Acct-0-Inprof-Octets-64 = 0x000100000000000000000
Alc-Acct-0-Outprof-Octets-64 = 0x000100000000000000000
Alc-Acct-0-Inprof-Pkts-64 = 0x000100000000000000000
Alc-Acct-0-Outprof-Pkts-64 = 0x000100000000000000000
Acct-Unique-Session-Id = "cc1701cf565fcdf5"
Timestamp = 1502727582
Request-Authenticator = Verified
```

Mon Aug 14 09:20:01 2017

```
Acct-Status-Type = Interim-Update
NAS-IP-Address = 192.0.2.2
Service-Type = Framed-User
Framed-Protocol = PPP
Framed-IP-Address = 10.10.1.2
NAS-Identifier = "BNG1"
Acct-Session-Id = "D896FF0000003559916D76"
Acct-Session-Time = 292
Acct-Multi-Session-Id = "D896FF0000003559916D70"
Event-Timestamp = "Aug 14 2017 02:34:18 PDT"
NAS-Port-Type = PPPoEoQinQ
NAS-Port-Id = "1/1/5:1.2"
ADSL-Agent-Circuit-Id = "sub-2"
ADSL-Agent-Remote-Id = "host-2"
Alc-Subsc-ID-Str = "sub-2"
Alc-Subsc-Prof-Str = "sub-prof-1"
Alc-SLA-Prof-Str = "sla-1"
Alc-Nat-Port-Range = "172.16.3.3 3024-5023 router base l2-aware-base"
Alc-Nat-Port-Range = "172.16.1.1 3024-5023 router 1 l2-aware-vrf1"
Alc-Nat-Port-Range = "172.16.2.2 3024-5023 router 2 l2-aware-vrf2"
Alc-Client-Hardware-Addr = "00:00:65:01:02:02"
Alc-Acct-Triggered-Reason = regular
Alc-Acct-I-Inprof-Octets-64 = 0x000100000000000000000
Alc-Acct-I-Outprof-Octets-64 = 0x00010000000000000000a5a0
Alc-Acct-I-Inprof-Pkts-64 = 0x000100000000000000000
Alc-Acct-I-Outprof-Pkts-64 = 0x000100000000000000001a8
Alc-Acct-0-Inprof-Octets-64 = 0x000100000000000000000
Alc-Acct-0-Outprof-Octets-64 = 0x000100000000000000000
Alc-Acct-0-Inprof-Pkts-64 = 0x000100000000000000000
Alc-Acct-0-Outprof-Pkts-64 = 0x000100000000000000000
Acct-Unique-Session-Id = "433924b3304e2d9e"
Timestamp = 1502727601
Request-Authenticator = Verified
```

Mon Aug 14 09:20:11 2017

```
Acct-Status-Type = Interim-Update
NAS-IP-Address = 192.0.2.2
Framed-IP-Address = 10.10.1.1
NAS-Identifier = "BNG1"
Acct-Session-Id = "D896FF0000002E59916D5E"
Acct-Session-Time = 326
Acct-Multi-Session-Id = "D896FF0000002F59916D5E"
Event-Timestamp = "Aug 14 2017 02:34:28 PDT"
NAS-Port-Type = Ethernet
NAS-Port-Id = "1/1/5:1.1"
ADSL-Agent-Circuit-Id = "sub-1"
ADSL-Agent-Remote-Id = "host-1"
Alc-Subsc-ID-Str = "sub-1"
Alc-Subsc-Prof-Str = "sub-prof-1"
Alc-SLA-Prof-Str = "sla-1"
Alc-Nat-Port-Range = "172.16.3.3 1024-3023 router base l2-aware-base"
Alc-Nat-Port-Range = "172.16.1.1 1024-3023 router 1 l2-aware-vrf1"
Alc-Nat-Port-Range = "172.16.2.2 1024-3023 router 2 l2-aware-vrf2"
Alc-Client-Hardware-Addr = "00:00:65:01:01:01"
```

```
Alc-Acct-Triggered-Reason = regular
Alc-Acct-I-Inprof-Octets-64 = 0x00010000000000000000
Alc-Acct-I-Outprof-Octets-64 = 0x00010000000000000b220
Alc-Acct-I-Inprof-Pkts-64 = 0x00010000000000000000
Alc-Acct-I-Outprof-Pkts-64 = 0x0001000000000000001c8
Alc-Acct-0-Inprof-Octets-64 = 0x00010000000000000000
Alc-Acct-0-Outprof-Octets-64 = 0x00010000000000000000
Alc-Acct-0-Inprof-Pkts-64 = 0x00010000000000000000
Alc-Acct-0-Outprof-Pkts-64 = 0x00010000000000000000
Acct-Unique-Session-Id = "fe5e55d7102d3f81"
Timestamp = 1502727611
Request-Authenticator = Verified

Mon Aug 14 09:20:16 2017
Acct-Status-Type = Interim-Update
NAS-IP-Address = 192.0.2.2
Framed-IP-Address = 10.10.1.1
NAS-Identifier = "BNG1"
Acct-Session-Id = "D896FF0000003259916D70"
Acct-Session-Time = 313
Acct-Multi-Session-Id = "D896FF0000003359916D70"
Event-Timestamp = "Aug 14 2017 02:34:33 PDT"
NAS-Port-Type = Ethernet
NAS-Port-Id = "1/1/5:1.2"
ADSL-Agent-Circuit-Id = "sub-2"
ADSL-Agent-Remote-Id = "host-1"
Alc-Subsc-ID-Str = "sub-2"
Alc-Subsc-Prof-Str = "sub-prof-1"
Alc-SLA-Prof-Str = "sla-1"
Alc-Nat-Port-Range = "172.16.3.3 3024-5023 router base l2-aware-base"
Alc-Nat-Port-Range = "172.16.1.1 3024-5023 router 1 l2-aware-vrf1"
Alc-Nat-Port-Range = "172.16.2.2 3024-5023 router 2 l2-aware-vrf2"
Alc-Client-Hardware-Addr = "00:00:65:01:02:01"
Alc-Acct-Triggered-Reason = regular
Alc-Acct-I-Inprof-Octets-64 = 0x00010000000000000000
Alc-Acct-I-Outprof-Octets-64 = 0x00010000000000000b734
Alc-Acct-I-Inprof-Pkts-64 = 0x00010000000000000000
Alc-Acct-I-Outprof-Pkts-64 = 0x0001000000000000001d5
Alc-Acct-0-Inprof-Octets-64 = 0x00010000000000000000
Alc-Acct-0-Outprof-Octets-64 = 0x00010000000000000000
Alc-Acct-0-Inprof-Pkts-64 = 0x00010000000000000000
Alc-Acct-0-Outprof-Pkts-64 = 0x00010000000000000000
Acct-Unique-Session-Id = "57659d99038f7c0d"
Timestamp = 1502727616
Request-Authenticator = Verified

Mon Aug 14 09:30:34 2017
Acct-Status-Type = Stop
NAS-IP-Address = 192.0.2.2
Framed-IP-Address = 10.10.1.1
NAS-Identifier = "BNG1"
Acct-Session-Id = "D896FF0000002E59916D5E"
Acct-Session-Time = 949
Acct-Terminate-Cause = User-Request
Acct-Multi-Session-Id = "D896FF0000002F59916D5E"
Event-Timestamp = "Aug 14 2017 02:44:51 PDT"
NAS-Port-Type = Ethernet
NAS-Port-Id = "1/1/5:1.1"
ADSL-Agent-Circuit-Id = "sub-1"
ADSL-Agent-Remote-Id = "host-1"
Alc-Subsc-ID-Str = "sub-1"
Alc-Subsc-Prof-Str = "sub-prof-1"
Alc-SLA-Prof-Str = "sla-1"
```

```
Alc-Nat-Port-Range = "172.16.3.3 1024-3023 router base l2-aware-base"  
Alc-Nat-Port-Range = "172.16.1.1 1024-3023 router 1 l2-aware-vrf1"  
Alc-Nat-Port-Range = "172.16.2.2 1024-3023 router 2 l2-aware-vrf2"  
Alc-Client-Hardware-Addr = "00:00:65:01:01:01"  
Alc-Acct-I-Inprof-Octets-64 = 0x00010000000000000000  
Alc-Acct-I-Outprof-Octets-64 = 0x000100000000000037974  
Alc-Acct-I-Inprof-Pkts-64 = 0x00010000000000000000  
Alc-Acct-I-Outprof-Pkts-64 = 0x0001000000000000008e5  
Alc-Acct-0-Inprof-Octets-64 = 0x00010000000000000000  
Alc-Acct-0-Outprof-Octets-64 = 0x00010000000000000000  
Alc-Acct-0-Inprof-Pkts-64 = 0x00010000000000000000  
Alc-Acct-0-Outprof-Pkts-64 = 0x00010000000000000000  
Acct-Unique-Session-Id = "fe5e55d7102d3f81"  
Timestamp = 1502728234  
Request-Authenticator = Verified  
  
Mon Aug 14 09:30:34 2017  
Acct-Status-Type = Stop  
NAS-IP-Address = 192.0.2.2  
Framed-IP-Address = 10.10.1.1  
NAS-Identifier = "BNG1"  
Acct-Session-Id = "D896FF0000003259916D70"  
Acct-Session-Time = 931  
Acct-Terminate-Cause = User-Request  
Acct-Multi-Session-Id = "D896FF0000003359916D70"  
Event-Timestamp = "Aug 14 2017 02:44:51 PDT"  
NAS-Port-Type = Ethernet  
NAS-Port-Id = "1/1/5:1.2"  
ADSL-Agent-Circuit-Id = "sub-2"  
ADSL-Agent-Remote-Id = "host-1"  
Alc-Subsc-ID-Str = "sub-2"  
Alc-Subsc-Prof-Str = "sub-prof-1"  
Alc-SLA-Prof-Str = "sla-1"  
Alc-Nat-Port-Range = "172.16.3.3 3024-5023 router base l2-aware-base"  
Alc-Nat-Port-Range = "172.16.1.1 3024-5023 router 1 l2-aware-vrf1"  
Alc-Nat-Port-Range = "172.16.2.2 3024-5023 router 2 l2-aware-vrf2"  
Alc-Client-Hardware-Addr = "00:00:65:01:02:01"  
Alc-Acct-I-Inprof-Octets-64 = 0x00010000000000000000  
Alc-Acct-I-Outprof-Octets-64 = 0x0001000000000000378ac  
Alc-Acct-I-Inprof-Pkts-64 = 0x00010000000000000000  
Alc-Acct-I-Outprof-Pkts-64 = 0x0001000000000000008e3  
Alc-Acct-0-Inprof-Octets-64 = 0x00010000000000000000  
Alc-Acct-0-Outprof-Octets-64 = 0x00010000000000000000  
Alc-Acct-0-Inprof-Pkts-64 = 0x00010000000000000000  
Alc-Acct-0-Outprof-Pkts-64 = 0x00010000000000000000  
Acct-Unique-Session-Id = "57659d99038f7c0d"  
Timestamp = 1502728234  
Request-Authenticator = Verified  
  
Mon Aug 14 09:30:47 2017  
Acct-Status-Type = Stop  
NAS-IP-Address = 192.0.2.2  
Service-Type = Framed-User  
Framed-Protocol = PPP  
Framed-IP-Address = 10.10.1.2  
NAS-Identifier = "BNG1"  
Acct-Session-Id = "D896FF0000003159916D66"  
Acct-Session-Time = 954  
Acct-Terminate-Cause = User-Request  
Acct-Multi-Session-Id = "D896FF0000002F59916D5E"  
Event-Timestamp = "Aug 14 2017 02:45:04 PDT"  
NAS-Port-Type = PPPoEoQinQ  
NAS-Port-Id = "1/1/5:1.1"  
ADSL-Agent-Circuit-Id = "sub-1"
```

```
ADSL-Agent-Remote-Id = "host-2"  
Alc-Subsc-ID-Str = "sub-1"  
Alc-Subsc-Prof-Str = "sub-prof-1"  
Alc-SLA-Prof-Str = "sla-1"  
Alc-Nat-Port-Range = "172.16.3.3 1024-3023 router base l2-aware-base"  
Alc-Nat-Port-Range = "172.16.1.1 1024-3023 router 1 l2-aware-vrf1"  
Alc-Nat-Port-Range = "172.16.2.2 1024-3023 router 2 l2-aware-vrf2"  
Alc-Client-Hardware-Addr = "00:00:65:01:01:02"  
Alc-Acct-I-Inprof-Octets-64 = 0x00010000000000000000  
Alc-Acct-I-Outprof-Octets-64 = 0x000100000000000037974  
Alc-Acct-I-Inprof-Pkts-64 = 0x00010000000000000000  
Alc-Acct-I-Outprof-Pkts-64 = 0x000100000000000000008e5  
Alc-Acct-0-Inprof-Octets-64 = 0x00010000000000000000  
Alc-Acct-0-Outprof-Octets-64 = 0x00010000000000000000  
Alc-Acct-0-Inprof-Pkts-64 = 0x00010000000000000000  
Alc-Acct-0-Outprof-Pkts-64 = 0x00010000000000000000  
Acct-Unique-Session-Id = "cc1701cf565fcd5"  
Timestamp = 1502728247  
Request-Authenticator = Verified  
  
Mon Aug 14 09:30:47 2017  
Acct-Status-Type = Stop  
NAS-IP-Address = 192.0.2.2  
Service-Type = Framed-User  
Framed-Protocol = PPP  
Framed-IP-Address = 10.10.1.2  
NAS-Identifier = "BNG1"  
Acct-Session-Id = "D896FF0000003559916D76"  
Acct-Session-Time = 938  
Acct-Terminate-Cause = User-Request  
Acct-Multi-Session-Id = "D896FF0000003359916D70"  
Event-Timestamp = "Aug 14 2017 02:45:04 PDT"  
NAS-Port-Type = PPPoEoQinQ  
NAS-Port-Id = "1/1/5:1.2"  
ADSL-Agent-Circuit-Id = "sub-2"  
ADSL-Agent-Remote-Id = "host-2"  
Alc-Subsc-ID-Str = "sub-2"  
Alc-Subsc-Prof-Str = "sub-prof-1"  
Alc-SLA-Prof-Str = "sla-1"  
Alc-Nat-Port-Range = "172.16.3.3 3024-5023 router base l2-aware-base"  
Alc-Nat-Port-Range = "172.16.1.1 3024-5023 router 1 l2-aware-vrf1"  
Alc-Nat-Port-Range = "172.16.2.2 3024-5023 router 2 l2-aware-vrf2"  
Alc-Client-Hardware-Addr = "00:00:65:01:02:02"  
Alc-Acct-I-Inprof-Octets-64 = 0x00010000000000000000  
Alc-Acct-I-Outprof-Octets-64 = 0x0001000000000000378ac  
Alc-Acct-I-Inprof-Pkts-64 = 0x00010000000000000000  
Alc-Acct-I-Outprof-Pkts-64 = 0x000100000000000000008e3  
Alc-Acct-0-Inprof-Octets-64 = 0x00010000000000000000  
Alc-Acct-0-Outprof-Octets-64 = 0x00010000000000000000  
Alc-Acct-0-Inprof-Pkts-64 = 0x00010000000000000000  
Alc-Acct-0-Outprof-Pkts-64 = 0x00010000000000000000  
Acct-Unique-Session-Id = "433924b3304e2d9e"  
Timestamp = 1502728247  
Request-Authenticator = Verified  
  
root@linux:/var/log/freeradius/radacct/192.168.114.2#
```

L2TP Network Server

This chapter provides information about L2TP network servers (LNS).

Topics in this chapter include:

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

Applicability

Initially, this chapter was initially written based on SR OS Release 11.0.R7, but the CLI in this edition is based on SR OS Release 16.0.R7.

Overview

The Layer 2 Tunneling Protocol (L2TP) is a tunneling protocol commonly used to transport PPP sessions from an initiator known as an L2TP Access Concentrator (LAC) to an L2TP Network Server (LNS). L2TP is typically used for wholesaling residential broadband services. In this scenario, the LAC resides in the wholesaler's network and has a Layer-2 connection to an access concentrator such as a DSLAM. The LAC responds during the discovery phase (if PPPoE is used) and during PPP Link Control Protocol (LCP) negotiation. The LAC also performs an initial authentication of the subscriber. A successful authentication, typically from RADIUS, indicates to the LAC that PPP frames from this subscriber should be tunneled to an LNS at the indicated IP address. The LAC then tunnels the PPP frames from this subscriber over an L2TP tunnel to the LNS, where the PPP session is actually terminated. In other words, PPP **sessions** require L2TP **tunnels** for the sessions to get carried over.

L2TP uses two types of messages; control messages and data messages. Control messages are used in the establishment, maintenance, and tearing down of tunnels and sessions. In order to provide extensibility and maximize interoperability, the payloads of control messages are encoded using Attribute Value Pairs (AVPs), some of which are applicable to all control messages, and some of which are specific to particular control messages. The L2TP header contains sequence number fields that must be present in control messages to allow for a reliable L2TP control channel that guarantees delivery. Data messages are used to encapsulate PPP frames being carried over the tunnel. Data messages are not retransmitted if packet loss occurs.


L2TP has a common fixed header format for both control and data messages, and a Type (T) bit in the header is used to indicate whether the packet is a control (1) or data (0) message. The L2TP packet is then carried in a transport protocol, and although the specifications allow L2TP to be directly encapsulated over Frame Relay, ATM, and UDP/IP, the latter is used almost exclusively.

The objective of this chapter is to provide a generic overview of how to configure the 7750 SR to support the LNS and LTS (L2TP Tunnel Switching) functions.

Example Topology

The example topology used through this chapter is shown in [Figure 76: Example Topology](#). Both the LAC and the LNS participate in IS-IS and LDP, together with PE-1. All three devices form part of AS 64496 and peer using iBGP for the VPN-IPv4 address family. None of these protocols is mandatory for supporting LNS functionality; L2TP packets can ingress the system over any network interface as native IP or encapsulated as IP in MPLS, or through an IES/VPRN IP interface (SAP) as native IP. The MPLS data-plane within the example topology is chosen purely because of its simplicity and flexibility. Tester T1 simulates a DSLAM and one or more PPP clients, and is connected directly to the LAC. Although the LAC in this topology is a 7750 SR router, the configuration requirements of that device are beyond the scope of this example. Tester T2 provides a traffic source/sink capability and is connected directly to PE-1.

Figure 76: Example Topology



Hardware Configuration

To support LNS (and LTS) functionality, at least one MS-ISA card is required, which must be configured as MDA type isa-bb, and must be housed in the carrier IOM. The MS-ISA performs L2TP data-plane encapsulation and de-capsulation, whereas the subscriber processing (Enhanced Subscriber Management or ESM) for PPP sessions is implemented within the carrier IOM.

```

configure
  card 1
    card-type iom3-xp
    mda 1
      mda-type m4-10gb-xp-xfp
      no shutdown
    exit
    mda 2
      mda-type isa-bb
      no shutdown
    exit
  no shutdown
exit
exit

```

The MS-ISA is then configured to become a member of an **lns-group**. Up to six MS-ISAs can be configured to belong to one or more lns-groups. When two or more MS-ISAs belong to the same lns-group, by default PPP sessions are load-balanced over those MS-ISAs on a per-session basis.

```

configure
  isa
    lns-group 1 create
      mda 1/2
      no shutdown
    exit

```

```
    exit
exit
```

Configuration

ESM Base Configuration

For completeness, the following outputs contain the base ESM configuration that is applied to subscribers instantiated at the LNS throughout this chapter. Deviations from these base parameters are mentioned explicitly.

The SLA-Profile and Sub-Profile configurations have a minimal set of parameters. The SLA-Profile uses the default ingress/egress QoS policy of 1, while the **no qos-marking-from-sap** command ensures that any subsequent marking is inherited from the egress QoS policy referenced in the SLA-profile, and not taken from the egress SAP. In order to do on-line accounting through RADIUS, the Sub-Profile also calls the relevant RADIUS accounting policy. Finally, the **sub-ident-policy** is configured with **use-direct-map-as-default** for the **sub-profile-map** and **sla-profile-map**, which means that the strings passed from RADIUS in the **Alc-Subs-Prof-Str** and **Alc-SLA-Prof-Str** Vendor Specific Attributes (VSAs) are interpreted verbatim so they are not used as string input to a mapping function.

```
configure
  subscriber-mgmt
    sla-profile "sla-profile-1" create
      egress
        no qos-marking-from-sap
      exit
    exit
    sub-profile "sub-profile-1" create
      radius-accounting
        policy "sm-acct-1"
      exit
    exit
    sub-ident-policy "all-subscribers" create
      sub-profile-map
        use-direct-map-as-default
      exit
      sla-profile-map
        use-direct-map-as-default
      exit
    exit
  exit
exit
exit
```

Whilst it is entirely possible to authenticate subscribers locally using a local user database (LUDB), the more widely deployed approach is to use RADIUS, and this approach is therefore used throughout this chapter. The next output shows the authentication policy *sm-auth-1* and the RADIUS accounting policy *sm-acct-1*. Both policies reference the radius-server-policy *aaa-radius-1*, which provides the context to configure the source-address to use for RADIUS messages and an associated routing context. The RADIUS server policy then references a RADIUS server *radius-172.16.1.11*, which in turn allows for configuration of the server IP address, the secret key to be used for message exchanges, and any other optional port configuration. In this example also Change of Authorization (COA) is configured at RADIUS server level, through the **accept-coa** command.

The intention is not to provide a complete description of all of the RADIUS parameters as this would distract from the objective of this chapter.

```
configure
  router
    radius-server
      server "radius-172.16.1.11" address 172.16.1.11 secret vsecret1 create
      accept-coa
      pending-requests-limit 1024
    exit
  exit
exit
aaa
  radius-server-policy "aaa-radius-1" create
  servers
    router "Base"
    source-address 192.0.2.2
    server 1 name "radius-172.16.1.11"
  exit
exit
subscriber-mgmt
  authentication-policy "sm-auth-1" create
  pppoe-access-method pap-chap
  include-radius-attribute
    nas-port-id
    nas-identifier
    access-loop-options
    calling-station-id remote-id
  exit
  radius-server-policy "aaa-radius-1"
exit
radius-accounting-policy "sm-acct-1" create
no queue-instance-accounting
session-accounting interim-update host-update
update-interval 120
include-radius-attribute
  circuit-id
  framed-ip-addr
  nas-identifier
  nas-port-id
  nas-port-type
  sla-profile
  sub-profile
  subscriber-id
  std-acct-attributes
exit
session-id-format number
radius-server-policy "aaa-radius-1"
exit
exit
exit
```

Basic LNS Configuration

To illustrate the building blocks that are required to implement LNS functionality, a VPRN is used on the LAC and the LNS supporting an L2TP tunnel and terminating PPP sessions at the LNS side of the L2TP

tunnel. The required configuration for this VPRN at the LNS is shown in the following output. The unicast VPRN parameters such as **route-distinguisher** and **vrf-import/vrf-export** are not discussed here, only the parameters that are relevant to subscriber termination, which are equally applicable to VPRN and/or IES services.

The interface called *system* is a logical loopback interface and is used as the LNS endpoint address for L2TP signaling. The name of this interface is not important; this interface only must be a loopback interface. The LAC has a corresponding interface with IP address 192.168.0.1. The address of the interface *system* is also used in the **subscriber-interface** context as argument to the **unnumbered** command, meaning this IP address is used for the purpose of IPCP negotiation with incoming PPP sessions. Within the **subscriber-interface** context, the **group-interface** has a different definition than a conventional ESM group interface. A conventional group interface contains one or more SAPs belonging to the same port or LAG. However, in the context of LNS, there are no SAPs. The group interface also might terminate sessions within the same L2TP tunnel which are anchored on different MS-ISAs in a common *Ins-group*. To accommodate this, the **group-interface** has the creation-time attribute **Ins**. This attribute essentially means that the group interface can terminate subscribers from more than one port/LAG; where port/LAG is interpreted as different MS-ISAs.

The **group-interface** then provides a **sap-parameters** context that allows for configuration of **sub-sla-mgmt** parameters that would typically be found under a SAP. These parameters apply to all subscribers terminated on this group interface. In the example shown, only the **sub-ident-policy** is configured; meaning that other ESM parameters such as **sla-profile**, **sub-profile**, and **subscriber-id** must be obtained from a different source. In this chapter, they are obtained through RADIUS.

The static route black-holes prefix 10.48.127.0/24 ensures this prefix is added to the route-table. Subscribers are allocated /32 addresses from this range, which must be advertised upstream to PE-1 to ensure end-to-end IP connectivity. This is implemented through the **vrf-export policy** (not shown for conciseness).

Within the **l2tp** context, an hierarchy of groups and tunnels is defined. Groups reside directly under the **l2tp** context, and tunnels reside within the **group** context. Groups are intended to administratively organize tunnels that may share a common use or contain common parameters. The L2TP tunnel characteristics can be inherited from the group context, or overridden within the **tunnel** context. In the **group** context shown in the following output, the **Ins-group 1** command refers to the **Ins-group** previously configured at the ISA level. This is followed by the **local-address** command that defines the IP address to be used as a source address for L2TP signaling. The **ppp** context then defines the characteristics to be used when PPP sessions are established. In this case, the authentication mechanism is CHAP, and the previously configured RADIUS **authentication-policy** is used to authenticate the user. During the PPP session setup, the LAC negotiates LCP and authentication parameters with the subscriber. Two AVPs, the **Proxy LCP** AVP and the **Proxy Authentication** AVP allow this information to be forwarded by the LAC to the LNS. This information can be accepted by the LNS, allowing PPP to continue with negotiation of IPCP, or it can be rejected, in which case the LNS initiates a new round of NCP and PPP authentication. The **proxy-authentication** and **proxy-lcp** commands allow the information contained in these AVPs to be accepted.

Finally, the **tunnel** context provides the context for defining L2TP tunnel specific parameters. The **peer** command defines the far-end (LAC) IP address to which L2TP messages are addressed. The **password** is used to authenticate the far-end tunnel initiator, and is used in conjunction with the **challenge** parameter to implement a CHAP-like authentication mechanism. The default behavior is to never challenge the initiator (LAC); the **challenge always** command is the reverse of this behavior. The **remote-name** is used to provide an additional level of security. When the Start Control Connection Request (SCCRQ) is received from the LAC initiating the tunnel setup, it carries a mandatory **Host Name** AVP. The value of this AVP is compared with the name configured in the **remote-name**, and only tunnels with matching names are accepted. In a similar way, the **local-name** parameter is used to populate the Host Name AVP sent by the LNS in the SCCRQ, and can be used as a similar security feature at the LAC.

When two or more MS-ISAs belong to the same lns-group, PPP sessions are load-balanced over those MS-ISAs on a per-session basis by default. Although it is not shown in the following configuration example, it is worth mentioning that within each L2TP group context, an option exists to load-balance the sessions on a per-L2TP tunnel basis using the **load-balance-method per-tunnel** command. This can be useful, for example, when multiple sessions are received from a single subscriber (for example, MLPPP member links) which must be handled within the same MS-ISA.

```
configure
  service
    vprn 1 customer 1 create
      vrf-import "vrf1-import"
      vrf-export "vrf1-export"
      route-distinguisher 64496:1
      auto-bind-tunnel
        resolution-filter
          ldp
        exit
      resolution filter
    exit
  interface "system" create
    address 192.168.0.2/32
    loopback
  exit
  subscriber-interface "LNS-SUB-INT" create
    unnumbered 192.168.0.2
    group-interface "LNS-GROUP-INT" lns create
      sap-parameters
        sub-sla-mgmt
          sub-ident-policy "all-subscribers"
        exit
      exit
    exit
  exit
  static-route-entry 10.48.127.0/24
    black-hole
    no shutdown
  exit
  exit
  l2tp
    group "L2TP-GROUP-1" create
      hello-interval 60
      idle-timeout 600
      lns-group 1
      local-address 192.168.0.2
      ppp
        authentication chap
        authentication-policy "sm-auth-1"
        default-group-interface "LNS-GROUP-INT" service-id 1
        keepalive 10 hold-up-multiplier 3
        proxy-authentication
        proxy-lcp
      exit
      tunnel "L2TP-TUNNEL-1" create
        challenge always
        local-name "LNS"
        peer 192.168.0.1
        remote-name "LAC"
        password tunnelpwd
        no shutdown
      exit
    no shutdown
```

```

        exit
        no shutdown
    exit
    no shutdown
exit
exit
exit

```

As previously described, RADIUS is used to authenticate the subscriber, which upon successful authentication returns the ESM parameters, Subscriber-ID (**Alc-Subsc-ID-Str**), SLA-Profile (**Alc-SLA-Prof-Str**), and Sub-Profile (**Alc-Subsc-Prof-Str**) as needed for instantiating the subscriber in SR OS. These parameters could be obtained locally on the LNS using the **def-sub-id**, **def-sla-profile** and **def-sub-profile** commands under the **group-interface>sap-parameters** context. This enables a mechanism to provide default parameters in the absence of another source. However, passing them from RADIUS has some benefits, such as:

- It is comparatively easy to provide different SLA- and Sub-Profiles to different users, which can be used to differentiate service levels.
- If RADIUS infrastructure is available and used to provide ESM parameters, it is relatively easy to extend that infrastructure to provide for mid-session changes of those parameters (such as **sla-profile** and **sub-profile**) using a Change of Authorization (CoA).

The following provides an example of a RADIUS users file entry for the test subscriber. In addition to the previously defined ESM parameters, the Alc-Serv-ID VSA is used to define the service number in which this subscriber must be terminated (in this case, VPRN 1 as previously configured), while the Alc-Interface VSA is used to define the relevant group interface within that service. If it is intended that all PPP sessions ingressing on a particular L2TP group are all to be terminated within a common service and group interface, it is not necessary for the RADIUS server to send the Alc-Serv-ID and Alc-Interface VSAs defining the service and group interface. Instead, a default service and group interface can be defined within the **ppp** context of the l2tp group using the parameter **default-group-interface <name> service-id <number>**. The remainder of the attributes in the output are well-known standard attributes.

```

subscriber1@isp.net      Cleartext-Password := "letmein"
                        Alc-Subsc-ID-Str = "subscriber1@isp.net",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-Serv-Id = "1",
                        Alc-Interface = "LNS-GROUP-INT",
                        Service-Type = Framed-User,
                        Framed-Protocol = PPP,
                        Framed-IP-Address = 10.48.127.27,

```

L2TP Tunnel Setup

Before the PPP session can be terminated at the LNS, an L2TP tunnel must be established between the LAC and LNS. This is achieved using a three-way control message exchange of Start Control Connection Request (SCCRQ), Start Control Connection Reply (SCCRP), and Start Control Connection Connected (SCCN). All of these messages are explicitly acknowledged by the peer using the sequence numbers (number sent, number received) in the L2TP header, thereby creating a reliable control channel. The acknowledgment can be piggybacked in a corresponding control message, or it can be an explicit acknowledgment using a control packet with only an L2TP header, known as a Zero Length Body (ZLB) message.

The SCCRQ is used to initiate the tunnel between LAC and LNS, and although it can be sent by either the LAC or LNS, it is typically sent by the LAC toward the LNS (as in this example). The SCCRQ contains a number of mandatory AVPs, denoted by the M-bit in the AVP header (set to 1), including Message Type, Protocol Version, Host Name, Framing Capabilities, and Assigned Tunnel ID. It can also contain a number of optional AVPs, such as Vendor Name, and Firmware Revision, which can be ignored by the recipient if they are unrecognized.

```
1 2019/05/23 16:27:13.662 CEST MINOR: DEBUG #2001 vprn1 L2TP(v2, ctrl, ingress)
"L2TP(v2, ctrl, ingress): UDP 192.168.0.1:1701 -> 192.168.0.2:1701
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"
  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
    4096
  AVP VendorName(0,8), flags:, reserved=0
    "Nokia"
  AVP AssignedTunnelId(0,9), flags: mandatory, reserved=0
    10007"
```

The SCCRP is sent in response to the SCCRQ and is used to indicate that the parameters in the SCCRQ were acceptable and that the establishment of the L2TP tunnel can continue. The SCCRP contains the same mandatory AVPs and can contain the same optional AVPs as the SCCRQ, but an additional optional AVP is the Challenge AVP which is included as a result of the **challenge always** and **password** parameters configured within the **tunnel** context.

```
2 2019/05/23 16:27:13.662 CEST MINOR: DEBUG #2001 vprn1 L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.168.0.2:1701 -> 192.168.0.1:1701
tunnel 10007 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"
  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
    4096
  AVP VendorName(0,8), flags:, reserved=0
    "Nokia"
  AVP AssignedTunnelId(0,9), flags: mandatory, reserved=0
    9265
  AVP Challenge(0,11), flags: mandatory, reserved=0
    0d 12 69 86 37 7b 35 7f d7 00 22 bd ca 25 ad ab
    59 d7 33 d5 a1 01 92 4f 22 a9 9a d8 b3 68 55 60
    98 96 8c f9 e4 0a 9c ce e3 b9 ed 48 d8 "
```

The response to the SCCRQ, and the completion of the three-way message exchange is the SCCN. The only mandatory AVP for the SCCN is the Message Type, and since the SCCRQ contained a Challenge AVP, the SCCN also contains an AVP Challenge Response. If this response is not satisfactory to the LNS, it generates a Stop Control Connection Notification (StopCCN) with a result code indicating that the requester is not authorized, and subsequently removes any associated tunnel state.

```
3 2019/05/23 16:27:13.664 CEST MINOR: DEBUG #2001 vprn1 L2TP(v2, ctrl, ingress)
"L2TP(v2, ctrl, ingress): UDP 192.168.0.1:1701 -> 192.168.0.2:1701
tunnel 9265 session 0, ns 1 nr 1, flags:, reserved=0
  AVP MessageType(0,0), flags: mandatory, reserved=0
    StartControlConnectionConnected(3)
  AVP ChallengeResponse(0,13), flags: mandatory, reserved=0
    e6 10 7b 11 36 7d ba fc e6 14 4d 30 3d f0 9a ba "
```

With a successful three-way exchange completed, the L2TP tunnel is established. A snapshot view of all L2TP tunnels within the relevant routing context can be displayed using the command in the following output. The Loc-Tu-ID and Rem-Tu-ID are the local and remote tunnel IDs passed in the Assigned Tunnel Id AVP in the SCCRQ and SCCRQ respectively. The Conn ID, or connection Id, is a locally significant parameter used for identifying a particular tunnel, and is a 32-bit representation of the local tunnel Id (1442 * 65536 = 946470912). The connection ID is for example used in event log entries for this tunnel. If the state is shown as *established*, then one or more PPP sessions are running over the tunnel. The state can also be *establishedIdle* meaning that although the tunnel is up and established, there are no PPP sessions active within the tunnel.

```
*A:LNS# show router 1 l2tp tunnel
=====
Conn ID   Loc-Tu-ID Rem-Tu-ID State           Blacklist-state   Ses Active
Group                                          Ses Total
Assignment
-----
607191040 9265      10007     establishedIdle  not-blacklisted  0
L2TP-GROUP-1
L2TP-TUNNEL-1
-----
No. of tunnels: 1
=====
*A:LNS#
```

Once a tunnel is established, maintenance and health-checking of that tunnel is achieved using a keepalive mechanism that employs Hello control messages. The Hello message contains only one AVP, the **Message Type** AVP, which indicates it is a Hello message. The Hello messages operate asynchronously between the peers. There is no echo request and echo response function, but simply a Hello followed by an acknowledgment. The Hello is acknowledged in the same way as other control messages, using either piggybacking or ZLB acknowledgments. This asynchronous behavior allows for either end of the tunnel to be configured for different Hello intervals (they are not negotiated), or even for one end not send Hellos at all.

```
5 2019/05/23 16:28:11.374 CEST MINOR: DEBUG #2001 vprn1 L2TP(v2, ctrl, ingress)
"L2TP(v2, ctrl, ingress): UDP 192.168.0.1:1701 -> 192.168.0.2:1701
tunnel 9265 session 0, ns 2 nr 1, flags:, reserved=0
  AVP MessageType(0,0), flags: mandatory, reserved=0
    Hello(6)"

6 2019/05/23 16:28:11.374 CEST MINOR: DEBUG #2001 vprn1 L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.168.0.2:1701 -> 192.168.0.1:1701
tunnel 10007 session 0, ns 1 nr 3, flags:, reserved=0"
```

The Hello interval at the LNS is configurable under the l2tp, group, or tunnel contexts using the **hello-interval** parameter. The range is 60 to 3600 seconds, with the default being 60 seconds. The **hello-interval infinite** option suppresses sending of Hellos. If the system sends a Hello message and does not get an acknowledgment, it will retransmit the Hello message as many times as the value of the **max-retries-estab** parameter, each time with an exponential back-off. The **max-retries-estab** parameter can be configured in the l2tp, group, or tunnel contexts. The default value is 5, and if no acknowledgment is received before this value is exceeded, the tunnel is declared down and a StopCCN is sent toward the peer.

The retry interval starts with 1 second and doubles on each retry with a maximum-interval of 8 seconds. For example, using a max-retries-estab value of 7 results in a retry of [1, 2, 4, 8, 8, 8 seconds]

The StopCCN is a message that can be generated by either LAC or LNS and is used to inform its peer that the tunnel is being closed. This implicitly means that all PPP sessions carried within that tunnel are also being closed without any associated control messages for those sessions. The StopCCN must contain the **Message Type** and **Tunnel ID** AVPs, and additionally carries a **Result Code** AVP with result code and error code fields to indicate to the peer the reason for the tunnel closure.

```

9 2017/06/07 15:19:00.56 CEST MINOR: DEBUG #2001 vprn1 L2TP(v2, ctrl, ingress)
"L2TP(v2, ctrl, ingress): UDP 192.168.0.1:1701 -> 192.168.0.2:1701
tunnel 14442 session 0, ns 2 nr 3, 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
    11458"

10 2017/06/07 15:19:00.56 CEST MINOR: DEBUG #2001 vprn1 L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.168.0.2:1701 -> 192.168.0.1:1701
tunnel 11458 session 0, ns 3 nr 3, flags:, reserved=0"

```

The tunnel **Connection Id** can be used as an additional argument to display the details of a particular tunnel when multiple tunnels are present. The following output is an example of this taken just after the L2TP tunnel has been closed by the LAC peer, and is intentionally taken at this time to illustrate the purpose of some of the fields shown in the output. The State is moved to *closedByPeer*, and the StopCCN Result field and Error Message field respectively contain the result code and error code of the Result Code AVP received from the LAC in the StopCCN. Because the tunnel is now in a closedByPeer state, all state and information related to this tunnel is removed from the system after a period defined by the Destruct Timeout (shown in the output as Destruct TO). The intention of the Destruct Timeout is to retain information about the tunnel closure which might aid operational communities. The default value as shown is 60 seconds, but it can be configured using the **destruct-timeout** parameter in the l2tp, group, or tunnel contexts. The remainder of the fields in the output are the operational parameters of the tunnel and are self-explanatory.

```
*A:LNS# show router 1 l2tp tunnel detail
```

```

=====
L2TP Tunnel 607191040
=====

Connection ID: 607191040
Protocol      : v2
State        : closedByPeer
IP           : 192.168.0.2
UDP          : 1701

```

```

Peer IP      : 192.168.0.1
Peer UDP     : 1701
Tx dst-IP    : 192.168.0.1
Tx dst-UDP   : 1701
Rx src-IP    : 192.168.0.1
Rx src-UDP   : 1701
Name         : LNS
Remote Name  : LAC
Assignment ID: L2TP-TUNNEL-1
Group Name   : L2TP-GROUP-1
Acct. Policy : N/A
Error Message: operator request

Tunnel ID    : 9265
Preference   : 50
Hello Interval (s): 300
Idle TO (s)  : 600
Max Retr Estab : 5
Cfg'd Sess Limit : unlimited
Transport Type : udpIp
Time Started  : 05/23/2019 16:27:14
Time Established : 05/23/2019 16:27:14
Stop CCN Result : generalReq
Blacklist-state : not-blacklisted
Set Dont Fragment : true

Remote Conn ID : 655818752
Remote Tunnel ID : 10007
Receive Window : 64
AVP Hiding     : never
Destruct TO (s) : 60
Max Retr Not Estab : 5
Oper Session Limit: 32767
Challenge      : always
Time Idle      : N/A
Time Closed    : 05/23/2019 16:30:15
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:LNS#

```

PPP Session Setup

Once the L2TP tunnel is created, the process of establishing a PPP session can start. Once again, a three-way control message exchange is used for establishing a session within an L2TP tunnel, consisting of the Incoming Call Request (ICRQ), Incoming Call Reply (ICRP), and Incoming Call Connected (ICCN). Given that they are control messages, they are all explicitly acknowledged using piggybacking or ZLB acknowledgments.

The ICRQ is sent from the LAC to the LNS to indicate that it has received an incoming call (PPP session) and that a session needs to be established between the two peers for this call. The ICRQ provides enough information about the call for the LNS to make a decision about whether it should answer the call or not. The ICRQ contains the Message Type and Assigned Session ID AVPs as well as a Call Serial Number AVP, which can be used on both the LAC and LNS as an easy reference to the call for troubleshooting purposes. The ICRQ can also carry optional AVPs including Calling Number and Access Line Information

AVPs (RFC 5515) such as Circuit ID, Remote ID, Actual Data Rate Upstream, and Actual Data Rate Downstream.

```
5 2019/05/23 16:38:49.730 CEST MINOR: DEBUG #2001 vprn1 L2TP(v2, ctrl, ingress)
"L2TP(v2, ctrl, ingress): UDP 192.168.0.1:1701 -> 192.168.0.2:1701
tunnel 12229 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
    15859
  AVP CallSerialNumber(0,15), flags: mandatory, reserved=0
    15550
  AVP CallingNumber(0,22), flags: mandatory, reserved=0
    "LAC 1/1/3:2"
  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 ICRP is sent by the LNS toward the LAC in response to the ICRQ to indicate that the parameters in the ICRQ were acceptable, and that the LAC should go ahead and proceed with the call. The ICRP contains only two AVPs; the Message Type and the Assigned Session ID. The Assigned Session ID values are local to each peer as opposed to a negotiated or agreed-upon value.

```
7 2019/05/23 16:38:49.731 CEST MINOR: DEBUG #2001 vprn1 L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.168.0.2:1701 -> 192.168.0.1:1701
tunnel 8826 session 15859, 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
    14338"
```

The final message in the three-way exchange used for establishing sessions within the tunnel is the ICCN. It is sent by the LAC to the LNS to indicate that the call has been answered, so the L2TP session is moved to the *established* state. It also provides additional information on parameters that were used to answer the call which may not have been available when the ICRQ was sent (although it is likely that in most cases they were available). At a minimum, the ICCN must contain the Message Type, Framing Type, and TX Connect Speed AVPs. The TX Connect Speed defines the speed in bits per second from the perspective of traffic flowing from the LAC toward the subscriber (i.e. the LAC downstream rate) and, for best accuracy, can be derived by the LAC from the PPP Broadband Forum Access Line Characteristic tags inserted by the access node (Appendix C TR-101). The TX Connect Speed can be useful for indirect setting of a Hierarchical QoS (H-QoS) aggregate rate. It is indirect because the LNS cannot infer and set an aggregate rate based directly on the TX Connect Speed AVP, but rather the TX Connect Speed is passed to the RADIUS server (using the **include-radius-attribute access-loop-option** parameter in the authentication policy), which in turn may pass the aggregate rate to the LNS in a QoS override VSA. This is described further in the QoS section.

A number of optional AVPs can also be present providing information from the LCP negotiation between the LAC and client. These include Initial Receive, Last Transmit and Last Receive LCP Config Requests, together with Proxy Authentication Type, Name, Challenge, and Response. These parameters allow the LNS to either force a renegotiation of LCP, or to continue with the PPP session and move to the IPCP phase. The final AVP present in the ICCN shown is the RX Connect Speed AVP, which is the opposite of

the TX Connect Speed and defines the speed in bits per second from the perspective of traffic flowing from the subscriber toward the LAC.

```

9 2019/05/23 16:38:49.732 CEST MINOR: DEBUG #2001 vprn1 L2TP(v2, ctrl, ingress)
"L2TP(v2, ctrl, ingress): UDP 192.168.0.1:1701 -> 192.168.0.2:1701
tunnel 12229 session 14338, 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 3c 32 9f a6
    [1] MRU: 1492
    [3] Authentication-Protocol: 0xc223 (CHAP), Algorithm = 5 (MD5)
    [5] Magic-Number: 0x3c329fa6
  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
    "subscriber1@isp.net"
  AVP ProxyAuthenChallenge(0,31), flags:, reserved=0
    d2 cb 6f 10 90 61 f5 bf 59 72 e0 d2 4b 8c c5 05
    02 5c c5 9e d3 c5 65 d9 f0 08 38 71 d0 a2 75 23
    6e 64 b3 7f 45 a9 3f 9e 9b a0 f0 e6 ac
  AVP ProxyAuthenId(0,32), flags:, reserved=0
    id=1, reserved=0
  AVP ProxyAuthenResponse(0,33), flags:, reserved=0
    d5 3e 86 2f 2b 50 0d 0b 01 15 5b 0f 6d ec aa fb
  AVP RxConnectSpeed(0,38), flags:, reserved=0
    4294967295"

```

On completion of the three-way control message exchange required for session setup, the LNS authenticates the user in the incoming call. In this example, RADIUS is used, which returns the standard and vendor-specific attributes previously defined in the users file. A successful authentication allows the LNS to move to the IPCP phase with the subscriber. In this example, RADIUS returns the IP address in the standard attribute Framed-IP-Address, but equally local pooling with a DHCP server could be used. For conciseness, the IPCP phase is not detailed within this example because the process is reasonably well-known and understood. However, on completion of the IPCP phase, the subscriber is instantiated and the L2TP session becomes active. The Tunnel-ID and Session-ID parameters are locally generated numbers that are passed in L2TP control messages. As previously described, the Connection Id is a locally significant parameter that is a 32-bit representation of the local tunnel Id ($6482 * 65536 = 424804352$). The ID field is again a locally significant parameter used to identify the L2TP session, and is again represented as a 32-bit number. It is derived from a sum of the Control Connection ID plus the Session ID ($424804352 + 26255 = 424830607$).

```
*A:LNS# show router 1 l2tp session
```

```
=====
L2TP Session Summary
=====
```

| ID | Control Conn ID | Tunnel-ID | Session-ID | State |
|-----------|-----------------|-----------|------------|-------------|
| 801454082 | 801439744 | 12229 | 14338 | established |

```
subscriber1@isp.net
interface: LNS-GROUP-INT
service-id: 1
10.48.127.27
-----
No. of sessions: 1
=====
*A:LNS#
```

The PPP session is also recorded in the subscriber-host table of VPRN 1 and a forwarding state of *Fwding* indicates that all attributes and resources associated with this subscriber are correctly installed and activated within the system. The subscriber username is shown, as is its MAC address and IP address. The IP address has an origin of IPCP. The fact that a MAC address is displayed here is somewhat misleading because this is a PPP over L2TP session, which does not have a MAC address present in any of its headers. When the MS-ISA removes the L2TP header, it converts the PPP packet to PPPoE for ease of subsequent processing. As a result of this, the MS-ISA generates a dummy MAC address, and this is the MAC address shown. The displayed SAP 1/2/lns-esm:1.259 is automatically generated by the system. Each operational MS-ISA that is part of the lns-group creates two internal objects, known as lns-net and lns-esm. These objects are the network-side (lns-net) and the subscriber-side (lns-esm) of each MS-ISA.

When the first L2TP session within this service is established, the system creates one lns-esm SAP where the first two digits indicate the MDA slot (1/2) where the MS-ISA is installed, and the last two numbers are the internal Q-in-Q tags used through the MS-ISA (1.259). This internal Q-in-Q tag value is of little relevance, but for informational purposes is derived from the group interface If index. If there is more than one MS-ISA active in the lns-group, a second session would be load-balanced onto this MS-ISA, and a second lns-esm SAP would be created, until a maximum of six SAPs is reached, which corresponds to the maximum number of supported active MS-ISA boards.

```
*A:LNS# show service id 1 subscriber-hosts
=====
Subscriber Host table
=====
Sap          Subscriber
IP Address
MAC Address  PPPoE-SID Origin  Fwding State
-----
[1/2/lns-esm:1.259]  subscriber1@isp.net
10.48.127.27
00:00:2f:c5:38:02  1          IPCP          Fwding
-----
Number of subscriber hosts : 1
=====
*A:LNS#
```

It is also possible to view the internal lns-net object, shown in the next output as interface name *_tmnx_lns-in-1/2* with port number *1/2/lns-net:1**. To further clarify (and reiterate), the lns-esm and lns-net are simply internal objects used to route L2TP traffic through the MS-ISA board. Upstream traffic (subscriber to LNS) ingresses through lns-net into the MS-ISA where the L2TP header is decapsulated before PPP packets are presented to the service group interface through lns-esm. Downstream traffic (LNS to subscriber) passes through lns-esm into the MS-ISA where the PPP packets are encapsulated in L2TP before egressing through lns-net and being routed toward the destination.

```
*A:LNS# show service id 1 interface
=====
Interface Table
=====
Interface-Name          Adm      Opr(v4/v6)  Type      Port/SapId
IP-Address              PfxState
```

```

-----
system          Up      Up/Down  VPRN  loopback
 192.168.0.2/32 n/a
LNS-SUB-INT     Up      Up/Down  VPRN S* subscriber
  Unnumbered If[192.168.0.2] n/a
LNS-GROUP-INT  Up      Up/Down  VPRN G* bbg-5.lns-esm
 _tmnx_lns-in-1/2 Up      Up/Down  VPRN  1/2/lns-net:1*
 -
-----
Interfaces : 4
=====

```

Wholesale/Retail

In the example configuration used so far, the L2TP tunnel/session is terminated in VPRN 1, and the subscriber is also terminated in the same VPRN 1. However, a common requirement is to build per-customer VRFs (VPRNs), particularly for business users. To meet this requirement, the so-called 'Wholesale/Retail' model is used, which provides a mechanism to terminate the subscriber in a different service context from the service which actually terminated the L2TP tunnel/session.

To achieve this, a second service is created which becomes the 'Retail VRF', or customer-specific VRF, and the previously defined VPRN 1 becomes the Wholesale VRF (which actually requires no further configuration). The necessary configuration for the Retail VRF is as follows and its parameters have been previously explained. Although they may seem obvious, there are a couple of points that are worth revisiting. The **vrf-import** and **vrf-export** parameters are used to reference policies to import/export VPN-IPv4/v6 prefixes with the customer-specific Route-Target Extended Communities. Given that a different routing context and unique Route-Targets are used for this Retail VRF, it is perfectly feasible to re-use the same IP address in VPRN 2 as was used in VPRN 1 for the unnumbered subscriber interface. The group interface has a different name from the group interface in VPRN 1, but this is simply for illustration purposes and both group interfaces can have the same name if a standard naming convention is required. More importantly, the group interface must have the creation-time attribute **lns** to allow subscriber termination without SAPs. The static route black-holes prefix 10.10.148.0/24, ensuring this prefix is added to the route-table. This IP address range is used to allocate addresses to subscribers, and is therefore advertised in VPN-IPv4.

```

configure
  service
    vprn 2 customer 1 create
      vrf-import "vrf2-import"
      vrf-export "vrf2-export"
      route-distinguisher 64496:2
      auto-bind-tunnel
        resolution-filter
          ldp
        exit
      resolution filter
    exit
  interface "loopback" create
    address 192.168.0.2/32
    loopback
  exit
  subscriber-interface "VPRN2-SUB-INT" create
    unnumbered 192.168.0.2
    group-interface "VPRN2-GROUP-INT" lns create
      sap-parameters
        sub-sla-mgmt
        sub-ident-policy "all-subscribers"

```

```

                exit
            exit
        exit
    static-route-entry 10.10.48.0/24
        black-hole
        no shutdown
    exit
    exit
    no shutdown
exit
exit
exit

```

In addition to the Retail VRF configuration, the RADIUS entry for the subscriber returns **Alc-Serv-Id** VSA with a value of 2 to indicate the Retail VRF Service Id, while the **Alc-Interface** VSA refers to the group interface name within that Retail VRF.

```

subscriber2@isp.net  Cleartext-Password := "letmein"
                    Alc-Subsc-ID-Str = "subscriber2@isp.net",
                    Alc-Subsc-Prof-Str = "sub-profile-1",
                    Alc-SLA-Prof-Str = "sla-profile-1",
                    Alc-Serv-Id = "2",
                    Alc-Interface = "VPRN2-GROUP-INT",
                    Service-Type = Framed-User,
                    Framed-Protocol = PPP,
                    Framed-IP-Address = 10.10.148.22

```

In this Wholesale/Retail scenario, the high-level functions are as follows:

- The L2TP tunnel and session are terminated in the Wholesale VRF (in this example, VPRN 1).
- When the LNS receives the ICCN for the session, it authenticates the user (in this example via RADIUS).

RADIUS returns the Retail VRF service Id and group interface. If RADIUS returns IP address information this address is used for the purpose of IPCP negotiation with the subscriber within the Retail VRF (in this example, VPRN 2). If RADIUS does not return IP address information, it can be derived from either of the following:

- A DHCP client function within the group interface, which is used to obtain an IP address from a local or remote DHCP server.
- The local-address-assignment feature, which directly accesses a local DHCP server through an internal procedure call (the server pool name must be obtained through RADIUS, LUDB, or default-pool-name).

Once the subscriber is activated, the PPP session and subscriber-host can be seen in VPRN 2. The description field of the **show service id 2 ppp session** command is however somewhat misleading. It is automatically concatenated from the VPRN that terminated the L2TP tunnel, the tunnel Connection Id, the local tunnel Id, and the L2TP session Id. It should not be misinterpreted as meaning that the subscriber has been terminated in VPRN 1.

```

*A:LNS# show service id 2 ppp session

=====
PPP sessions for service 2
=====
User-Name
  Descr.
      Up Time      Type  Termination      IP/L2TP-Id/Interface-Id MC-Stdby
-----

```

```
subscriber2@isp.net
vprn:1 connid:449518148 tid:6859 sid:6724
      0d 00:00:10 oL2tp local          10.10.148.22
-----
No. of PPP sessions: 1
=====
*A:LNS#
```

QoS

In the preceding examples, the subscriber PPP sessions terminated by the LNS have been instantiated using the default SAP-ingress/egress QoS policies (policy 1), with a single queue and no use of H-QoS. This section demonstrates the use of slightly more complex QoS policies that employ H-QoS, with the intention of providing an overview of those capabilities.

For subscriber termination in broadband networks, it is fairly commonplace to use one or more policers on ingress, and not apply an aggregate rate limit on ingress (upstream) traffic. Whilst this is possible in SR OS for general ESM subscriber termination, policers are not supported when the system is functioning as an LNS. It is therefore necessary to use one or more queues on ingress with the usual considerations with regard to the use of service-queuing or shared-queuing. Conversely, on egress (downstream) it is common to see more than one queue in use for different services, particularly for business services, with an aggregate rate applied to the subscriber through the use of H-QoS. For example, assume that there are three classes in use; Best-Effort (BE), Assured-Forwarding (AF), and Expedited Forwarding (EF). This section will look at two ways to achieve this. Firstly using a conventional H-QoS scheduler, and secondly using an egress Port-Scheduler.

The SAP-ingress QoS policy classifies traffic into three Forwarding Classes (FCs) and maps those FCs to a single queue. Ingress traffic is not rate-limited (default PIR in queue 1 is max), and queue 1 is not mapped to a parent H-QoS scheduler.

```
configure
qos
  sap-ingress 10 create
  queue 1 create
  exit
  queue 11 multipoint create
  exit
  fc "af" create
  queue 1
  exit
  fc "be" create
  queue 1
  exit
  fc "ef" create
  queue 1
  exit
  dscp be fc "be"
  dscp ef fc "ef"
  dscp af31 fc "af"
  exit
exit
exit
exit
```

A scheduler policy is created having a single a tier 1 scheduler with a rate-limit of 8Mb/s.

```
configure
qos
```

```
    scheduler-policy "Subscriber-Aggregate-Policy" create
      tier 1
        scheduler "Aggregate-Rate" create
          rate 8000
        exit
      exit
    exit
  exit
exit
```

The SAP-egress QoS policy performs egress classification and maps classified traffic to the relevant FC, which in turn is mapped to its own queue. All queues are mapped to the previously configured tier 1 scheduler *Aggregate-Rate* such that queue 3 (EF) is allocated bandwidth first, and queue 1 (BE) and 2 (AF) are allocated bandwidth next in a 1:4 ratio.

```
configure
  qos
    sap-egress 10 create
      queue 1 create
        parent "Aggregate-Rate" level 2 weight 20
      exit
      queue 2 best-effort create
        parent "Aggregate-Rate" level 2 weight 80
      exit
      queue 3 expedite create
        parent "Aggregate-Rate" cir-level 3
        rate 1024 cir 1024
      exit
      fc af create
        queue 2
      exit
      fc be create
        queue 1
      exit
      fc ef create
        queue 3
      exit
      dscp be fc "be"
      dscp ef fc "ef"
      dscp af31 fc "af"
    exit
  exit
exit
```

To this point, the QoS configuration is no different from a typical SAP-level QoS application. To make it applicable to ESM, the previously configured SAP-ingress and SAP-egress QoS policies must be referenced in the ingress/egress contexts of the sla-profile, respectively. Equally, the H-QoS scheduler-policy must be referenced in the ingress/egress contexts of the sub-profile. In this example, H-QoS is only used on egress, and as a result the scheduler-policy is referenced only in the egress context.

```
configure
  subscriber-mgmt
    sla-profile "ESM-SLA-PROF" create
      ingress
        qos 10
      exit
    exit
    egress
      qos 10
    exit
  no qos-marking-from-sap
```

```

        exit
    exit
    sub-profile "ESM-SUB-PROF" create
        collect-stats
        radius-accounting
            policy "AAA-ACCT-POLICY"
        exit
    egress
        scheduler-policy "Subscriber-Aggregate-Policy"
    exit
    exit
    exit
    exit
    exit
    exit

```

The queues assigned to the subscriber through the preceding SAP-ingress/egress QoS policies, together with accumulative statistics can be viewed using the **show service active-subscribers subscriber <name> detail** command (real time rates can be seen using the **monitor** command). The H-QoS scheduler hierarchy, with the SAP-egress queues mapped as child queues to a parent scheduler can be validated using the command **show qos scheduler-hierarchy subscriber <name> egress**. The **detail** argument as an extension of this command provides a significant amount of detail on real-time bandwidth allocated to each queue by the scheduler in the within-CIR and above-CIR passes. It also provides a useful snapshot on offered traffic loads in Kb/s on a per-queue basis.

```

*A:LNS# show qos scheduler-hierarchy subscriber "subscriber2@isp.net" egress
=====
Scheduler Hierarchy - Subscriber subscriber2@isp.net
=====
Egress Scheduler Policy : Subscriber-Aggregate-Policy
-----
Root (Egr)
| slot(1)
|--(S) : Aggregate-Rate (Port 1/2/lns-esm Orphan)
|
|   |--(Q) : Sub=subscriber2@isp.net:sla-profile-1 2->1/2/lns-esm:1.263->3
|   |--(Q) : Sub=subscriber2@isp.net:sla-profile-1 2->1/2/lns-esm:1.263->2
|   |--(Q) : Sub=subscriber2@isp.net:sla-profile-1 2->1/2/lns-esm:1.263->1
|
|
=====
*A:LNS#

```

The advantage of using conventional H-QoS schedulers is that they can be applied universally on ingress and egress to provide a subscriber aggregate rate capability. The disadvantage of this approach is that the aggregate rate defined in the scheduler-policy (or overridden in the sub-profile) cannot be dynamically overridden from RADIUS using the QoS-override VSA (**Alc-Subscriber-QoS-Override**). If ingress H-QoS is not a requirement, but the ability to override the subscriber egress aggregate-rate is, then H-QoS should be implemented using an egress port-scheduler.

The egress port-scheduler is functionally the same as a conventional H-QoS scheduler in the manner with which it arbitrates bandwidth across its child queues. However, it has some notable differences:

- It is applied at the egress port level. Any queue that uses that egress port to which it is applied that is not explicitly mapped to a port-scheduler is considered an orphan queue. Orphan queues are not serviced by the port-scheduler until all of its child queues have been serviced.

- Unlike conventional H-QoS schedulers that include only Ethernet overhead, the port-scheduler includes Preamble and Inter-Frame Gap for each packet.
- It is supported only on Ethernet ports, and only on egress.
- The egress aggregate rate applied to the subscriber can be overridden from RADIUS.

The first bullet point above is significant from an LNS perspective. In general, after ESM handling, downstream traffic for subscribers egresses the system over a physical port. This is not the case for L2TP subscribers, for which the traffic is passed through to the MS-ISA for L2TP encapsulation before egressing the LNS (and in fact could egress the system on any number of physical ports). It is therefore not possible to apply the port-scheduler policy to the egress port in the conventional manner, and what is needed is a mechanism to apply the port-scheduler policy to the logical internal ports that interface to the MS-ISA. To achieve this, an intermediate object known as a **port-policy** is used, which, when configured, references the **port-scheduler** policy, and which subsequently is applied to the relevant **lms-group**.

Create the port-scheduler-policy.

```
configure
  qos
    port-scheduler-policy "egress-port-scheduler" create
  exit
exit
exit
```

Create the port-policy and reference the previously configured port-scheduler policy.

```
configure
  port-policy "isa-port-policy" create
    egress-scheduler-policy "egress-port-scheduler"
  exit
exit
```

Attach the port-policy to the lms-group containing the MS-ISA.

```
configure
  isa
    lms-group 1 create
      shutdown
      port-policy "isa-port-policy"
      no shutdown
    exit
  exit
exit
```

Once the **port-scheduler** policy and **port-policy** are in place, the subscriber QoS can reference it. The QoS configuration previously used for conventional H-QoS schedulers differs in both the **sap-egress** policy and **sub-profile** when an egress **port-scheduler** is used. The queues within the **sap-egress** policy are each configured to be parented to the egress port-scheduler using the **port-parent** keyword (as opposed the **parent** keyword used for conventional H-QoS schedulers).

```
configure
  qos
    sap-egress 10 create
      queue 1 create
        port-parent level 2 weight 20
      exit
      queue 2 best-effort create
```



```

        port-parent level 2 weight 80
    exit
    queue 3 expedite create
        port-parent cir-level 3
        rate 1024 cir 1024
    exit
    fc af create
        queue 2
    exit
    fc be create
        queue 1
    exit
    fc ef create
        queue 3
    exit
    dscp be fc "be"
    dscp ef fc "ef"
    dscp af31 fc "af"
    exit
exit
exit
exit

```

The sub-profile contains no reference to scheduler policies, but instead contains a per-subscriber egress aggregate rate in Kb/s, defined through the **agg-rate-limit** parameter.

```

configure
  subscriber-mgmt
    sub-profile "sub-profile-1"
      egress
        no scheduler-policy
        agg-rate-limit 8000
      exit
    exit
  exit
exit
exit

```

Once again, the queues assigned to the subscriber through the preceding SAP-ingress/egress QoS policies, together with accumulative statistics, can be viewed using the **show service active-subscribers subscriber <name> detail** command (real time rates can be seen using the **monitor** command). The scheduler SAP-egress queues mapped as child queues to a port-scheduler can be validated using the **show qos scheduler-hierarchy subscriber <name> egress** command. The **detail** argument provides a significant amount of detail on bandwidth allocated to each queue by the scheduler in the within-CIR and above-CIR passes. It also provides a useful snapshot on offered traffic loads in Kb/s on a per-queue basis. Alternatively, all of the child queues and orphans mapped to the port-scheduler can be displayed using the **show qos scheduler-hierarchy port <slot/mda/lms-esm>** command, again with the optional **detail** argument.

```

*A:LNS# show qos scheduler-hierarchy subscriber "subscriber2@isp.net" egress
=====
Scheduler Hierarchy - Subscriber subscriber2@isp.net
=====
Egress Scheduler Policy :
-----
Root (Egr)
| slot(1)
|--(Q) : Sub=subscriber2@isp.net:sla-profile-1 2->1/2/lms-esm:1.263->3 (Port
| 1/2/lms-esm)
|
|--(Q) : Sub=subscriber2@isp.net:sla-profile-1 2->1/2/lms-esm:1.263->2 (Port

```

```
1/2/lms-esm)
|
|--(Q) : Sub=subscriber2@isp.net:sla-profile-1 2->1/2/lms-esm:1.263->1 (Port
1/2/lms-esm)
|
```

```
=====
*A:LNS#
```

With the previously configured QoS policies and schedulers available, the aggregate rate limit in use for the subscriber can be viewed using the **show service active-subscribers subscriber <name> detail** command. There are three fields in this output that are of interest here. The **E. Agg Rate Limit** field shows the configured rate-limit in the sub-profile and is therefore relatively static. The **RADIUS Rate-Limit** field shows the aggregate rate received by RADIUS using the **Alc-Subscriber-QoS-Override** VSA, which overrides any rate limit statically configured in the sub-profile. Finally, the **Oper-Rate-Limit** shows the static or RADIUS-received rate-limit, minus any other H-QoS adjustments, such as Multicast H-QoS adjustment (snooping on IGMP joins) or ANCP line-rate adjustments.

```
*A:LNS# show service active-subscribers subscriber "subscriber2@isp.net" detail
      | match expression " E. Agg Rate Limit|RADIUS Rate-Limit|Oper-Rate-Limit"
E. Sched. Policy : N/A                E. Agg Rate Limit: 8000
RADIUS Rate-Limit: N/A
Oper-Rate-Limit : 8000
Hs-Oper-Rate-Limit : Maximum
*A:LNS#
```

Overriding the **agg-rate-limit** defined in the sub-profile can be done as part of the RADIUS Access-Accept, or through a Change of Authorization (CoA), and as previously outlined uses the **Alc-Subscriber-QoS-Override** VSA. This override function can be used, for example, to reconcile the LNS aggregate rate with the subscriber downstream rate learned through the TxConnectSpeed AVP in the ICCN message from the LAC. This ensures that the LNS does not overwhelm any downstream access node, and ensures that the LNS is responsible for all QoS scheduling in the event of congestion. In the following example, an override of the aggregate rate to 10Mb/s is sent as a CoA.

```
63 2019/05/24 13:47:00.340 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
Change of Authorization(43) id 178 len 66 from 172.16.1.11:38256 vrid 1
SESSION ID [44] 22 020DFF000000225CE7D31B
VSA [26] 16 Nokia(6527)
SUBSCRIBER QOS OVERRIDE [126] 14 e:r:rate=10000
"

64 2019/05/24 13:47:00.341 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
Change of Authorization Ack(44) 172.16.1.11:38256 id 178 len 20 vrid 1
"
```

Re-issuing the **show service active-subscribers subscriber <name> detail** command after the CoA shows that the **RADIUS Rate-Limit** field and the **Oper-Rate-Limit** field both correctly show 10Mb/s.

```
*A:LNS# show service active-subscribers subscriber "subscriber2@isp.net" detail
      | match expression " E. Agg Rate Limit|RADIUS Rate-Limit|Oper-Rate-Limit"
E. Sched. Policy : N/A                E. Agg Rate Limit: 8000
RADIUS Rate-Limit: 10000
Oper-Rate-Limit : 10000
```


```
Hs-Oper-Rate-Limit : Maximum
*A:LNS#
```

Propagating QoS Markings to L2TP/MPLS Headers

It is often desirable to mark the L2TP header (DCSP) or MPLS header (EXP) based on the class of service that is carried in the encapsulated subscriber IP payload. In general, when a packet is classified and mapped to an FC on ingress, that FC value is carried in the internal switch fabric header and is present when the packet is egressing the node. In the case of L2TP traffic however, the operation becomes a little more complex due to the fact that traffic transits the MS-ISA board with subscriber QoS implemented 'mid-chassis'.

In the upstream direction (from subscriber to LNS), traffic arrives encapsulated in L2TP at the ingress IOM, and is passed through the MS-ISA via the internal object Ins-net. When traffic exits the MS-ISA through Ins-esm as native IP, it is subject to subscriber ingress QoS implemented on the (MS-ISA) carrier IOM. Traffic is classified and mapped to an FC at this point, and that FC mapping is maintained in the switch fabric header. As a result, marking of traffic is effected by the network egress QoS policy.

Figure 77: Ingress/Egress QoS Processing



In the downstream direction (from LNS to subscriber), traffic arrives at the ingress IOM as native IP and is diverted to the MS-ISA via the Ins-esm internal object. At the Ins-esm, the traffic is subject to subscriber egress QoS. When the traffic is passed through Ins-esm to the MS-ISA for L2TP encapsulation, internal Q-in-Q VLAN tags are attached as previously described. As the Ins-esm is effectively a SAP-egress, the internal switch fabric header containing the FC marking is removed at this point, and, as a result FC information is lost.

Therefore, in order to allow for FC-continuity through the MS-ISA, the system implements a queue-group at the ingress of Ins-net that has a dot1p to FC mapping as shown in [Figure 77: Ingress/Egress QoS Processing](#). Assuming a SAP-egress QoS policy that employs FCs BE, AF and EF, the QoS policy would include the additional configuration to implement the appropriate dot1p marking as shown in the following output. When traffic arrives at Ins-net, it is classified and mapped into the appropriate FCs, and

the associated FC mapping included in the switch fabric header. At network egress, the L2TP packet is then subject to marking as defined in the network egress QoS policy.

```
configure
  qos
    sap-egress 10 create
      queue 1 create
      exit
      queue 2 best-effort create
      exit
      queue 3 expedite create
      exit
      fc af create
        queue 2
        dot1p 2
      exit
      fc be create
        queue 1
        dot1p 0
      exit
      fc ef create
        queue 3
        dot1p 5
      exit
      dscp be fc "be"
      dscp ef fc "ef"
      dscp af31 fc "af"
    exit
  exit
exit
```

Framed-Route

The majority of residential services in broadband networks have a single registered 32-bit IPv4 address on the WAN side of the RG and a private (RFC 1918) network on the LAN side. Traffic from the LAN toward the BNG (and Internet) is thereafter subject to Network Address and Port Translation (NAPT). However, a common requirement for delivery of business services is the ability for the BNG to recognize one or more IP subnets on the LAN side of the RG that is not subject to NAT, and the subscriber prefix is a route to a network. This is achieved using the standard RADIUS **Framed-Route** attribute, or dynamic BGP peering. Both serve the function of allowing one or more subnets to be learned at the LNS with a next-hop IP address of the RG WAN.

To provide an example of the use of Framed-Route, the Retail VRF VPRN 2 is again used, and in fact requires no modification in order to support subscribers with Framed-Routes. In general ESM, where Framed-Route is used, there is a requirement to configure **anti-spoof type nh-mac**, but for LNS SAPs this is the default. The RADIUS users file is updated to also return a Framed-Route attribute for prefix 10.128.46.0/24 with a next-hop determined by the subscriber IP prefix. The prefix has a metric of 10, and has a tag of value 200, which may be used for example, for routing policy.

```
subscriber2@isp.net    Cleartext-Password := "letmein"
                      Alc-Subsc-ID-Str = "subscriber2@isp.net",
                      Alc-Subsc-Prof-Str = "sub-profile-1",
                      Alc-SLA-Prof-Str = "sla-profile-1",
                      Alc-Serv-Id = "2",
                      Alc-Interface = "VPRN2-GROUP-INT",
                      Service-Type = Framed-User,
                      Framed-Protocol = PPP,
```

```
Framed-IP-Address = 10.10.148.22,
Framed-Route = "10.128.46.0/24 0.0.0.0 10 tag 200",
```

In SR OS, a prefix learned through the Framed-Route attribute is known internally as a *Managed Route*. Once the subscriber is instantiated, the presence of the Managed Route can be verified as installed.

```
*A:LNS# show service id 2 ppp session detail | match "Managed Routes"
                                         post-lines 5
Managed Routes
-----
IP Address                               Status      Metric Tag      Pref
-----
10.128.46.0/24                           installed   10      200      0
-----
*A:LNS#
```

The Managed Route can also be seen present in the VPRN routing-table, learned through protocol *Managed*.

```
*A:LNS# show router 2 route-table protocol managed
=====
Route Table (Service: 2)
=====
Dest Prefix[Flags]                      Type   Proto   Age      Pref
  Next Hop[Interface Name]              Metric
-----
10.128.46.0/24                           Remote Managed 00h01m33s 0
  10.10.148.22                             10
-----
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
      B = BGP backup route available
      L = LFA nexthop available
      S = Sticky ECMP requested
=====
*A:LNS#
```

L2TP Tunnel Switching (LTS)

In general, L2TP tunnels are established directly between LAC and LNS. However, if there are a large number of LAC devices (and therefore a large number of L2TP tunnels), it may be desirable to perform some aggregation of these tunnels before presenting them to the LNS. This is implemented by one or more LNS devices performing the function of an L2TP Tunnel Switch (LTS). The LTS terminates multiple L2TP tunnels from the LAC(s), and sources a single L2TP tunnel toward the target LNS, switching L2TP sessions from one tunnel to another tunnel accordingly.

SR OS supports LTS functionality, and from a configuration perspective, it requires no more than L2TP being enabled (no shutdown) in the required routing context if the relevant attributes are returned from RADIUS. As with LNS functions, at least one MS-ISA is required to support LTS functions. In this example, VPRN 1 is used with the previously defined configuration. To recap, this VPRN has a single L2TP group "L2TP-GROUP-1", and within that group, a single tunnel defined "L2TP-TUNNEL-1" that terminates the tunnel from the LAC. To demonstrate LTS functionality, the LAC continues to function as a LAC, the LNS functions as an LTS, and PE-1 becomes the LNS.

The RADIUS users file for subscriber1@isp.net is modified to include a number of additional attributes and VSAs. The **Aic-Serv-Id** and **Aic-Interface** define the service ID and group interface where the subscriber

is terminated, and this can be any IES or VPRN service. The **Alc-Tunnel-Serv-Id** VSA identifies the service where the L2TP tunnel is initiated. It can be different from the service where the subscriber is terminated, but id does not need to. If it is a different service, then the minimum requirement is that L2TP is placed in a no shutdown state. In this example VPRN 1 terminates the subscriber, and also initiates the L2TP tunnel. The other attributes are standard attributes defined in RFC 2868 for L2TP tunnel setup. The Tunnel-Assignment-Id attribute is used to maintain the concept of groups and tunnels, where Tunnel-Assignment-Id:0 is used to indicate the group name and Tunnel-Assignment-Id:1 is used to indicate the tunnel name. This provides sufficient information for the LTS to initiate an L2TP tunnel without any additional nodal configuration.

```
subscriber1@isp.net      Cleartext-Password := "letmein"
                        Alc-Subsc-ID-Str = "subscriber1@isp.net",
                        Alc-Subsc-Prof-Str = "sla-profile-1",
                        Alc-SLA-Prof-Str = "sub-profile-1",
                        Alc-Serv-Id = "1",
                        Alc-Interface = "LNS-GROUP-INT",
                        Alc-Tunnel-Serv-Id = 1,
                        Tunnel-Assignment-Id:0 = "RADIUS-returned-Tunnel-Group",
                        Tunnel-Type:1 += L2TP,
                        Tunnel-Medium-Type:1 += IP,
                        Tunnel-Server-Endpoint:1 += 192.168.0.3,
                        Tunnel-Password:1 += "letmein",
                        Tunnel-Assignment-Id:1 += "RADIUS-returned-Tunnel-Name",
                        Tunnel-Client-Auth-Id = "LTS",
```

The LAC forwards the PPP session into the LAC to LTS tunnel, and after the LTS receives the ICCN from the LAC, it proceeds in authenticating the subscriber. RADIUS returns the preceding attributes with sufficient information for the LTS to instantiate the subscriber and initiates an L2TP tunnel/session with PE-1, the target LNS. The LNS then authenticates the user once more, this time providing it with IP address information through IPCP negotiation. This interaction between PPP client and LNS is transparent to the LTS, which is responsible for switching PPP packets between L2TP sessions. However, the user is instantiated in the system as a full-fledged subscriber.

```
*A:LNS# 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
-----
subscriber1@isp.net
  vprn:1 connid:345328060 tid:5269 sid:18876
      0d 00:01:42   oL2tp lac          905456423
-----
No. of PPP sessions: 1
=====
*A:LNS#
```

Within VPRN 1, two L2TP tunnels are active. The entry with Connection Id 937033728 belongs to group *RADIUS-returned-Tunnel-Group* (obtained from the Tunnel-Assignment-Id:0 RADIUS attribute) and has tunnel name *RADIUS-returned-Tunnel-Name* (obtained from the Tunnel-Assignment-Id:1 RADIUS attribute). This is the tunnel from LTS to LNS, and it is in the *established* state and has one session active. The entry with Connection Id 560005120 is the statically defined tunnel from the LAC, belonging to the

CLI-configured group L2TP-GROUP-1 with tunnel name L2TP-TUNNEL-1. This tunnel is also in the *established* state, with one session active.

```
*A:LNS# show router 1 l2tp tunnel
=====
Conn ID   Loc-Tu-ID Rem-Tu-ID State           Blacklist-state  Ses Active
  Group                                     Ses Total
  Assignment
-----
345309184 5269      12695     established     not-blacklisted  1
L2TP-GROUP-1
L2TP-TUNNEL-1
905445376 13816     12071     established     not-blacklisted  1
RADIUS-returned-Tunnel-Group
RADIUS-returned-Tunnel-Name
-----
No. of tunnels: 2
=====
*A:LNS#
```

Equally, within VPRN 1, two L2TP sessions are active for subscriber subscriber1@isp.net. Session 937037337 is carried in Tunnel-ID 14298, which, as shown in the previous output, is the tunnel toward the LNS, while session 560030762 is carried in Tunnel-ID 8545, which is the tunnel toward the LAC.

```
*A:LNS# show router 1 l2tp session
=====
L2TP Session Summary
=====
ID           Control Conn ID   Tunnel-ID   Session-ID   State
-----
345328060    345309184        5269       18876        established
subscriber1@isp.net
interface: LNS-GROUP-INT
service-id: 1
905456423    905445376        13816      11047        established
-----
No. of sessions: 2
=====
*A:LNS#
```

IPv6

The deployment of IPv6 into residential broadband networks dictates some design choices, or perhaps even some enforced IPv6 address allocation mechanisms:

- Bridged or Routed Residential Gateways (RGs).
- Numbered or unnumbered WAN.
- Stateful (DHCPv6) or stateless (Stateless Address Auto-Configuration, or SLAAC) address assignment.

The purpose of this example is not to show every possibility, but simply to demonstrate that enabling IPv6 is possible at the LNS, just as if this were a conventional BNG doing PPP Termination and Aggregation (PTA). This example uses a widely adopted approach of dual-stack Routed RG with DHCPv6 Prefix Delegation.

The configuration of VPRN 2 is modified to include some IPv6 parameters. In the **subscriber-interface** context, the **delegated-prefix-len** command is set to **variable** indicating that prefixes delegated to subscribers may be of varying length (the default delegated prefix length is /64). The **allow-unmatching-prefixes** command tells the subscriber interface to operate in an IPv6 unnumbered mode, allowing IPv6 addresses to be allocated to subscribers that do not fall within the range of any IPv6 subnet defined under the subscriber interface. Within the **group-interface** context, the **ipv6** context places router-advertisements into a no shutdown state and has the **managed-configuration** flag set indicating that stateful (DHCPv6) address configuration is to be used.

Also a **dhcp6 proxy-server** is enabled, providing an interworking function between the RADIUS server (where the Delegated Prefix is obtained from) and the DHCPv6 client. The proxy will take the RADIUS-provided prefix and responds to the clients Solicit message with an DHCPv6 Advertise message containing the delegated prefix (IA_PD). Because the DHCPv6 messages from the client need to be received over the subscriber PPP session, the proxy-server is configured to allow this using the **client-applications ppp** command. Finally, there is a static-route for black-holing the /48 IPv6 prefix. The client is allocated a /64 prefix from this range and this static-route is used to provide an aggregated upstream prefix advertisement.

```
configure
  service
    vprn 2
      subscriber-interface "VPRN2-SUB-INT" create
        ipv6
          default-dns 2001:db8:2c41::56
          delegated-prefix-len variable
          allow-unmatching-prefixes
        exit
      group-interface "VPRN2-GROUP-INT" lns create
        ipv6
          router-advertisements
            managed-configuration
            no shutdown
          exit
        dhcp6
          proxy-server
            client-applications ppp
            no shutdown
          exit
        exit
      exit
    exit
  static-route-entry 2a00:8010:1b00::/48
    black-hole
    no shutdown
  exit
  no shutdown
exit
exit
exit
```

The RADIUS users file entry for subscriber2@isp.net is also modified to return the IPv6 Delegated Prefix using the standard attribute **Delegated-IPv6-Prefix**.

```
subscriber2@isp.net    Cleartext-Password := "letmein"
                      Alc-Subsc-ID-Str = "subscriber2@isp.net",
                      Alc-Subsc-Prof-Str = "sub-profile-1",
                      Alc-SLA-Prof-Str = "sla-profile-1",
                      Alc-Serv-Id = "2",
```



```
Alc-Interface = "VPRN2-GROUP-INT",
Service-Type = Framed-User,
Framed-Protocol = PPP,
Framed-IP-Address = 10.10.148.22,
Delegated-IPv6-Prefix = 2001:db8:1b00:100::/64
```

After the PPP LCP phase and RADIUS authentication, the LNS is aware that the subscriber also has IPv6 enabled (in this case because it received the **Delegated-IPv6-Prefix** attribute). As a result, the LNS begins to negotiate both IPCP and IPv6CP with the client. For IPv6CP, only an Interface-ID is negotiated, for which the LNS uses an EUI-64 extended version of the chassis MAC address. Once IPv6CP negotiation is completed, the client can initiate a DHCPv6 Solicit for a delegated prefix (IA_PD option). After a successful Advertise/Request/Reply exchange the subscriber is instantiated as dual-stack IPv4/IPv6.

```
*A:LNS# show service active-subscribers subscriber "subscriber2@isp.net"

=====
Active Subscribers
=====
Subscriber subscriber2@isp.net (sub-profile-1)
-----
(1) SLA Profile Instance sap:[1/2/lns-esm:1.263] - sla:sla-profile-1
-----
IP Address          MAC Address          Session          Origin          Svc          Fwd
-----
10.10.148.22        00:00:14:95:29:18   PPP 1           IPCP            2            Y
2001:db8:1b00:100::/64
                    00:00:14:95:29:18   PPP 1           DHCP6-PD        2            Y
-----
=====
*A:LNS#
```

Conclusion

SR OS offers a comprehensive feature set for LNS implementations. The MS-ISA provides the hardware-assist for L2TP encapsulation/de-capsulation while the carrier IOM implements conventional subscriber management functions.

Multi-Chassis IPsec Redundancy

This chapter provides information about multi-chassis IPsec redundancy configurations.

Topics in this chapter include:

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

Applicability

This initial version of this chapter was based on SR OS Release 10.0.R8, but the CLI in the current edition corresponds to SR OS Release 22.10.R2.

Overview

Multi-Chassis IPsec redundancy (MC-IPsec) is a stateful inter-chassis IPsec failover mechanism. IPsec tunnel states are synchronized between the primary and standby chassis. A tunnel group failure on the primary chassis or a primary chassis failure could trigger MC-IPsec failover to the standby chassis.

The following are some highlights of this feature:


- Internet Key Exchange version 2 (IKEv2) only
- Multi-active tunnel group only
- The granularity of failover is tunnel group, which means a specific tunnel group could failover to the standby chassis independent of other tunnel groups on the primary chassis
- Both static and dynamic LAN-to-LAN tunnels are supported

This feature has the following building blocks:

- Primary chassis election: MC-IPsec mastership protocol (MIMP) runs between the chassis to elect a primary chassis with independent MIMP runs for each tunnel group
- Synchronization: multi-chassis synchronization (MCS) synchronizes the IPsec states between chassis
- Routing:
 - MC-IPsec-aware routing attracts traffic to the primary chassis
 - Shunting support
 - MC-IPsec-aware virtual router redundancy protocol (VRRP)

Figure 78: MC-IPsec architecture shows two redundant IPsec chassis in the middle: a primary chassis and a standby chassis.

Figure 78: MC-IPSec architecture




The fundamentals of MC-IPSec are:

- Only the primary chassis processes encapsulating security payload (ESP) and IKE traffic. If the standby chassis receives traffic, it shunts it to the primary chassis, if possible. The traffic is discarded if the standby chassis fails to shunt the traffic.
- The same local gateway address must be provisioned on both chassis.
- MC-IPSec does not synchronize configurations.
- MC-IPSec-aware routing attracts traffic to the primary chassis for both public and private services, which is achieved by exporting the corresponding IPsec routes to the routing protocol using a route policy and setting a different routing metric according to the MC-IPSec state.
- In case of a Layer 2 public network, MC-IPSec-aware VRRP can be used to trigger VRRP switchover upon MC-IPSec switchover.
- MCS synchronizes IPsec states between chassis so that existing IPsec tunnels do not need to be re-established upon switchover.
- MIMP elects mastership between two chassis, and it can also detect chassis failure and tunnel group failure; a central BFD session can be associated with MIMP to achieve fast chassis failure detection.

Configuration

The example topology is shown in [Figure 79: Example topology](#).

Figure 79: Example topology



The example setup includes:

- an IPSec tunnel initiated by CE-1 and terminated on the primary chassis of the two SeGWs.
- a public IES service "IES-1" and a private VPRN service "VPRN-2" configured on CE-1, SeGW-3, and SeGW-4.
- VPRN 2 (also) configured on P-5.
- a static LAN-to-LAN tunnel with pre-shared key.
- a local VPLS service "VPLS-3" on S-2 to simulate a Layer 2 switch.
- VRRP 10 between SeGW-3 and SeGW-4 to provide a backup address 192.168.1.254, which is the default next hop for CE-1.
- VRRP policy 1 bound to VRRP 10 on the primary chassis SeGW-3 to change the in-use priority upon MC-IPSec switchover.
- OSPF as IGP running in the base routing instance between SeGW-3, SeGW-4, and P-5.
- MP-BGP running between SeGW-3, SeGW-4, and P-5 for the VPN-IPv4 address family.

A ping in VPRN 2 between loopback interface address 192.168.1.1 on CE-1 and 192.168.1.5 on P-5 is used to verify the connectivity over the IPSec tunnel.

The MC-IPSec configuration commands are shown below.

```
config>redundancy>multi-chassis>
  peer <ip-address> [create]
  sync
    ipsec
      tunnel-group <tunnel-group-id> sync-tag <tag-name> [create]
  mc-ipsec
    bfd-enable
```

```

discovery-interval <interval-1> [boot <interval-2>]
hold-on-neighbor-failure <multiplier>
keep-alive-interval <interval>
tunnel-group <tunnel-group-id> [create]
  peer-group <tunnel-group-id>
  priority <priority>
  shutdown

```

```

config>router>policy-options>policy-statement>entry>from>
state ipsec-master-with-peer|ipsec-non-master|ipsec-master-without-peer
protocol ipsec

```

```

config>service>ies>if>
config>service>vprn>if>
static-tunnel-redundant-next-hop <ip-address>
dynamic-tunnel-redundant-next-hop <ip-address>

```

```

config>isa>tunnel-grp>
ipsec-responder-only

```

```

config>vrrp>policy>priority-event>
mc-ipsec-non-forwarding <tunnel-grp-id>
hold-clear <seconds>
hold-set <seconds>
priority <priority-level> explicit

```

The parameters are the following:

- in the **configure redundancy multi-chassis** context:
 - **peer <ip-address> [create]** — This command creates or enters a multi-chassis peer. The peer address is by default the system address. This can be changed on the peer using the **configure redundancy multi-chassis peer source-address** command.
 - **sync** — This command enters the sync configuration context.
 - **ipsec** — This command enables MCS to synchronize IPsec states.
 - **tunnel-group <tunnel-group-id> sync-tag <tag-name>[create]** — This command enables MCS to synchronize the IPsec states of the specified tunnel group. The **sync-tag** parameter is used to match the tunnel group of the peer. The tunnel group states with the same **sync-tag** on both chassis will be synchronized.
 - **mc-ipsec** — This command enters the multi-chassis IPsec configuration context.
 - **bfd-enable** — This command enables tracking a central BFD session; if the BFD session goes down, then the system considers the peer as down and changes the MC-IPsec status of the configured tunnel group accordingly.

The BFD session uses the source address of MCS as its source address and the MCS peer address as the destination address. Other BFD parameters are configured with the **bfd** command on the interface that the MCS source address resides on.

The configuration of BFD is optional for MC-IPsec.
 - **discovery-interval <interval-1> [boot <interval-2>]** — This command specifies the time interval that the tunnel group stays in **discovery** state. Interval 1 is used as discovery interval when a new tunnel group is added to multi-chassis redundancy (**mp-ipsec**); interval 2 is used

as discovery interval after system boot-up. Interval 2 is optional, and when it is not specified, the value for interval 1 is used. Both intervals have a default value of 300 seconds.

- **hold-on-neighbor-failure** <2..25> — This command specifies the number of keep-alive failures before considering the peer to be down. The default value is 3.
- **keep-alive-interval** <5..500> — This command specifies the time interval of the mastership election protocol keep-alive packets in deciseconds. The default value is 10 deciseconds (1 s).
- **tunnel-group** <tunnel-group-id> [**create**] — This command enables multi-chassis redundancy for the specified tunnel group, or enters an already configured tunnel group context. The configured tunnel groups can failover independently.

peer-group <peer-group-id> [**peer-tunnel-group-id**] — This command specifies the corresponding tunnel group ID on the peer node. The peer tunnel group ID is not necessarily equal to local tunnel group ID.

priority <priority> — This command specifies the local priority of the tunnel group, this is used to elect a primary chassis, where the higher number prevails. If the priorities are the same, then the peer which has more active ISAs wins; if the priority and the number of active ISAs are same, then the peer with higher IP address wins. The range is from 0 to 255 and the default value is 100.

- in a **from** statement of a route policy entry:
 - **state ipsec-master-with-peer | ipsec-non-master | ipsec-master-without-peer** — These commands specify the MC-IPSec state in a **from** statement of a route policy entry:
 - **ipsec-master-with-peer**: The tunnel group is the primary chassis with a peer reachable.
 - **ipsec-master-without-peer**: The tunnel group is the primary chassis with peer unreachable.
 - **ipsec-non-master**: The tunnel group is not the primary chassis.
 - **protocol ipsec** — This command specifies IPSec as protocol in a **from** statement of a route policy entry. **protocol ipsec** refers to the /32 local gateway routes (of both static and dynamic tunnels) and reverse route of dynamic tunnel.
- on a public or private IPSec interface in an IES or VPRN service:
 - **static-tunnel-redundant-next-hop** <ip-address> and **dynamic-tunnel-redundant-next-hop** <ip-address> — These commands specify the redundant next hop address on a public or private IPSec interface (with public or private tunnel SAP) for a static and dynamic IPSec tunnel respectively. The specified next hop address is used by the standby chassis to shunt traffic to the primary chassis in case it receives any traffic. The next hop address is resolved in the routing table of the corresponding service.



Note:

- Shunting is supported over:
 - directly connected SAPs
 - spoke SDP terminated IP interfaces
- Shunting over auto-bind tunnel is not supported.
- Shunting does not work if the tunnel group is down.

- in the **isa tunnel-group** <id> context:

- **ipsec-responder-only** — With this command enabled, the system only acts as IKE responder except for the automatic CHILD_SA rekey upon MC-IPSec switchover.

This command is required for MC-IPSec support of static LAN-to-LAN tunnels.

- in the **vrrp policy <id> priority-event** context:
 - **mc-ipsec-non-forwarding <tunnel-grp-id>** — This command creates a VRRP policy priority event: *mc-ipsec-non-forwarding*, which is triggered whenever the specified tunnel group enters the non-forwarding state.
 - **hold-clear <seconds>** — This command configures the hold time before clearing the event. The range is from 0 to 86400 seconds and the default value is 0 s.
 - **hold-set <seconds>** — This command configures the hold time before setting the event. The range is from 0 to 86400 seconds and the default value is 0 s.
 - **priority <priority-level> explicit** — This command sets the VRRP in-use priority to the configured value upon the event. The range is from 0 to 254 and the default value is 0.

The initial configuration must include the following:

- The system time of SeGW-3 and SeGW-4 must be the same for the feature to work. Nokia recommends to use a time synchronization protocol such as NTP or SNTP.
- SeGW-3 and SeGW-4 must be IP reachable in the base routing instance because both MCS and MIMP run in the base routing instance.

Configuration of MC-IPSec

In this section, the following steps are described:

- configure CE-1
- configure S-2
- configure P-5
- configure IPsec tunnel on SeGW-3
- enable MC-IPSec for tunnel group on SeGW-3
- configure MC-IPSec-aware routing on SeGW-3
- configure MC-IPSec-aware VRRP on SeGW-3
- configure SeGW-4

Configure CE-1

On CE-1, the following is configured:

- a public IES service "IES-1" and a private VPRN service "VPRN-2".
- a static default route pointing to the VRRP backup address 172.16.1.254.
- a static IPsec tunnel "tunnel-1" with local address 10.10.10.1 and remote address 10.10.20.1.
- a loopback interface in VPRN 2 with address 192.168.1.1/32 to be used as source address for the ping command to verify the connectivity between CE-1 and P-5 over the IPsec tunnel.

The following base router configuration on CE-1 includes a static route with next hop 172.16.1.254, which is the VRRP backup address.

```
# on CE-1
configure
  router Base
    interface "int-CE-1-S-2"
      address 172.16.1.100/24
      port 1/1/1:1000
    exit
    interface "system"
      address 172.31.2.1/32
    exit
    autonomous-system 64496
    static-route-entry 0.0.0.0/0
      next-hop 172.16.1.254          # VRRP backup address
      no shutdown
    exit
  exit
```

IPSec is configured as follows:

```
configure
  ipsec
    ike-transform 1 create
    exit
    ike-policy 1 create
      ike-version 2
      dpd          # dead peer detection (on peer side; not on MC-IPSec chassis)
      ike-transform 1
    exit
  ipsec-transform 1 create
  exit
```

Tunnel group 1 is configured as follows:

```
configure
  isa
    tunnel-group 1 isa-scale-mode tunnel-limit-2k create
      primary 1/2
      no shutdown
    exit
```

The public IES service is configured as follows:

```
configure
  service
    ies 1 name "IES-1" customer 1 create
      interface "int-IPsec-Public-1" create
        address 10.10.10.254/24
        tos-marking-state untrusted
        sap tunnel-1.public:1 create
      exit
    exit
  no shutdown
  exit
```


The private VPRN service on CE-1 is configured as follows. Instead of configuring **delivery-service 1** for the IPsec tunnel, it is possible to configure **delivery-service-name "IES-1"**.

```
configure
  service
    vprn 2 name "VPRN-2" customer 1 create
      ipsec
        security-policy 1 create
          entry 10 create
            local-ip 192.168.1.1/32
            remote-ip 192.168.1.5/32
          exit
        exit
      exit
    interface int-loopback-1 create
      address 192.168.1.1/32
      loopback
    exit
    interface int-IPsec-private-1 tunnel create
      sap tunnel-1.private:1 create
      ipsec-tunnel "tunnel-1" create
        security-policy 1
        local-gateway-address 10.10.10.1 peer 10.10.20.1 delivery-service 1
        dynamic-keying
          ike-policy 1
          pre-shared-key "pass"
          transform 1
        exit
      no shutdown
    exit
  exit
  static-route-entry 192.168.1.5/32
    ipsec-tunnel "tunnel-1"
    no shutdown
  exit
exit
no shutdown
```

Configure S-2

On S-2, a local VPLS service 3 simulates a Layer 2 switch between CE-1, SeGW-3, and SeGW-4:

```
# on S-2
configure
  service
    vpls 3 name "VPLS-3" customer 1 create
      sap 1/1/c1/1:1 create
        description "to SAP in IES 1 on SeGW-3"
      exit
      sap 1/1/c1/2:1000 create
        description "to router interface on CE-1"
      exit
      sap 1/1/c1/3:1 create
        description "to SAP in IES 1 on SeGW-4"
      exit
    no shutdown
  exit
```

Configure P-5

P-5 simulates the core network router, connecting to SeGW-3 and SeGW-4. The configuration on P-5 includes the following:

- a loopback interface with address 192.168.1.5/32 in VPRN 2 is the destination address of the ping traffic from CE-1.
- an MP-BGP session for the VPN-IPv4 address family between P-5, SeGW-3, and SeGW-4.
- GRE spoke SDPs to connect to SeGW-3 and SeGW-4.

On P-5, the following router interfaces are configured in the base router. OSPF is used as IGP.

```
# on P-5
configure
router Base
  interface "int-P-5-SeGW-3"
    address 192.168.35.2/30
    port 1/1/c1/2:1000
  exit
  interface "int-P-5-SeGW-4"
    address 192.168.45.2/30
    port 1/1/c1/1:1000
  exit
  interface "system"
    address 192.0.2.5/32
  exit
  ospf 0
    area 0.0.0.0
      interface system
      exit
      interface "int-P-5-SeGW-3"
      exit
      interface "int-P-5-SeGW-4"
      exit
    exit
  no shutdown
exit
```

On P-5, the following GRE SDPs are configured toward SeGW-3 and SeGW-4:

```
configure
service
  sdp 53 create
    description "GRE SDP toward SeGW-3"
    signaling off
    far-end 192.0.2.3
    no shutdown
  exit
  sdp 54 create
    description "GRE SDP toward SeGW-4"
    signaling off
    far-end 192.0.2.4
    no shutdown
  exit
```

VPRN 2 is configured on P-5, as follows:

```
configure
service
```

```
vrpn 2 name "VPRN-2" customer 1 create
  interface "int-loopback-1" create
    address 192.168.1.5/32
    loopback
  exit
  bgp-ipvpn
    mpls
    route-distinguisher 64496:2
    vrf-target target:64496:2
    no shutdown
  exit
  exit
  spoke-sdp 53:2 create
  exit
  spoke-sdp 54:2 create
  exit
  no shutdown
exit
```

The BGP configuration on P-5 is as follows:

```
configure
  router Base
    autonomous-system 64496
    bgp
      group "MPBGP"
        family vpn-ipv4
        type internal
        neighbor 192.0.2.3
        exit
        neighbor 192.0.2.4
        exit
      exit
    no shutdown
  exit
```

Configure IPsec tunnel on SeGW-3

The configuration on SeGW-3 is described in four consecutive sections. In this first section, the following is configured:

- the tunnel group, which must be in multi-active mode before MC-IPsec can be enabled.
- an interface "int-Redundant-1", which is a spoke-SDP terminated interface used for shunting.
- GRE SDP 34 toward SeGW-4 and GRE SDP 35 toward P-5.
- IPsec tunnel "tunnel-1" is the tunnel to CE-1; both SeGW-3 and SeGW-4 use the same local gateway address: 10.10.20.1.

The following configures tunnel group 1 on SeGW-3:

```
# on SeGW-3
configure
  isa
    tunnel-group 1 isa-scale-mode tunnel-limit-2k create
    ipsec-responder-only
    multi-active
    mda 1/2
    no shutdown
```

```
exit
```

On SeGW-3, the following router interfaces are configured in the base router. A static route is configured toward CE-1. OSPF is the IGP used between SeGW-3, SeGW-4, and P-5.

```
configure
router Base
  interface "int-SeGW-3-P-5"
    address 192.168.35.1/30
    port 1/1/1:1000
    no shutdown
  exit
  interface "int-SeGW-3-SeGW-4"
    address 192.168.34.1/30
    port 1/1/3:1000
    no shutdown
  exit
  interface "system"
    address 192.0.2.3/32
    bfd 100 receive 100 multiplier 3
    no shutdown
  exit
  static-route-entry 10.10.10.0/24
    next-hop 172.16.1.100
    no shutdown
  exit
  exit
  ospf 0
    area 0.0.0.0
      interface "system"
        no shutdown
      exit
      interface "int-SeGW-3-P-5"
        no shutdown
      exit
      interface "int-SeGW-3-SeGW-4"
        no shutdown
      exit
    exit
  exit
  no shutdown
exit
```

The IPSec settings are as follows:

```
configure
ipsec
  ike-transform 1 create
    isakmp-lifetime 172800
  exit
  ike-policy 1 create
    ike-version 2
    ipsec-lifetime 7200
    ike-transform 1
  exit
  ipsec-transform 1 create
  exit
exit
```

The GRE SDPs are configured as follows:

```
configure
service
```

```

sdp 34 create
  description "GRE SDP toward SeGW-4"
  signaling off
  far-end 192.0.2.4
  keep-alive
    shutdown
  exit
  no shutdown
exit
sdp 35 create
  description "GRE SDP toward P-5"
  signaling off
  far-end 192.0.2.5
  keep-alive
    shutdown
  exit
  no shutdown
exit

```

The public IES service is configured as follows. The VRRP configuration will be added in a later step.

```

configure
  service
    ies 1 name "IES-1" customer 1 create
      interface "int-SeGW-3-S-2" create
        address 172.16.1.252/24
        sap 1/1/2:1 create
          description "SAP to switch S-2"
        exit
      exit
      interface "int-IPsec-Public-1" create
        address 10.10.20.254/24
        tos-marking-state untrusted
        sap tunnel-1.public:1 create
        exit
        static-tunnel-redundant-next-hop 192.168.34.2
      exit
    no shutdown
  exit

```

The private VPRN service is configured as follows:

```

configure
  service
    vprn 2 name "VPRN-2" customer 1 create
      ipsec
        security-policy 1 create
          entry 10 create
            local-ip 192.168.1.5/32
            remote-ip 192.168.1.1/32
          exit
        exit
      exit
      interface "int-IPsec-Private-1" tunnel create
        sap tunnel-1.private:1 create
        ipsec-tunnel "tunnel-1" create
          security-policy 1
          local-gateway-address 10.10.20.1 peer 10.10.10.1 delivery-service 1
          dynamic-keying
            ike-policy 1
            pre-shared-key "pass"
          transform 1
        exit

```

```
        no shutdown
        exit
    exit
    static-tunnel-redundant-next-hop 192.168.20.2
exit
interface "int-Redundant-1" create
    description "interface used for shunting"
    address 192.168.20.1/30
    spoke-sdp 34:20 create
        ingress
            vc-label 2049
        exit
        egress
            vc-label 2048
        exit
    no shutdown
    exit
exit
static-route-entry 192.168.1.1/32
    ipsec-tunnel "tunnel-1"
        no shutdown
    exit
exit
bgp-ipvpn
    mpls
        route-distinguisher 64496:2
        vrf-target target:64496:2
        no shutdown
    exit
exit
spoke-sdp 34:2 create
    description "SDP to SeGW-4"
exit
spoke-sdp 35:2 create
    description "SDP to P-5"
exit
no shutdown
exit
```

Enable MC-IPSec for tunnel group 1 on SeGW-3

In this section, the following steps are described:

- Create a multi-chassis peer using the system address of SeGW-4.
- Enable MCS for IPsec and tunnel group 1.
- Enable MC-IPSec for the tunnel group with a configured priority 200.
- Bind a central BFD session to MC-IPSec from the system interface.

Create multi-chassis peer 192.0.2.4 and enable MCS and MC-IPSec for tunnel group 1:

```
# on SeGW-3
configure
    redundancy
        multi-chassis
            peer 192.0.2.4 create
            sync
                ipsec
                    tunnel-group 1 sync-tag "tag-1" create
```

```

        no shutdown
    exit
    mc-ipsec
        bfd-enable
        tunnel-group 1 create
            peer-group 1
            priority 200
            no shutdown
        exit
    exit
    no shutdown
exit
exit

```

BFD is enabled in the preceding configuration. On the system interface, the following BFD settings are configured:

```

configure
  router Base
    interface "system"
      address 192.0.2.3/32
      bfd 100 receive 100 multiplier 3
      no shutdown
    exit

```

Configure MC-IPSec-aware routing on SeGW-3

In this step, a route policy is defined and applied to VPRN 2.

Route policy "IPsec-to-MPBGP" exports static route 192.168.1.1/32 in VPRN 2 to P-5. This policy sets the local preference of the prefix 192.168.1.1/32 according to the MC-IPSec state:

- for the **ipsec-master-with-peer** state: local preference 200
- for the **ipsec-non-master** state: local preference 100
- for the **ipsec-master-without-peer** state: local preference 200

The state **ipsec-master-without-peer** can be used to attract traffic to the designated primary chassis in case of "dual master" (meaning two chassis lose the MIMP connection in the base routing instance). In this example, SeGW-3 has local preference 200 and SeGW-4 has local preference 100 for **ipsec-master-without-peer**.

The route policy is configured as follows:

```

# on SeGW-3:
configure
  router Base
    policy-options
      begin
        prefix-list "CE-1-Internal"
          prefix 192.168.1.1/32 exact
        exit
        community "vprn2"
          members "target:64496:2"
        exit
        policy-statement "IPsec-to-MPBGP"
          entry 10
            from
              prefix-list "CE-1-Internal"
              state ipsec-master-with-peer

```

```

        exit
        action accept
        community add "vprn2"
        local-preference 200
    exit
exit
entry 20
    from
        prefix-list "CE-1-Internal"
        state ipsec-non-master
    exit
    action accept
    community add "vprn2"
    local-preference 100
    exit
exit
entry 30
    from
        prefix-list "CE-1-Internal"
        state ipsec-master-without-peer
    exit
    action accept
    community add "vprn2"
    local-preference 200
    exit
exit
default-action accept
    community add "vprn2"
exit
exit
commit

```

The BGP configuration on SeGW-3 is as follows:

```

configure
router Base
    autonomous-system 64496
    bgp
        group "MPBGP"
        family vpn-ipv4
        type internal
        neighbor 192.0.2.4
        exit
        neighbor 192.0.2.5
        exit
    exit
    no shutdown
exit

```

The route policy is applied as **vrf-export** in VPRN 2:

```

configure
service
    vprn "VPRN-2"
    bgp-ipvpn
    mpls
        vrf-export "IPsec-to-MPBGP"
    exit
    exit
exit

```


Configure MC-IPSec-aware VRRP on SeGW-3

In this section, a VRRP policy is defined that uses the **mc-ipsec-non-forwarding** priority event to lower the in-use VRRP priority upon MC-IPSec switchover, which ensures VRRP and MC-IPSec have the same primary chassis. The VRRP instance needs to be in preempt mode.

This VRRP policy is only configured on the designated VRRP primary chassis SeGW-3, not on the standby chassis. The VRRP policy is applied to the interface "int-SeGW3-S-2" of IES 1.

VRRP policy 1 is configured as follows:

```
# on SeGW-3:
configure
  vrrp
    policy 1
      priority-event
        mc-ipsec-non-forwarding 1
        priority 50 explicit
      exit
    exit
  exit
```

The VRRP policy is applied in VRRP instance 10 in the IES service:

```
configure
  service
    ies "IES-1"
      interface "int-SeGW-3-S-2"
        address 172.16.1.252/24
        vrrp 10
          backup 172.16.1.254
          priority 200
          policy 1
          ping-reply
        exit
      exit
    exit
  ---snip---
```

Configure SeGW-4

The configuration on the standby chassis SeGW-4 is similar, but with different priorities and without the VRRP policy.

The tunnel group is configured in multi-active mode:

```
# on SeGW-4
configure
  isa
    tunnel-group 1 create
      ipsec-responder-only
      multi-active
      mda 1/2
      no shutdown
    exit
```

The MCS and MC-IPSec configuration is as follows:

```
configure
```

```
redundancy
  multi-chassis
    peer 192.0.2.3 create
    sync
      ipsec
        tunnel-group 1 sync-tag "tag-1" create
        no shutdown
      exit
    mc-ipsec
      bfd-enable
      tunnel-group 1 create
        peer-group 1
        priority 150
        no shutdown
      exit
    exit
  no shutdown
exit
exit
```

The base router configuration on SeGW-4 includes the following router interfaces and a static route to CE-1. OSPF is used as IGP between SeGW-3, SeGW-4, and P-5.

```
configure
  router Base
    interface "int-SeGW-4-P-5"
      address 192.168.45.1/30
      port 1/1/2:1000
      no shutdown
    exit
    interface "int-SeGW-4-SeGW-3"
      address 192.168.34.2/30
      port 1/1/3:1000
      no shutdown
    exit
    interface "system"
      address 192.0.2.4/32
      bfd 100 receive 100 multiplier 3
      no shutdown
    exit
    static-route-entry 10.10.10.0/24
      next-hop 172.16.1.100
      no shutdown
    exit
  exit
  ospf 0
    area 0.0.0.0
      interface "system"
        no shutdown
      exit
      interface "int-SeGW-4-SeGW-3"
        no shutdown
      exit
      interface "int-SeGW-4-P-5"
        no shutdown
      exit
    exit
  no shutdown
exit
```

The IPsec configuration is as follows:

```
configure
  ipsec
    ike-transform 1 create
      isakmp-lifetime 172800
    exit
    ike-policy 1 create
      ike-version 2
      ipsec-lifetime 7200
      ike-transform 1
    exit
    ipsec-transform 1 create
  exit
```

The following route policy is configured on SeGW-4, The local preference is lower for the **ipsec-master-without-peer** state.

```
configure
  router Base
    policy-options
      begin
        prefix-list "CE-1-Internal"
          prefix 192.168.1.1/32 exact
        exit
        community "vprn2"
          members "target:64496:2"
        exit
        policy-statement "IPsec-to-MPBGP"
          entry 10
            from
              prefix-list "CE-1-Internal"
              state ipsec-master-with-peer
            exit
            action accept
              community add "vprn2"
              local-preference 200
            exit
          exit
          entry 20
            from
              prefix-list "CE-1-Internal"
              state ipsec-non-master
            exit
            action accept
              community add "vprn2"
              local-preference 100
            exit
          exit
          entry 30
            from
              prefix-list "CE-1-Internal"
              state ipsec-master-without-peer
            exit
            action accept
              community add "vprn2"
              local-preference 100 # lower preference on standby SeGW
            exit
          exit
        default-action accept
          community add "vprn2"
        exit
      exit
```

```
commit
```

The BGP configuration on SeGW-4 is as follows:

```
configure
router Base
  autonomous-system 64496
  bgp
    group "MPBGP"
      family vpn-ipv4
      type internal
      neighbor 192.0.2.3
      exit
      neighbor 192.0.2.5
      exit
    exit
  no shutdown
exit
```

The following GRE SDPs are configured:

```
configure
service
  sdp 43 create
  description "GRE SDP toward SeGW-3"
  signaling off
  far-end 192.0.2.3
  keep-alive
  shutdown
  exit
  no shutdown
exit
  sdp 45 create
  description "GRE SDP toward P-5"
  signaling off
  far-end 192.0.2.5
  keep-alive
  shutdown
  exit
  no shutdown
exit
```

The public IES service is configured as follows:

```
configure
service
  ies 1 name "IES-1" customer 1 create
  interface "int-SeGW-4-S-2" create
  address 172.16.1.253/24
  vrrp 10
  backup 172.16.1.254
  ping-reply
  exit
  sap 1/1/1:1 create
  description "SAP toward switch S-2"
  exit
exit
interface "int-IPsec-Public-1" create
address 10.10.20.254/24
tos-marking-state untrusted
sap tunnel-1.public:1 create
exit
```

```
        static-tunnel-redundant-next-hop 192.168.34.1
    exit
    no shutdown
exit
```

The private VPRN service is configured as follows:

```
configure
  service
    vprn 2 name "VPRN-2" customer 1 create
    ipsec
      security-policy 1 create
        entry 10 create
          local-ip 192.168.1.5/32
          remote-ip 192.168.1.1/32
        exit
      exit
    exit
  interface "int-IPsec-Private-1" tunnel create
    sap tunnel-1.private:1 create
    ipsec-tunnel "tunnel-1" create
      security-policy 1
      local-gateway-address 10.10.20.1 peer 10.10.10.1 delivery-service 1
      dynamic-keying
        ike-policy 1
        pre-shared-key "pass"
        transform 1
      exit
    no shutdown
    exit
  exit
  static-tunnel-redundant-next-hop 192.168.20.1
exit
interface "int-Redundant-1" create
  description "interface used for shunting"
  address 192.168.20.2/30
  spoke-sdp 43:20 create
    ingress
      vc-label 2048
    exit
    egress
      vc-label 2049
    exit
  no shutdown
  exit
exit
static-route-entry 192.168.1.1/32
  ipsec-tunnel "tunnel-1"
  no shutdown
  exit
exit
bgp-ipvpn
  mpls
    route-distinguisher 64496:2
    vrf-export "IPsec-to-MPBGP"
    vrf-target target:64496:2
    no shutdown
  exit
exit
spoke-sdp 43:2 create
  description "SDP to SeGW-3"
exit
spoke-sdp 45:2 create
```

```

        description "SDP to P-5"
    exit
    no shutdown
exit
    
```

Verification

The following will be verified in this section:

- the MC-IPSec status and VRRP status on SeGW-3 and SeGW-4
- the status of the IPSec tunnel on CE-1
- the status of the IPSec tunnel on the SeGWs

Verify the MC-IPSec status on SeGW-3 and SeGW-4

The following is verified:

- SeGW-3 is the primary chassis and SeGW-4 is the standby for tunnel group 1 because SeGW-3 has the higher priority 200.
- SeGW-3 is the primary node for VRRP instance 10 and SeGW-4 is the backup.

SeGW-3 is the primary chassis in tunnel group 1 with priority 200:

```
*A:SeGW-3# show redundancy multi-chassis mc-ipsec peer 192.0.2.4
```

```
=====
Multi-Chassis MC-IPsec
=====
```

```
Peer Name      : (Not Specified)
Peer Addr      : 192.0.2.4
Keep Alive Intvl: 1.0 secs          Hold on Nbr Fail   : 3
Discovery Intvl : 300 secs          Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update    : 02/13/2023 09:37:51
```

```
=====
Multi-Chassis IPsec Multi Active Tunnel-Group Table
=====
```

| ID | Peer Group | Priority | Admin State | Mastership |
|----|------------|----------|-------------|------------|
| 1 | 1 | 200 | Up | master |

```
-----
Multi Active Tunnel Group Entries found: 1
=====
```

SeGW-4 is the standby chassis in tunnel group 1 with priority 150:

```
*A:SeGW-4# show redundancy multi-chassis mc-ipsec peer 192.0.2.3
```

```
=====
Multi-Chassis MC-IPsec
=====
```

```
Peer Name      : (Not Specified)
Peer Addr      : 192.0.2.3
Keep Alive Intvl: 1.0 secs          Hold on Nbr Fail   : 3
```

```
Discovery Intvl : 300 secs           Discovery Boot Intvl : 300 secs
BFD              : Enable
Last update     : 02/13/2023 09:38:21
```

```
=====
Multi-Chassis IPsec Multi Active Tunnel-Group Table
=====
```

| ID | Peer Group | Priority | Admin State | Mastership |
|----|------------|----------|-------------|------------|
| 1 | 1 | 150 | Up | standby |

```
-----
Multi Active Tunnel Group Entries found: 1
=====
```

SeGW-3 is the primary node for VRRP instance 10:

```
*A:SeGW-3# show router vrrp instance
```

```
=====
VRRP Instances
=====
```

| Interface Name | VR Id | Own | Adm | State | Base Pri | Msg Int |
|------------------------|-----------|-----|-----|---------------|-----------|---------|
| | IP | | Opr | Pol Id | InUse Pri | Inh Int |
| int- SeGW-3-S-2 | 10 | No | Up | Master | 200 | 1 |
| | IPv4 | | Up | 1 | 200 | No |

```
Backup Addr: 172.16.1.254
```

```
-----
Instances : 1
=====
```

SeGW-4 is backup for VRRP instance 10:

```
*A:SeGW-4# show router vrrp instance
```

```
=====
VRRP Instances
=====
```

| Interface Name | VR Id | Own | Adm | State | Base Pri | Msg Int |
|------------------------|-----------|-----|-----|---------------|-----------|---------|
| | IP | | Opr | Pol Id | InUse Pri | Inh Int |
| int- SeGW-4-S-2 | 10 | No | Up | Backup | 100 | 1 |
| | IPv4 | | Up | n/a | 100 | No |

```
Backup Addr: 172.16.1.254
```

```
-----
Instances : 1
=====
```

Verify the IPsec tunnel on CE-1

The following is verified in this section:

- the connectivity between CE-1 and P-5
- the IPsec tunnel information

A ping command is launched from the loopback interface in VPRN 2 on CE-1 to the loopback interface in VPRN 2 on P-5:

```
*A:CE-1# ping router 2 192.168.1.5
PING 192.168.1.5 56 data bytes
64 bytes from 192.168.1.5: icmp_seq=1 ttl=63 time=3.97ms.
64 bytes from 192.168.1.5: icmp_seq=2 ttl=63 time=7.50ms.
64 bytes from 192.168.1.5: icmp_seq=3 ttl=63 time=2.84ms.
64 bytes from 192.168.1.5: icmp_seq=4 ttl=63 time=2.61ms.
64 bytes from 192.168.1.5: icmp_seq=5 ttl=63 time=2.67ms.

---- 192.168.1.5 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.61ms, avg = 3.92ms, max = 7.50ms, stddev = 1.86ms
```

The following command shows the IPsec tunnel information.

```
*A:CE-1# show ipsec tunnel

=====
IPsec Tunnels
=====
TunnelName          LocalAddress      SvcId      Admn  Keying
 SapId              RemoteAddress    DlvrySvcId Oper   Sec
                   Plcy
-----
tunnel-1            10.10.10.1       2          Up    Dynamic
 tunnel-1.private:1 10.10.20.1       1          Up    1
-----
IPsec Tunnels: 1
=====
```

Verify the IPsec tunnel on the SeGWs

In this section, the following is verified:

- the MCS database is in-sync, so the tunnel status is up on both chassis.
- P-5 receives two VPN-IPv4 routes for prefix 192.168.1.1/32: the route from SeGW-3 has local preference 200; the route from SeGW-4 has local preference 100.

On both SeGWs, the IPsec tunnel with local address 10.10.20.1 and remote address 10.10.10.1 is up:

```
*A:SeGW-3# show ipsec tunnel

=====
IPsec Tunnels
=====
TunnelName          LocalAddress      SvcId      Admn  Keying
 SapId              RemoteAddress    DlvrySvcId Oper   Sec
                   Plcy
-----
tunnel-1            10.10.20.1       2          Up    Dynamic
 tunnel-1.private:1 10.10.10.1       1          Up    1
-----
IPsec Tunnels: 1
=====
```



```
*A:SeGW-4# show ipsec tunnel
```

```
=====
IPsec Tunnels
=====
TunnelName          LocalAddress      SvcId      Admn  Keying
  SapId             RemoteAddress    DlvrySvcId Oper   Sec
                   Plcy
-----
tunnel-1            10.10.20.1       2          Up    Dynamic
  tunnel-1.private:1 10.10.10.1       1          Up    1
-----
IPsec Tunnels: 1
=====
```

MCS is in sync on both SeGWs:

```
*A:SeGW-3# show redundancy multi-chassis sync
```

```
=====
Multi-chassis Peer Table
=====
Peer
-----
Peer IP Address      : 192.0.2.4
Description          : (Not Specified)
Authentication       : Disabled
Source IP Address    : 192.0.2.3
Admin State          : Enabled
Warm standby         : No
Remote warm standby  : No
Sub-mgmt options     :
  DHCP lease threshold : Inactive
  Local / Remote       : -- / --
-----
Sync-status
-----
Client Applications  : IPsec
Sync Admin State     : Up
Sync Oper State      : Up
Sync Oper Flags      :
DB Sync State        : inSync
Num Entries          : 2
Lcl Deleted Entries  : 0
Alarm Entries        : 0
OMCR Standby Entries : 0
OMCR Alarm Entries   : 0
Rem Num Entries      : 2
Rem Lcl Deleted Entries : 0
Rem Alarm Entries    : 0
Rem OMCR Standby Entries : 0
Rem OMCR Alarm Entries : 0
=====
```

```
*A:SeGW-4# show redundancy multi-chassis sync
```

```
=====
Multi-chassis Peer Table
=====
Peer
-----
```

```

Peer IP Address      : 192.0.2.3
Description         : (Not Specified)
Authentication      : Disabled
Source IP Address   : 192.0.2.4
Admin State         : Enabled
Warm standby        : No
Remote warm standby : No
Sub-mgmt options    :
  DHCP lease threshold : Inactive
  Local / Remote      : -- / --
-----
Sync-status
-----
Client Applications  : IPsec
Sync Admin State    : Up
Sync Oper State     : Up
Sync Oper Flags     :
DB Sync State       : inSync
Num Entries         : 2
Lcl Deleted Entries : 0
Alarm Entries       : 0
OMCR Standby Entries : 0
OMCR Alarm Entries  : 0
Rem Num Entries     : 2
Rem Lcl Deleted Entries : 0
Rem Alarm Entries   : 0
Rem OMCR Standby Entries : 0
Rem OMCR Alarm Entries : 0
=====
=====

```

The following command shows that P-5 received two VPN-IPv4 routes for prefix 192.168.1.1/32: one from SeGW-3 with local preference 200 and one from SeGW-4 with local preference 100:

```

*A:P-5# show router bgp routes vpn-ipv4
=====
BGP Router ID:192.0.2.5      AS:64496      Local AS:64496
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
               l - leaked, x - stale, > - best, b - backup, p - purge
Origin codes  : i - IGP, e - EGP, ? - incomplete
=====
BGP VPN-IPv4 Routes
=====

```

| Flag | Network Nexthop (Router) As-Path | LocalPref Path-Id | MED IGP Cost Label |
|------|--|----------------------|--------------------------|
| u*>i | 64496:2:192.168.1.1/32 192.0.2.3 No As-Path | 200 None | None 10 524287 |
| *i | 64496:2:192.168.1.1/32 192.0.2.4 No As-Path | 100 None | None 10 524287 |
| u*>i | 64496:2:192.168.20.0/30 192.0.2.3 No As-Path | 100 None | None 10 524287 |
| *>i | 64496:2:192.168.20.0/30 192.0.2.4 No As-Path | 100 None | None 10 524287 |
| u*>i | 64496:2:192.168.20.1/32 | 100 | 0 |

```

192.0.2.3          None          10
No As-Path        524287
u*>i 64496:2:192.168.20.2/32 100
192.0.2.4        None          10
No As-Path        524287
-----
Routes : 6
=====

```

MC-IPSec failover scenarios

Two MC-IPSec failover scenarios are described in this section:

- MC-IPSec failover when MS-ISA is disabled
- MC-IPSec failover when the primary chassis SeGW-3 reboots

Failover when MS-ISA is disabled

Initially, SeGW-3 is the primary chassis and SeGW-4 is the standby:

```

*A:SeGW-3# show redundancy multi-chassis mc-ipsec peer 192.0.2.4
=====
Multi-Chassis MC-IPsec
=====
Peer Name       : (Not Specified)
Peer Addr      : 192.0.2.4
Keep Alive Intvl: 1.0 secs           Hold on Nbr Fail   : 3
Discovery Intvl : 300 secs           Discovery Boot Intvl : 300 secs
BFD            : Enable
Last update    : 02/13/2023 09:37:51
=====
Multi-Chassis IPsec Multi Active Tunnel-Group Table
=====
ID          Peer Group  Priority  Admin State  Mastership
-----
1           1           200      Up           master
-----
Multi Active Tunnel Group Entries found: 1
=====

```

```

*A:SeGW-3# show router vrrp instance
=====
VRRP Instances
=====
Interface Name          VR Id  Own  Adm  State  Base Pri  Msg Int
                       IP      Opr  Pol Id  InUse Pri  Inh Int
-----
int-SeGW-3-S-2         10    No  Up   Master  200      1
                       IPv4    Up   1     200      No
Backup Addr: 172.16.1.254
-----
Instances : 1

```

```

=====
*A:SeGW-4# show redundancy multi-chassis mc-ipsec peer 192.0.2.3
=====
Multi-Chassis MC-IPsec
=====
Peer Name      : (Not Specified)
Peer Addr     : 192.0.2.3
Keep Alive Intvl: 1.0 secs           Hold on Nbr Fail      : 3
Discovery Intvl : 300 secs           Discovery Boot Intvl  : 300 secs
BFD           : Enable
Last update    : 02/13/2023 09:38:21
=====
Multi-Chassis IPsec Multi Active Tunnel-Group Table
=====
ID            Peer Group  Priority Admin State  Mastership
-----
1             1           150    Up           standby
-----
Multi Active Tunnel Group Entries found: 1
=====

```

```

*A:SeGW-4# show router vrrp instance
=====
VRRP Instances
=====
Interface Name          VR Id Own Adm State      Base Pri  Msg Int
                        IP      Opr Pol Id   InUse Pri  Inh Int
-----
int-SeGW-4-S-2         10  No  Up  Backup   100     1
                        IPv4   Up  n/a     100     No
Backup Addr: 172.16.1.254
-----
Instances : 1
=====

```

The following command disables the MS-ISA on the primary chassis SeGW-3, which will trigger an MC-IPSec failover.

```

configure
card 1
  mda 2
  shutdown

```

With MS-ISA disabled, the MC-IPSec state of tunnel group 1 on SeGW-3 becomes **notEligible**, which means that the tunnel group is down, see the *7450 ESS, 7750 SR, and VSR Multiservice ISA and ESA Guide* for details description of MIMP states.:

```

*A:SeGW-3# show redundancy multi-chassis mc-ipsec peer 192.0.2.4
=====
Multi-Chassis MC-IPsec
=====
Peer Name      : (Not Specified)
Peer Addr     : 192.0.2.4
Keep Alive Intvl: 1.0 secs           Hold on Nbr Fail      : 3

```

```
Discovery Intvl : 300 secs          Discovery Boot Intvl : 300 secs
BFD              : Enable
Last update      : 02/13/2023 09:37:51
```

```
=====
Multi-Chassis IPsec Multi Active Tunnel-Group Table
=====
```

| ID | Peer Group | Priority | Admin State | Mastership |
|----|------------|----------|-------------|--------------------|
| 1 | 1 | 200 | Up | notEligible |

```
-----
Multi Active Tunnel Group Entries found: 1
=====
```

SeGW-3 is backup for VRRP instance 10 with in-use priority 50, as per the VRRP policy 1:

```
*A:SeGW-3# show router vrrp instance
```

```
=====
VRRP Instances
=====
```

| Interface Name | VR Id | Own | Adm | State | Base Pri | Msg Int |
|----------------|-------|-----|-----|---------------|-----------|---------|
| | IP | | Opr | Pol Id | InUse Pri | Inh Int |
| int-SeGW-3-S-2 | 10 | No | Up | Backup | 200 | 1 |
| | IPv4 | | Up | 1 | 50 | No |

```
Backup Addr: 172.16.1.254
```

```
-----
Instances : 1
=====
```

SeGW-4 is now the primary chassis in tunnel group 1. This is triggered by MC-IPSec failover, as per the **mc-ipsec-non-forwarding** event in VRRP policy 1.

```
*A:SeGW-4# show redundancy multi-chassis mc-ipsec peer 192.0.2.3
```

```
=====
Multi-Chassis MC-IPsec
=====
```

```
Peer Name      : (Not Specified)
Peer Addr      : 192.0.2.3
Keep Alive Intvl: 1.0 secs          Hold on Nbr Fail      : 3
Discovery Intvl : 300 secs          Discovery Boot Intvl  : 300 secs
BFD            : Enable
Last update    : 02/13/2023 09:38:21
```

```
=====
Multi-Chassis IPsec Multi Active Tunnel-Group Table
=====
```

| ID | Peer Group | Priority | Admin State | Mastership |
|----|------------|----------|-------------|---------------|
| 1 | 1 | 150 | Up | master |

```
-----
Multi Active Tunnel Group Entries found: 1
=====
```

SeGW-4 is primary for VRRP instance 10;

```
*A:SeGW-4# show router vrrp instance
```

```

=====
VRRP Instances
=====
Interface Name          VR Id Own Adm State      Base Pri  Msg Int
                        IP      Opr Pol Id    InUse Pri  Inh Int
-----
int-SeGW-4-S-2         10   No  Up  Master    100      1
                        IPv4    Up  n/a      100      No
Backup Addr: 172.16.1.254
-----
Instances : 1
=====

```

The situation is restored by enabling MS-ISA on SeGW-3:

```

configure
  card 1
    mda 2
      no shutdown

```

MC-IPSec failover when primary chassis reboots

The following **tools** command on SeGW-3 triggers an MC-IPSec switchover:

```

tools perform redundancy multi-chassis mc-ipsec force-switchover tunnel-group 1

A:SeGW-3# tools perform redundancy multi-chassis mc-ipsec force-switchover tunnel-group 1
WARNING! Forcing a mastership switchover may significantly impact traffic. Are you sure (y/n)?
y

```

SeGW-3 is the primary chassis for tunnel group 1:

```

A:SeGW-3# show redundancy multi-chassis mc-ipsec peer 192.0.2.4

=====
Multi-Chassis MC-IPsec
=====
Peer Name      : (Not Specified)
Peer Addr      : 192.0.2.4
Keep Alive Intvl: 1.0 secs          Hold on Nbr Fail   : 3
Discovery Intvl : 300 secs          Discovery Boot Intvl : 300 secs
BFD            : Enable
Last update    : 02/13/2023 09:37:51

=====
Multi-Chassis IPsec Multi Active Tunnel-Group Table
=====
ID      Peer Group  Priority  Admin State  Mastership
-----
1       1           200      Up           master
-----
Multi Active Tunnel Group Entries found: 1
=====

```

SeGW-3 is primary for VRRP instance 10:

```

A:SeGW-3# show router vrrp instance

```

```
=====
VRRP Instances
=====
Interface Name          VR Id Own Adm State      Base Pri  Msg Int
                       IP      Opr Pol Id    InUse Pri  Inh Int
-----
int-SeGW-3-S-2         10   No  Up  Master    200      1
                       IPv4    Up   1      200      No
Backup Addr: 172.16.1.254
-----
Instances : 1
=====
```

SeGW-4 is the standby chassis for tunnel group 1:

```
*A:SeGW-4# show redundancy multi-chassis mc-ipsec peer 192.0.2.3
```

```
=====
Multi-Chassis MC-IPsec
=====
```

```
Peer Name      : (Not Specified)
Peer Addr      : 192.0.2.3
Keep Alive Intvl: 1.0 secs          Hold on Nbr Fail   : 3
Discovery Intvl : 300 secs          Discovery Boot Intvl : 300 secs
BFD            : Enable
Last update    : 02/13/2023 09:38:21
```

```
=====
Multi-Chassis IPsec Multi Active Tunnel-Group Table
=====
```

| ID | Peer Group | Priority | Admin State | Mastership |
|----|------------|----------|-------------|----------------|
| 1 | 1 | 150 | Up | standby |

```
Multi Active Tunnel Group Entries found: 1
=====
```

The VRRP state on SeGW-4 is backup:

```
*A:SeGW-4# show router vrrp instance
```

```
=====
VRRP Instances
=====
```

| Interface Name | VR Id | Own | Adm | State | Base Pri | Msg Int |
|----------------|-------|-----|-----|---------------|-----------|---------|
| | IP | | Opr | Pol Id | InUse Pri | Inh Int |
| int-SeGW-4-S-2 | 10 | No | Up | Backup | 100 | 1 |
| | IPv4 | | Up | n/a | 100 | No |

```
Backup Addr: 172.16.1.254
```

```
-----
Instances : 1
=====
```

The following command reboots the primary chassis SeGW-3:

```
admin reboot now
```

While SeGW-3 reboots, the IPsec state of SeGW-4 becomes **eligible**:

```
*A:SeGW-4# show redundancy multi-chassis mc-ipsec peer 192.0.2.3
=====
Multi-Chassis MC-IPsec
=====
Peer Name      : (Not Specified)
Peer Addr     : 192.0.2.3
Keep Alive Intvl: 1.0 secs      Hold on Nbr Fail      : 3
Discovery Intvl : 300 secs      Discovery Boot Intvl  : 300 secs
BFD           : Enable
Last update   : 02/13/2023 09:38:21
=====
Multi-Chassis IPsec Multi Active Tunnel-Group Table
=====
ID            Peer Group  Priority  Admin State  Mastership
-----
1             1           150      Up           eligible
-----
Multi Active Tunnel Group Entries found: 1
=====
```

The VRRP state on SeGW-4 is primary (**master**):

```
*A:SeGW-4# show router vrrp instance
=====
VRRP Instances
=====
Interface Name          VR Id Own  Adm  State      Base Pri  Msg Int
                        IP      Opr  Pol Id      InUse Pri  Inh Int
-----
int-SeGW-4-S-2         10   No   Up   Master    100     1
                        IPv4   Up   n/a      100     No
      Backup Addr: 172.16.1.254
-----
Instances : 1
=====
```

When SeGW-3 comes up, the IPsec state of tunnel group 1 is **discovery**, which means that the system has not established the MIMP session with its peer yet.

```
A:SeGW-3# show redundancy multi-chassis mc-ipsec peer 192.0.2.4
=====
Multi-Chassis MC-IPsec
=====
Peer Name      : (Not Specified)
Peer Addr     : 192.0.2.4
Keep Alive Intvl: 1.0 secs      Hold on Nbr Fail      : 3
Discovery Intvl : 300 secs      Discovery Boot Intvl  : 300 secs
BFD           : Enable
Last update   : 02/10/2023 12:17:46
=====
Multi-Chassis IPsec Multi Active Tunnel-Group Table
=====
ID            Peer Group  Priority  Admin State  Mastership
-----
```



```

1          1          200      Up          discovery
-----
Multi Active Tunnel Group Entries found: 1
=====

```

After a while, the preceding **show** command is repeated and the IPSec state for tunnel 1 on SeGW-3 is standby:

```

A:SeGW-3# show redundancy multi-chassis mc-ipsec peer 192.0.2.4
=====
Multi-Chassis MC-IPsec
=====
Peer Name      : (Not Specified)
Peer Addr     : 192.0.2.4
Keep Alive Intvl: 1.0 secs      Hold on Nbr Fail      : 3
Discovery Intvl: 300 secs      Discovery Boot Intvl  : 300 secs
BFD           : Enable
Last update   : 02/13/2023 09:46:03
=====
Multi-Chassis IPsec Multi Active Tunnel-Group Table
=====
ID            Peer Group  Priority  Admin State  Mastership
-----
1             1           200      Up           standby
-----
Multi Active Tunnel Group Entries found: 1
=====

```

The VRRP state on SeGW-3 is backup:

```

A:SeGW-3# show router vrrp instance
=====
VRRP Instances
=====
Interface Name          VR Id Own  Adm  State      Base Pri  Msg Int
                       IP      Opr  Pol Id      InUse Pri  Inh Int
-----
int-SeGW-3-5-2         10   No  Up  Backup    200      1
                       IPv4   Up  1           50       No
    Backup Addr: 172.16.1.254
-----
Instances : 1
=====

```

SeGW-4 is the primary chassis in MC-IPSec tunnel group 1:

```

*A:SeGW-4# show redundancy multi-chassis mc-ipsec peer 192.0.2.3
=====
Multi-Chassis MC-IPsec
=====
Peer Name      : (Not Specified)
Peer Addr     : 192.0.2.3
Keep Alive Intvl: 1.0 secs      Hold on Nbr Fail      : 3
Discovery Intvl: 300 secs      Discovery Boot Intvl  : 300 secs
BFD           : Enable
Last update   : 02/13/2023 09:38:21

```

```

=====
Multi-Chassis IPsec Multi Active Tunnel-Group Table
=====
ID          Peer Group    Priority  Admin State  Mastership
-----
1           1             150      Up           master
-----
Multi Active Tunnel Group Entries found: 1
=====

```

SeGW-4 is the primary node for VRRP instance 10:

```

*A:SeGW-4# show router vrrp instance
=====
VRRP Instances
=====
Interface Name          VR Id  Own  Adm  State      Base Pri  Msg Int
                       IP      Opr  Pol Id      InUse Pri  Inh Int
-----
int-SeGW-4-S-2         10     No   Up   Master    100      1
                       IPv4    Up   n/a         100      No
    Backup Addr: 172.16.1.254
-----
Instances : 1
=====

```

Configuration guidelines

The following is a list of guidelines for configuring MC-IPSec:

- To avoid high CPU load and issues in some complex cases, the following are suggestions for configuring the IKEv2 lifetime:
 - Both IKE_SA and CHILD_SA lifetime on MC-IPSec chassis (SeGW-3 and SeGW-4) should be around three times larger than on the IPsec peer (CE-1).
 - With the first rule, the lifetime of the side with smaller lifetime (IPSec peer CE-1) should not be too small (these being the default values):
 - IKE_SA: >= 86400 seconds
 - CHILD_SA: >= 3600 seconds
 - With the first rule, on the side with smaller lifetime (IPSec peer CE-1), the IKE_SA lifetime must be at least 3 times larger than CHILD_SA lifetime.
- The IKE protocol is the control plane of IPSec, so IKE packets must be treated as high QoS priority in the end-to-end path of the public service. On the public interface, a SAP ingress QoS policy must be configured to ensure that IKE packets get high QoS priority.
- Configure **ipsec-responder-only** under **tunnel-group** for static LAN-to-LAN tunnels.
- Enable dead peer detection (DPD) on the IPSec peer side (CE-1); disable DPD (default) on the MC-IPSec chassis side.
- The direct and redundant physical link between MC-IPSec chassis must be configured with sufficient bandwidth for MCS and shunting traffic, and proper QoS configuration to make sure the MIMP and MCS packets are treated as high priority traffic.

- The system time must be same on both MC-IPSec chassis.
- Make sure the protection status is **nominal** on both chassis before provoking a controlled switchover. The protection status can be displayed with the **show redundancy multi-chassis mc-ipsec peer <addr>** command.
- Wait at least five minutes between two consecutive switchovers if possible, to prevent a second switchover happening before the standby is ready to become the primary chassis.

Conclusion

MC-IPSec provides a stateful multi-chassis IPSec redundancy solution. This is very important in a carrier grade network, especially in applications such as mobile backhaul where high value mobile services run over IPSec tunnels.

NAT Stateless Dual-Homing

This chapter describes NAT stateless dual-homing.

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.R4.

Overview

With the IPv4 address space almost consumed, many operators are deploying Network Address Translation (NAT) at centralized or semi-centralized points in their IP/MPLS networks. The NAT function is implemented using Carrier Grade NAT (CGN) nodes, which typically support tens of thousands of clients/subscribers. Therefore, a failure of one of these nodes would be considered a significant event.

Many operators consider a stateful failover mechanism between CGN nodes to be too demanding with regard to control plane requirements and state synchronization of NAT bindings. A reasonable compromise appears to be a stateless failover mechanism, capable of providing failover between geo-redundant CGN devices, but with manageable control plane implications.

This chapter describes the NAT stateless dual-homing feature supported in SR OS. NAT stateless dual-homing is supported for Large-Scale NAT and NAT64, and this chapter describes how both can be supported, either independently or together.


Example Topology

The topology shown in [Figure 80: Example Topology](#) is an example of the use of NAT stateless dual-homing. PE-2, PE-4, RR-10, CGN-1, and CGN-2 form part of Autonomous System 64496 and run IS-IS level 2 and LDP. PE-2, PE-4, CGN-1, and CGN-2 are clients of Route Reflector RR-10 and peer with the VPN-IPv4 and VPN-IPv6 address families.

IAR-1 acts as an Internet Service Provider (ISP) edge router belonging to AS 65535, and provides Internet access to AS 65596. CGN-1 and CGN-2 are configured with a VPRN (1000) that serves as a NAT outside routing VPN Routing and Forwarding (VRF) instance. In this VRF, both CGN-1 and CGN-2 peer with IAR-1 in EBGP for the IPv4 address family. CGN-1 and CGN-2 advertise the relevant NAT outside pools to IAR-1, and IAR-1 advertises a default route to both CGN-1 and CGN-2. IAR-1 has IP connectivity to a web server at 10.1.128.43/24.

CE-2 is connected to PE-2 and is part of VPRN 200. CE-2 has an IPv4 host (172.31.102.3) that serves to test NAT44 connectivity to the web server, and an IPv6 host (2001:DB8:4001:200::3) that serves to test NAT64 connectivity to the web server.

Figure 80: Example Topology



26075

Configuration

To support NAT functionality, some form of Integrated Services Adaptor (ISA) or Integrated Services Module (ISM) is required, as listed in the Applicability section. In this example, CGN-1 and CGN-2 both have a single MS-ISA card, which is configured as MDA type **isa-bb**, as follows:

```
configure
card 1
  card-type iom3-xp-b
  mda 2
    mda-type isa-bb
    no shutdown
  exit
no shutdown
exit
```

The MS-ISA is then configured to become a member of one or more **nat-group**. Up to fourteen MS-ISAs can be configured to belong to up to four NAT groups. When more than one MS-ISA is configured in a NAT group, the MS-ISAs can work in active-standby mode or active-active mode.

In active-standby mode, one or more MS-ISAs act as standby, and in normal operation, are idle. If an active MS-ISA fails, one standby MS-ISA accepts the traffic from the failed card. In the active-standby scenario, the mapping between failed card and standby card can always be considered to be 1:1. In active-active mode, all of the MS-ISAs in the NAT group are active, and if one MS-ISA in the group fails, the load is distributed across the remaining active MS-ISAs in the group.

The default setting is active-standby. In both active-standby and active-active modes, the dynamically created NAT bindings are not synchronized between cards. Therefore, a failover will cause an interruption in traffic until the NAT bindings are re-initiated by the clients/subscribers behind the NAT.

The NAT group also requires that an **active-mds-limit** is configured, as follows, which allows the operator to specify how many MDAs (MS-ISAs) will be active in the group. Any operational MDAs above this configured limit will be considered spare MDAs. Finally, the **nat-group** must be placed in a **no shutdown** state.

```
configure
  isa
    nat-group 1 create
      active-mds-limit 1
      mda 1/2
      no shutdown
    exit
  exit
```

NAT Outside Context

The NAT **outside** function is responsible for creation of the NAT bindings, using outside IP addresses defined in outside pools (together with their associated ports), and for advertising the address ranges in those pools to upstream routers. The NAT stateless dual-homing redundancy mechanism is based on ownership of an outside pool, where each member of a redundant pair can assume either an active (master) or standby role for an outside pool. This active/standby role is determined by the presence of a **monitor** prefix. Both CGN nodes of a redundant pair implement the following:

1. Advertise a unique route into the routing instance that the NAT outside function resides in. This is known as the export route and may be advertised into the Global Routing Table (GRT) or a VPRN instance. For example, CGN-1 advertises (exports) prefix P1 while CGN-2 advertises prefix P2.
2. Check for the presence of a configured route in the routing instance that the NAT outside function resides in. This is known as the monitor route. Continuing the preceding example, CGN-1 monitors prefix P2 while CGN-2 monitors prefix P1.

Therefore, the export route of CGN-1 becomes the monitor route of CGN-2, and the export route of CGN-2 becomes the monitor route of CGN-1. The redundancy mechanism thereafter checks the (virtual) routing table of the NAT outside function for the presence of the monitor route. If it is not present, the redundancy state for the pool is set to active and the following occurs (subject to routing policy):

1. The redundancy export route is populated in the NAT outside routing instance and advertised externally in the relevant routing instance.
2. The outside pool address is populated in the NAT outside routing instance and advertised externally in the relevant routing instance.
3. Routes that need to become active in any associated NAT inside routing instances, to attract traffic to the active CGN node, are populated in the relevant routing tables. For example, NAT64 translator routes and/or NAT44 steering routes, both of which are described later in this chapter, are populated.

If the monitor route is present, the redundancy state for the pool is set to standby and the following occurs:

1. The redundancy export route is not populated in the NAT outside routing instance and is, therefore, not eligible to be advertised externally.
2. The outside pool address is not populated in the NAT outside routing instance and is, therefore, not eligible to be advertised externally.

3. Routes that need to become active in any associated NAT inside routing instances, to attract traffic to the active CGN node, are not populated in the relevant routing tables.

There are no configurable options for selection of active/standby CGN nodes. The status of a node is based on the presence of the monitor route. In the event of a collision during redundancy startup, hardcoded debounce timers ensure that only a single CGN node is selected as active.

The first of two following examples shows the configuration of VPRN 1000 (the NAT outside VRF) at CGN-2, with the second configuration showing the **vrf-export** and **vrf-import** policy statements. For advertising and importing routes from/to the VRF, there are two requirements: to advertise the export redundancy route and to import the monitor route. This is the purpose of the *"vrf1000-export"* and *"vrf1000-import"* policy statements. Other VPRN parameters are generic and, therefore, not discussed here.

The VPRN contains an interface, "to-IAR-1", which has an associated EBGP peering session to IAR-1. The export policy under the BGP **neighbor** context will be described later in this chapter. That policy contains sufficient logic to advertise the NAT outside pools.

The NAT pools are configured in the **nat>outside** context. The first configuration provides an example of a single pool, "4-to-4", which will be used for NAT44. As well as a name, the pool requires association with a NAT group, and definition of the **type** of NAT that will be configured; in this case, **large-scale**. The **mode** of the pool is set to **naft** to indicate N:1 NAT, and the **address-range** that will be used for outside addressing is 10.1.4.1 to 10.1.4.254.

The relevant and required parameters for stateless dual-homing are configured in the **redundancy** context. In this example, CGN-2 exports prefix 192.168.0.249/32 and monitors prefix 192.168.0.248/32. Conversely, CGN-1 exports prefix 192.168.0.248/32 and monitors prefix 192.168.0.249/32. The redundancy node and the pool must be placed into a **no shutdown** state.

```
*A:CGN-2#
configure
  service
    vprn 1000 customer 1 create
      vrf-import "vrf1000-import"
      vrf-export "vrf1000-export"
      autonomous-system 64496
      route-distinguisher 64496:1000
      auto-bind-tunnel
        resolution any
      exit
    interface "to-IAR-1" create
      address 172.31.19.2/30
      sap 1/1/2:200 create
      exit
    exit
  aggregate 10.1.4.0/24 summary-only
  bgp
    group "EBGP"
      family ipv4
      peer-as 65535
      split-horizon
      neighbor 172.31.19.1
        authentication-key <password>
        export "vrf1000-ebgp-export"
      exit
    exit
  no shutdown
exit
nat
  outside
    pool "4-to-4" nat-group 1 type large-scale create
      redundancy
```

```

        export 192.168.0.249/32
        monitor 192.168.0.248/32
        no shutdown
    exit
    mode napt
    address-range 10.1.4.1 10.1.4.254 create
    exit
    no shutdown
    exit
    exit
    exit
    service-name "NAT-Outside"
    no shutdown
    exit
    exit

```

```

*A:CGN-2#
configure
router
    policy-options
    begin
    prefix-list "vrf1000-nat-export"
        prefix 192.168.0.249/32 exact
    exit
    community "vrf1000-export" members "target:64496:1000"
    community "vrf1000-import" members "target:64496:1000"
    policy-statement "vrf1000-export"
        entry 10
            from
                protocol nat
                prefix-list "vrf1000-nat-export"
            exit
            to
                protocol bgp-vpn
            exit
            action accept
                community add "vrf1000-export"
            exit
        exit
    exit
    policy-statement "vrf1000-import"
        entry 10
            from
                community "vrf1000-import"
            exit
            action accept
            exit
        exit
        default-action drop
    exit
    exit

```

After the redundancy node and pool are placed into a **no shutdown** state, the master and standby can be elected, based on the previously described criteria. In this example, CGN-2 becomes the active CGN node for the pool "4-to-4". This is shown using the following command, where the Active field shows true. Conversely, the same output at CGN-1 shows the Active field as false.

```

*A:CGN-2# show router 1000 nat pool "4-to-4"

```

```

=====
NAT Pool 4-to-4
=====

```



```

Description                : (Not Specified)
ISA NAT Group              : 1
Pool type                  : largeScale
Applications               : (None)
Admin state                : inService
Mode                       : napt
Port forwarding dyn blocks reserved : 0
Port forwarding range      : 1 - 1023
Port reservation           : 128 blocks
Block usage High Watermark (%) : 90
Block usage Low Watermark (%) : 20
Subscriber limit per IP address : 65535
Active
Deterministic port reservation : (Not Specified)
Last Mgmt Change           : 10/06/2016 11:29:41
=====

NAT address ranges of pool 4-to-4
=====
Range                                Drain Num-blk
-----
10.1.4.1 - 10.1.4.254                0
-----
No. of ranges: 1
=====

NAT members of pool 4-to-4 ISA NAT group 1
=====
Member                                Block-Usage-% Hi
-----
1                                      < 1          N
-----
No. of members: 1
=====

Dual-Homing
=====
Type                : Leader
Export route        : 192.168.0.249/32
Monitor route       : 192.168.0.248/32
Admin state         : inService
Dual-Homing State : Active
=====

Dual-Homing fate-share-group
=====
Router    Pool                Type
-----
vprn1000  4-to-4                Leader
-----
No. of pools: 1
=====

```

Although the entire 10.1.4.0/24 block is allocated for NAT outside addressing, the address range shown in the preceding output does not include the network address (10.1.4.0/24) or broadcast address (10.1.4.255/24). Therefore, the address range does not include the entire /24 prefix and has to be fragmented into a number of longer prefixes, known through protocol NAT, to cover the 10.1.4.1-10.1.4.255 range. This is shown in the route table of VPRN 1000, following, and is due to the whole subnet not being

defined in the address-range configuration. (If the address range was 10.1.4.0-10.1.4.255, there would be a single entry of 10.1.4.0/24 in the route table of VPRN 1000.)

```
*A:CGN-2# show router 1000 route-table 10.1.4.0/24 longer

=====
Route Table (Service: 1000)
=====
Dest Prefix[Flags]
Next Hop[Interface Name]
Type Proto Age Metric Pref
-----
10.1.4.0/24
Black Hole Blackh* Aggr 00h01m34s 130
0
10.1.4.1/32
NAT outside to mda 1/2 Remote NAT 00h01m34s 0
0
10.1.4.2/31
NAT outside to mda 1/2 Remote NAT 00h01m34s 0
0
10.1.4.4/30
NAT outside to mda 1/2 Remote NAT 00h01m34s 0
0
10.1.4.8/29
NAT outside to mda 1/2 Remote NAT 00h01m34s 0
0
10.1.4.16/28
NAT outside to mda 1/2 Remote NAT 00h01m34s 0
0
10.1.4.32/27
NAT outside to mda 1/2 Remote NAT 00h01m34s 0
0
10.1.4.64/26
NAT outside to mda 1/2 Remote NAT 00h01m34s 0
0
10.1.4.128/26
NAT outside to mda 1/2 Remote NAT 00h01m34s 0
0
10.1.4.192/27
NAT outside to mda 1/2 Remote NAT 00h01m34s 0
0
10.1.4.224/28
NAT outside to mda 1/2 Remote NAT 00h01m34s 0
0
10.1.4.240/29
NAT outside to mda 1/2 Remote NAT 00h01m34s 0
0
10.1.4.248/30
NAT outside to mda 1/2 Remote NAT 00h01m34s 0
0
10.1.4.252/31
NAT outside to mda 1/2 Remote NAT 00h01m36s 0
0
10.1.4.254/32
NAT outside to mda 1/2 Remote NAT 00h01m36s 0
0
-----
No. of Routes: 15
Flags: n = Number of times nexthop is repeated
B = BGP backup route available
L = LFA nexthop available
S = Sticky ECMP requested
=====
* indicates that the corresponding row element may have been truncated.
```

CGN-2 advertising all of these longer prefixes to the edge router of the ISP is not wanted. The ISP may even enforce a minimum /24 prefix length. Therefore, the preceding configuration of VPRN 1000 shows an **aggregate** command for the 10.1.4.0/24 prefix with the argument **summary-only**. When at least one of the more-specific prefixes in the 10.1.4.0/24 range is populated in the route table of VPRN 1000, the aggregate becomes active and can be used by the route policy for exporting to IAR-1, while suppressing the more-specific routes.

The following output shows that outside pool prefixes are not populated in the NAT outside routing context at the standby CGN node (CGN-1). Even with the aggregate command configured at both CGN nodes, the

aggregate prefix will only become active at the active CGN node. Therefore, only the active CGN node will advertise that aggregate prefix upstream.

```
*A:CGN-1# show router 1000 route-table 10.1.4.0/24 longer

=====
Route Table (Service: 1000)
=====
Dest Prefix[Flags]                               Type   Proto   Age      Pref
  Next Hop[Interface Name]                       Metric
-----
No. of Routes: 0
Flags: n = Number of times nexthop is repeated
      B = BGP backup route available
      L = LFA nexthop available
      S = Sticky ECMP requested
=====
```

So far, only one pool has been defined that could be used for stateless dual-homing for both NAT44 and NAT64. However, to allow for independent address management of each of these functions, a separate outside pool is created for each. The following configuration shows two pools in the **nat>outside** context of VPRN 1000: the pool "4-to-4", which is for NAT44 purposes, and the pool "6-to-4", which is for NAT64 purposes.

The address range defined in the "6-to-4" pool is 10.1.6.1 to 10.1.6.254. As with the range 10.1.4.1 to 10.1.4.254 in the "4-to-4" pool, an aggregate command is configured to advertise the 10.1.6.0/24 prefix, while suppressing the more-specific prefixes.

With the redundancy node in the "6-to-4" pool, there are instances where traffic from a NAT inside routing context may be mapped to multiple outside pools, which in a stateless dual-homed environment may cause the NAT function to fail. For example, assume a NAT outside context has two pools, P1 and P2, where pool P1 is active and pool P2 is standby. An active pool can trigger the advertisement of inside and outside prefixes, and traffic will be attracted to this CGN node. When traffic arrives, it may be mapped to pool P1 on the active CGN node, due to the NAT mapping criteria. However, traffic may also be mapped to pool P2, due to the mapping criteria; this traffic will fail because the pool P2 is active on the redundant CGN node.

To ensure that this traffic failure does not happen, a group of pools accessed by the same inside routing context must all be active on the same CGN node simultaneously. To achieve this, SR OS uses a Pool Fate Sharing Group (PFSG). The PFSG ensures that all co-located pools accessed by the same inside routing context are either active or standby; not a combination of both. This is achieved by having a *leader* pool and *follower* pools.

If the leader pool is active, all follower pools are active. If the leader pool is standby, all follower pools are standby. This is enabled in the **redundancy** context of the "6-to-4" pool using the **follow** command. The **follow** command configures the pool as a follower and allows the user to access the routing context and outside pool of the leader pool. In the following example, pool "4-to-4" is a leader pool and pool "6-to-4" is a follower pool, which always assumes the same state as that of the leader.

```
*A:CGN-2#
configure
  service
    vprn 1000 customer 1 create
      aggregate 10.1.4.0/24 summary-only
      aggregate 10.1.6.0/24 summary-only
      nat
        outside
          pool "4-to-4" nat-group 1 type large-scale create
          redundancy
```

```

        export 192.168.0.249/32
        monitor 192.168.0.248/32
        no shutdown
    exit
    mode napt
    address-range 10.1.4.1 10.1.4.254 create
    exit
    no shutdown
exit
pool "6-to-4" nat-group 1 type large-scale create
    redundancy
        follow router 1000 pool "4-to-4"
    exit
    mode napt
    address-range 10.1.6.1 10.1.6.254 create
    exit
    no shutdown
exit
exit
exit

```

The following output shows a PFSG with leaders and followers in the operational state of pool "6-to-4" at CGN-2. The pool state is active, but the Dual-Homing Type field indicates that this pool is a follower. Therefore, the state is derived from the state of the leader pool, which is pool "4-to-4" in router 1000. The output also contains a list of all the pools that are part of the same PFSG.

```
*A:CGN-2# show router 1000 nat pool "6-to-4"
```

```
=====
NAT Pool 6-to-4
=====
```

```

Description                : (Not Specified)
ISA NAT Group               : 1
Pool type                   : largeScale
Applications                : (None)
Admin state                 : inService
Mode                        : napt
Port forwarding dyn blocks reserved : 0
Port forwarding range       : 1 - 1023
Port reservation            : 128 blocks
Block usage High Watermark (%) : 90
Block usage Low Watermark (%)  : 20
Subscriber limit per IP address : 65535
Active                    : true
Deterministic port reservation : (Not Specified)
Last Mgmt Change            : 10/06/2016 13:09:07
=====

```

```
=====
NAT address ranges of pool 6-to-4
=====
```

| Range | Drain Num-blk |
|-----------------------|---------------|
| 10.1.6.1 - 10.1.6.254 | 0 |

```
-----
No. of ranges: 1
=====
```

```
=====
NAT members of pool 6-to-4 ISA NAT group 1
=====
```

| Member | Block-Usage-% Hi |
|--------|------------------|
| ----- | ----- |

```

1                                     < 1      N
-----
No. of members: 1
=====
Dual-Homing
=====
Type                                : Follower
Follow-pool                          : "4-to-4" router 1000
Dual-Homing State                    : Active
=====
Dual-Homing fate-share-group
-----
Router      Pool                      Type
-----
vprn1000    4-to-4                    Leader
vprn1000  6-to-4                    Follower
-----
No. of pools: 2
=====

```

NAT Policies

NAT policies allow for definition of NAT attributes such as:

- filtering behavior (endpoint-independent or address-and-port-dependent)
- NAT mapping timeouts
- per-user session/flow limits
- configuration of Application Level Gateway (ALG) protocols
- high/low resource watermarks.

These attributes are generic NAT configuration parameters that are beyond the scope of this chapter.

A NAT policy also references the routing context and name of the outside pool used for the creation of NAT bindings associated with the policy. Therefore, if multiple outside pools are needed, multiple NAT policies must also be used. The following shows the configuration of the required NAT policies at CGN-2.

Because two outside pools exist in VPRN 1000 (the pool "4-to-4" for NAT44 and the pool "6-to-4" for NAT64), two policies are created using the **nat-policy** parameter. The **nat-policy** "NAT44" uses the **pool** keyword to access the "4-to-4" pool in **router** 1000, while the **nat-policy** "NAT64" uses the same **pool** keyword to access the "6-to-4" pool in **router** 1000. In this example, the same outside routing context is used for both NAT policies, but the outside routing contexts can also be different for each policy.

```

configure
  service
    nat
      nat-policy "NAT44" create
        pool "4-to-4" router 1000
      exit
      nat-policy "NAT64" create
        pool "6-to-4" router 1000
      exit
    exit

```

NAT Inside Context

The NAT inside routing context is the interface toward the customer or end user. There can be multiple NAT inside routing contexts mapped to a single NAT outside context (the relationship can be 1:1 or N:1). This is possible even if overlapping addresses are used in two or more NAT inside routing contexts because the NAT flow mapping tuple consists of the parameters {routing-instance, inside-IP, inside-port} mapped to {routing-instance, outside-IP, outside-port}.

The NAT inside routing context is responsible for two main functions:

1. Diverting some or all traffic toward the NAT function (ISA board).
2. Attracting traffic that should be subject to NAT toward it. This should be conditional because only the master CGNAT node should attract traffic toward itself.

For diverting traffic toward the NAT function, there are two approaches:

1. The first approach is to use IP filters with **action nat** to divert matched traffic into the ISA. Traffic subject to NAT can have a different inside and outside routing context, or the same routing context can be used for both inside and outside.
2. The second approach is a routing-based approach using a **destination-prefix** in the **nat>inside** context. Any traffic with a destination address matching the defined destination-prefix is forwarded to the ISA for NAT. When the destination-prefix approach is used, different routing contexts must be used for inside and outside.

For NAT44, both the IP filter and destination-prefix approaches are permitted. For NAT64, the diversion to NAT is only supported using IPv6 filters. The example setup in this chapter consists of both NAT44 and NAT64. Therefore, for the purpose of standardization across the NAT44 and NAT64 functions, the IP filter-based approach is used for both.

In [Figure 80: Example Topology](#), CE-2 is connected to PE-2 and is part of VPRN 200. To provide Internet access with stateless NAT dual-homing to VPRN 200, it is extended to both CGN-1 and CGN-2 as a NAT inside VRF. The following shows the configuration of VPRN 200 at CGN-2 with a similar configuration also applied at CGN-1.

Because one of the main functions of the NAT inside routing context is to attract traffic toward the (active) CGN node, this is configured in the **nat>inside** context. The **redundancy** parameter provides a context for the configuration of NAT44 redundancy when the diversion to NAT is implemented with IP filters. In this **redundancy** context, the **peer** command is used to configure the address of the redundant peer (in this case CGN-1). If upstream traffic that is subject to NAT inadvertently arrives at the CGN node that is standby for the outside pool used for the NAT mapping, this parameter provides a forwarding address for that traffic. However, if destination-prefix based redirect to NAT is used instead of IP filters, only a **nat-policy** and **destination-prefix** need to be configured in the NAT inside routing context.

The **steering-route** command is optional. It allows for configuration of a (non-default) prefix/length that is only active in the routing table of the NAT inside routing context of the active CGN node. When this steering route is active in the routing table, it can either be advertised directly using the route-policy framework, or it can be used as an indirect next-hop to advertise some other prefix. This latter approach is used in the following configuration example, where the **steering-route** of 192.168.203.1/32 is used as an indirect next-hop for the **static-route-entry** of 0.0.0.0/0. This creates the following dependencies:

- If the CGN node is active, the steering route of 192.168.203.1/32 becomes active in the routing table of VPRN 200.
- When 192.168.203.1/32 is active, it becomes a valid indirect next-hop for 0.0.0.0/0, so this route also becomes active in the routing table.

- When the default route is active, it can be exported to the rest of the VPN using the route-policy framework.

In the first of three following configurations, the **vrf-export** command accesses a policy with the name *vrf200-export*. The second configuration shows the contents of that policy, where entry 10 accesses a prefix-list (*vrf200-lsn44-default*) containing the default route and advertises it into BGP-VPN, with the relevant Route-Target (*target:64496:200*) and Origin (*origin:64496:200*) Extended Communities attached.

The Origin Extended Community is used by the redundant CGN peer to drop the default route, as shown in entry 10 of the corresponding *vrf200-import* policy in the same configuration. The reason for dropping the default route at the standby CGN node is that the *vrf200-export* policy only requires that a default route is present in the routing table in order to source/advertise a default route itself. If the standby CGN node imports the default route from the active CGN node into the routing table, the standby will also attempt to advertise a default route, which is not wanted.

The **nat64** command provides the context to configure the NAT64 redundancy parameters. In the case of NAT64, the CGN node becomes a translator between the IPv6 and IPv4 address families and needs to advertise the NAT64 translator address that will be used by IPv6 clients to embed IPv4 addresses. In the first configuration, the address *2001:DB8:122:344::/96* is used as the NAT64 translator address. As with NAT44, the advertisement of this address is conditional and must only be advertised by the active CGN node. This is ensured because the prefix is only present in the routing table of the active CGN, not the standby CGN.

Entry 20 of the *vrf200-export* policy shown in the second configuration provides the relevant policy rules to ensure that this IPv6 prefix is advertised into BGP-VPN with the relevant Route-Target value when the prefix is present in the routing table.

The last parameter in the **nat>inside** context is the **nat-policy**, which is known as the default NAT policy and must exist in the **nat>inside** context. When multiple NAT policies are used in a single NAT inside routing context, the default NAT policy is used for any traffic that is not matched (using the destination-prefix for NAT44 or IPv4/IPv6 filters for NAT44/NAT64) and associated with an explicit NAT policy. The default NAT policy can reference a separate NAT policy, or it can reference a NAT policy that is already in use.

As previously described, the intention is to use IP filters to implement the diversion to NAT. The relevant IPv4 and IPv6 filter IDs (ID number 200 in both cases) are shown in the third configuration. The IPv4 filter has no match criteria in this example; it has **action nat** using **nat-policy "NAT44"**, which accesses the outside pool "4-to-4" in VPRN 1000. The IPv6 filter also has no match criteria and has **action nat**, but distinguishes between DSLite (DSLite is not supported for NAT stateless dual-homing) and NAT64, using the **nat-type** argument. The **nat-policy** that should be used is the policy "NAT64", which accesses the outside pool "6-to-4" in VPRN 1000. These IP filters need to match traffic that ingresses the redundant CGN nodes from the MPLS side of VPRN 200 and are, therefore, applied in the **network>ingress** context in the first configuration.

The remainder of the VPRN parameters are generic and are not explained here.

```
*A:CGN-2#
configure
  service
    vprn 200 customer 1 create
      vrf-import "vrf200-import"
      vrf-export "vrf200-export"
      route-distinguisher 64496:200
      auto-bind-tunnel
        resolution any
      exit
    exit
  static-route-entry 0.0.0.0/0
```

```

        indirect 192.168.203.1
        no shutdown
    exit
exit
nat
    inside
        nat-policy "NAT44"
        nat64
            prefix 2001:db8:122:344::/96
            no shutdown
        exit
        redundancy
            peer 192.0.2.8
            steering-route 192.168.203.1/32
        exit
    exit
exit
network
    ingress
        filter ip 200
        filter ipv6 200
    exit
exit
no shutdown
exit

```

```

*A:CGN-2#
configure
router
    policy-options
        begin
        prefix-list "vrf200-lsn44-default"
            prefix 0.0.0.0/0 exact
        exit
        prefix-list "vrf200-nat64-translator"
            prefix 2001:db8:122:344::/96 exact
        exit
        community "vrf200-soo" members "origin:64496:200"
        community "vrf200-export" members "target:64496:200"
        community "vrf200-import" members "target:64496:200"
        policy-statement "vrf200-export"
            entry 10
                from
                    prefix-list "vrf200-lsn44-default"
                exit
                to
                    protocol bgp-vpn
                exit
                action accept
                    community add "vrf200-soo" "vrf200-export"
                exit
            exit
            entry 20
                from
                    prefix-list "vrf200-nat64-translator"
                exit
                to
                    protocol bgp-vpn
                exit
                action accept
                    community add "vrf200-export"
                exit
        exit
    exit

```




```
exit
policy-statement "vrf200-import"
  entry 10
    from
      community "vrf200-soo"
    exit
    action reject
  exit
  entry 20
    from
      community "vrf200-import"
    exit
    action accept
  exit
exit
```

```
*A:CGN-2#
configure
  filter
    ip-filter 200 create
      entry 10 create
        action
          nat nat-policy "NAT44"
        exit
      exit
    exit
    ipv6-filter 200 create
      entry 10 create
        action
          nat nat-type nat64 nat-policy "NAT64"
        exit
      exit
    exit
  exit
exit
```

Verification of the Active CGN Node

After the configuration of the inside and NAT outside routing contexts, with the associated NAT policies, the state of the stateless redundant CGN nodes can be verified; see [Figure 81: Redundancy Status](#).

Figure 81: Redundancy Status



26100

The following two outputs show that the pool "4-to-4" is a leader at both CGN-1 and CGN-2 (for which pool "6-to-4" is a follower), and that CGN-2 is the active CGN node.

```
*A:CGN-2# show router 1000 nat pool "4-to-4" | match "Dual-Homing State" pre-lines 6
Dual-Homing
=====
Type                : Leader
Export route        : 192.168.0.249/32
Monitor route       : 192.168.0.248/32
Admin state         : inService
Dual-Homing State   : Active
```

```
*A:CGN-1# show router 1000 nat pool "4-to-4" | match "Dual-Homing State" pre-lines 6
Dual-Homing
=====
Type                : Leader
Export route        : 192.168.0.248/32
Monitor route       : 192.168.0.249/32
Admin state         : inService
Dual-Homing State   : Standby
```

Although the redundancy export route is populated in the NAT outside routing table and advertised externally by the active CGN node (if permitted by route-policy), it is not populated in the NAT outside routing table of the standby CGN, which is not advertised externally. The following two outputs show that CGN-2 advertises its export route (192.168.0.249/32) into IPv4 BGP-VPN, but CGN-1 does not advertise its own export route because the monitor route (192.168.0.249/32) is present.

```
*A:CGN-2# show router bgp routes vpn-ipv4 rd 64496:1000 hunt
---snip---
```

```

-----
RIB Out Entries
-----
Network       : 192.168.0.249/32
Nexthop       : 192.0.2.9
Route Dist.   : 64496:1000          VPN Label      : 262128
Path Id       : None
To            : 192.0.2.10
Res. Nexthop  : n/a
Local Pref.   : 100
Aggregator AS : None              Interface Name : NotAvailable
Atomic Aggr.  : Not Atomic        Aggregator     : None
AIGP Metric   : None              MED            : 0
Connector     : None
Community     : target:64496:1000
Cluster       : No Cluster Members
Originator Id : None              Peer Router Id : 192.0.2.10
Origin        : IGP
AS-Path       : No As-Path
Route Tag     : 0
Neighbor-AS   : N/A
Orig Validation: N/A
Source Class  : 0                  Dest Class     : 0
-----

```

```

*A:CGN-1# show router 1000 route-table 192.168.0.249/32
=====
Route Table (Service: 1000)
=====
Dest Prefix[Flags]          Type   Proto   Age      Pref
Next Hop[Interface Name]   Metric
-----
192.168.0.249/32           Remote BGP VPN 06d03h21m 170
192.0.2.9 (tunneled)      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:CGN-1# show router bgp routes vpn-ipv4 rd 64496:1000 hunt
---snip---
-----
RIB Out Entries
-----
-----

```

As well as the export/monitor routes, the outside pools are populated in the NAT outside routing context and advertised by the active CGN node. The outside pools are summarized as 10.1.4.0/24 (pool "4-to-4") and 10.1.6.0/24 (pool "6-to-4") using the **aggregate** command in VPRN 1000. Because the (more-explicit) NAT outside pool addresses are only populated in the route table of the active CGN node, the aggregate will also only be populated in the routing table of the active CGN node. Therefore, the following policy is applied to the EBGP peering session with IAR-1 at both CGN-1 and CGN-2. The output following the policy

example shows that CGN-2 advertises both of the NAT outside pools to IAR-1, while CGN-1 advertises no outside pool prefixes to IAR-1.

```
configure
router
  policy-options
  begin
    policy-statement "vrf1000-ebgp-export"
      entry 10
        from
          protocol aggregate
        exit
      to
        protocol bgp
      exit
      action accept
        origin igp
      exit
    exit
  default-action drop
  exit
exit
commit
```

```
*A:CGN-2# show router 1000 bgp neighbor 172.31.19.1 advertised-routes
```

```
=====
BGP Router ID:192.0.2.9      AS:64496      Local AS:64496
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
               l - leaked, x - stale, > - best, b - backup, p - purge
Origin codes  : i - IGP, e - EGP, ? - incomplete
```

```
=====
BGP IPv4 Routes
```

| Flag | Network Nextthop (Router) As-Path | LocalPref Path-Id | MED Label |
|------|---|----------------------|--------------|
| i | 10.1.4.0/24 172.31.19.2 64496 | n/a None | None - |
| i | 10.1.6.0/24 172.31.19.2 64496 | n/a None | None - |

```
-----
Routes : 2
=====
```

```
*A:CGN-1# show router 1000 bgp neighbor 172.31.18.1 advertised-routes
```

```
=====
BGP Router ID:192.0.2.8      AS:64496      Local AS:64496
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
               l - leaked, x - stale, > - best, b - backup, p - purge
Origin codes  : i - IGP, e - EGP, ? - incomplete
```

```
=====
BGP IPv4 Routes
```

| Flag | Network | LocalPref | MED |
|------|---------|-----------|-----|
|------|---------|-----------|-----|

| Nexthop (Router) As-Path | Path-Id | Label |
|-----------------------------|---------|-------|
| ----- | | |
| No Matching Entries Found | | |
| ===== | | |

In the NAT inside routing instance, any routes that are used to attract traffic are populated in the relevant route tables of the active CGN node and must be advertised externally. For the NAT44 function, the steering route 192.168.203.1/32 populates the route table of the active CGN node, and this route is used as an indirect next-hop for a static-route-entry to a default route. For the NAT64 function, the NAT64 translator address 2001:DB8:122:344::/96 is used to attract IPv6 traffic with IPv4-embedded addresses.

The first of the two following outputs shows that CGN-2 is advertising the NAT44 default route as a VPN-IPv4 prefix and the NAT64 translator address as a VPN-IPv6 prefix. Both routes are advertised with the relevant Route-Target for VPRN 200 (target:64496:200) and, therefore, will be imported by PE-2 and subsequently advertised to CE-2. The second output shows the same commands entered at CGN-1 and verifies that because CGN-1 is the standby CGN node, it is not advertising either VPN-IPv4/VPN-IPv6 prefix.

```
*A:CGN-2# show router bgp routes vpn-ipv4 rd 64496:200 hunt
=====
BGP Router ID:192.0.2.9      AS:64496      Local AS:64496
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
               l - leaked, x - stale, > - best, b - backup, p - purge
Origin codes  : i - IGP, e - EGP, ? - incomplete
=====
BGP VPN-IPv4 Routes
=====
---snip---
-----
RIB Out Entries
-----
Network       : 0.0.0.0/0
Nexthop       : 192.0.2.9
Route Dist.   : 64496:200          VPN Label    : 262141
Path Id       : None
To            : 192.0.2.10
Res. Nexthop  : n/a
Local Pref.   : 100
Aggregator AS : None              Interface Name : NotAvailable
Atomic Aggr.  : Not Atomic       Aggregator    : None
AIGP Metric   : None             MED           : None
Connector     : None
Community     : target:64496:200
Cluster       : No Cluster Members
Originator Id : None              Peer Router Id : 192.0.2.10
Origin        : IGP
AS-Path       : No As-Path
Route Tag     : 0
Neighbor-AS   : N/A
Orig Validation: N/A
Source Class  : 0                 Dest Class    : 0
-----

*A:CGN-2# show router bgp routes vpn-ipv6 rd 64496:200 hunt
=====
```

```

BGP Router ID:192.0.2.9      AS:64496      Local AS:64496
=====
Legend -
Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                l - leaked, x - stale, > - best, b - backup, p - purge
Origin codes  : i - IGP, e - EGP, ? - incomplete
=====
BGP VPN-IPv6 Routes
=====
---snip---
-----
RIB Out Entries
-----
Network       : 2001:db8:122:344::/96
NextHop       : ::ffff:192.0.2.9
Route Dist.   : 64496:200                VPN Label    : 262141
Path Id       : None
To            : 192.0.2.10
Res. NextHop  : n/a
Local Pref.   : 100                      Interface Name : NotAvailable
Aggregator AS : None                    Aggregator    : None
Atomic Aggr.  : Not Atomic              MED           : 0
AIGP Metric   : None
Connector     : None
Community     : target:64496:200
Cluster       : No Cluster Members
Originator Id : None                    Peer Router Id : 192.0.2.10
Origin        : IGP
AS-Path       : No As-Path
Route Tag     : 0
Neighbor-AS   : N/A
Orig Validation: N/A
Source Class  : 0                        Dest Class    : 0
-----

```

```

*A:CGN-1# show router bgp routes vpn-ipv4 rd 64496:200 hunt
=====
BGP Router ID:192.0.2.8      AS:64496      Local AS:64496
=====
Legend -
Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                l - leaked, x - stale, > - best, b - backup, p - purge
Origin codes  : i - IGP, e - EGP, ? - incomplete
=====
BGP VPN-IPv4 Routes
=====
---snip---
-----
RIB Out Entries
-----

*A:CGN-1# show router bgp routes vpn-ipv6 rd 64496:200 hunt
=====
BGP Router ID:192.0.2.8      AS:64496      Local AS:64496
=====
Legend -
Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                l - leaked, x - stale, > - best, b - backup, p - purge

```

```
Origin codes : i - IGP, e - EGP, ? - incomplete
```

```
=====
BGP VPN-IPv6 Routes
=====
```

```
---snip---
```

```
-----
RIB Out Entries
-----
-----
```

Verification of Data Path

The host connected to CE-2 and the web server accessible from IAR-1 are used to verify the end-to-end data path for both NAT44 and NAT64.

NAT44

The host connected to CE-2 initiates an IPv4 UDP session toward the web server with a source address of 172.31.102.3 and a destination address of 10.1.128.43. The source port used is 1357, and the destination port is 80.

A two-way data transfer is verified as successful. The following output shows the details of the NAT44 binding at CGN-2, the active CGN node. The inside IP address and port are as described, while the allocated outside IP address is 10.1.4.254 using outside port 1047.

```
*A:CGN-2# tools dump nat sessions inside-ip 172.31.102.3
```

```
=====
Matched 1 session on Slot #1 MDA #2
=====
```

```
Owner           : LSN-Host@172.31.102.3
Router          : 200
Policy          : NAT44
FlowType       : UDP           Timeout (sec)      : 300
Inside IP Addr  : 172.31.102.3
Inside Port    : 1357
Outside IP Addr : 10.1.4.254
Outside Port   : 1047
Foreign IP Addr : 10.1.128.43
Foreign Port   : 80
Dest IP Addr   : 10.1.128.43
Nat Group      : 1
Nat Group Member : 1
-----
=====
```

NAT64

The host connected to CE-2 also initiates an IPv6 UDP session toward the web server with a source address of 2001:DB8:4001:200::3 and a destination address of 2001:DB8:122:344::A01:802B. The destination address represents the NAT64 translator address (2001:DB8:122:344::/96) and the embedded IPv4 address (10.1.128.43) translated into colon-hexidecimal format (A01:802B). The source port is 2468 and the destination port is 80.

A two-way data transfer is verified as successful. The following output shows the details of the NAT64 binding at CGN-2. Again, the inside IPv6 address and port are as described, while the allocated outside IP address is 10.1.6.254 using outside port 1032.

```
*A:CGN-2# tools dump nat sessions inside-ip 2001:db8:4001:200::3

=====
Matched 1 session on Slot #1 MDA #2
=====
Owner           : NAT64-Sub@2001:db8:4001:200::3
Router          : 200
Policy          : NAT64
FlowType       : UDP                               Timeout (sec)   : 300
Inside IP Addr : 2001:db8:4001:200::3                Inside Port     : 2468

Outside IP Addr : 10.1.6.254
Outside Port    : 1032
Foreign IP Addr : 10.1.128.43
Foreign Port    : 80
Dest IP Addr   : 10.1.128.43
Nat Group      : 1
Nat Group Member : 1
-----
=====
```

Failover

Before simulating a failover test, an IPv4 UDP session is established between the host connected to CE-2 and the web server, to ensure data-path continuity during the failure.

CGN-2 is the active CGN node; to simulate a failure of the MS-ISA board, it is placed into a shutdown state.

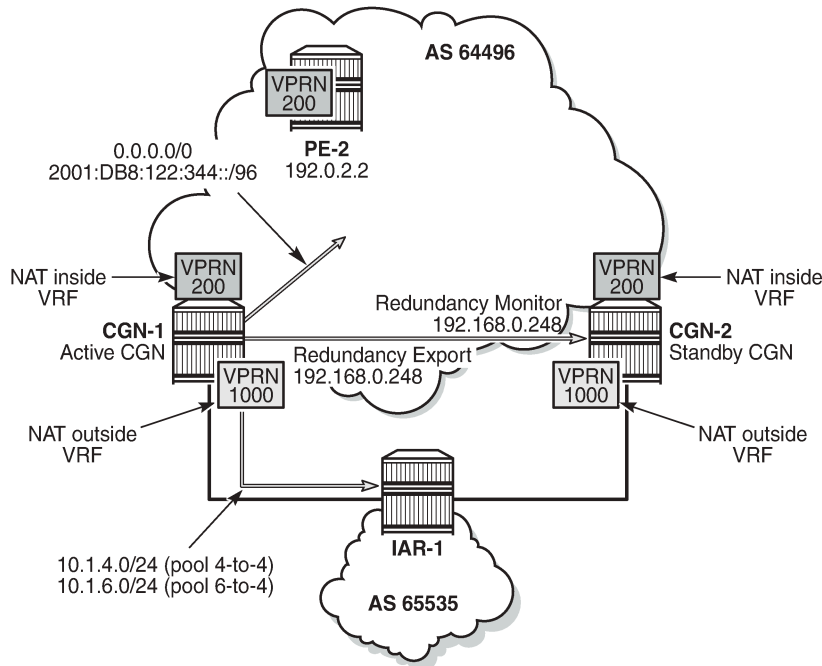
```
*A:CGN-2# configure card 1 mda 2 shutdown

1 2016/10/13 15:10:21.39 UTC WARNING: SNMP #2004 Base 1/2/nat-in-ip
"Interface 1/2/nat-in-ip is not operational"

2 2016/10/13 15:10:21.39 UTC MINOR: NAT #2024 Base NAT
"The state of NAT group 1 changed to out-of-service."
```

Figure 82: Post-Failover Redundancy State shows the example topology in the post-failover redundancy state.

Figure 82: Post-Failover Redundancy State



26101

Because the MS-ISA in slot 1/2 is the only MDA in **nat-group 1**, it is sufficient to force the NAT group down. After **nat-group 1** is declared down at CGN-2, the following actions take place:

1. CGN-2 transitions the active state to false for the leader outside pool and any follower pools because its admin state changed to down, as follows:

```
22 2016/10/13 15:10:21.40 UTC WARNING: NAT #2017 vprn1000 NAT redundancy
"The Large Scale NAT activity changed to false for pool "4-to-4" - state changed
to "Down"."

23 2016/10/13 15:10:21.40 UTC WARNING: NAT #2017 vprn1000 NAT redundancy
"The Large Scale NAT activity changed to false for pool "6-to-4" - state changed
to "Down"."
```

2. CGN-2 withdraws the redundancy export route (192.168.0.249/32) from the NAT outside routing context (VPRN 1000). This means that the monitor route is no longer present in the routing table of CGN-1. Therefore, CGN-1 transitions to an active state and advertises its own export route. In this example, where BGP is used to advertise monitor routes, the Minimum Route Advertisement Interval is configured for 1 second, to reduce re-convergence times. However, when the monitor route is withdrawn at the standby CGN node, the system will wait for 10 seconds to ensure that this is not a route flap before declaring itself active. This is a non-configurable timer.

```
6 2016/10/13 15:10:31.53 UTC WARNING: NAT #2017 vprn1000 NAT redundancy
"The Large Scale NAT activity changed to true for pool "4-to-4" - state changed
to "Active"."

8 2016/10/13 15:10:31.53 UTC WARNING: NAT #2017 vprn1000 NAT redundancy
"The Large Scale NAT activity changed to true for pool "6-to-4" - state changed
to "Active"
```

3. Because CGN-2 is now standby, the outside pool addresses are no longer present in the routing table of the outside routing context (VPRN 1000). Therefore, they are withdrawn by CGN-2 in the EBGP peering session to IAR-1. Conversely, because CGN-1 is now active, it advertises the outside pools to IAR-1.
4. The NAT44 redundancy steering route (192.168.203.1/32) is no longer active in the NAT inside (VPRN 200) routing table at CGN-2. Therefore, the default route no longer has a valid next-hop, so the route is withdrawn. Conversely, the steering route is now present in the routing table of VPRN 200 at CGN-1. Therefore, the default route becomes active, and is advertised into VPRN 200 as a VPN-IPv4 prefix.
5. The NAT64 translator address is no longer active in the NAT inside (VPRN 200) IPv6 routing table at CGN-2, so the address is withdrawn. Conversely, the NAT64 translator address is now present in the IPv6 routing table of VPRN 200 at CGN-1 and is advertised into VPRN 200 as a VPN-IPv6 prefix.

The operational state of the outside pools can be verified at CGN-1, as follows:

```
*A:CGN-1# show router 1000 nat pool "4-to-4" | match "Dual-Homing" post-lines 12
Dual-Homing
=====
Type                               : Leader
Export route                       : 192.168.0.248/32
Monitor route                      : 192.168.0.249/32
Admin state                        : inService
Dual-Homing State                  : Active
=====

Dual-Homing fate-share-group
=====
Router      Pool                Type
-----
vprn1000    4-to-4                  Leader
vprn1000    6-to-4                  Follower
-----
No. of pools: 2
=====
```

Finally, the integrity of the IPv4 UDP session between the host connected to CE-2 and the web server is verified, and the associated NAT binding is shown at CGN-1, as follows:

```
*A:CGN-1# tools dump nat sessions inside-ip 172.31.102.3
=====
Matched 1 session on Slot #1 MDA #2
=====
Owner           : LSN-Host@172.31.102.3
Router          : 200
Policy          : NAT44
FlowType        : UDP                Timeout (sec)      : 300
Inside IP Addr  : 172.31.102.3
Inside Port     : 1357
Outside IP Addr : 10.1.4.254
Outside Port    : 1049
Foreign IP Addr : 10.1.128.43
Foreign Port    : 80
Dest IP Addr    : 10.1.128.43
Nat Group       : 1
Nat Group Member : 1
-----
=====
```

When the failure is resolved at CGN-2 and the MS-ISA comes back online, the failover mechanism is non-revertive. This is because CGN-2 already has the CGN-1 export route present in the routing table of the NAT outside routing context (VPRN 1000) as its monitor route. The following output at CGN-2 shows the MS-ISA and NAT group 1 transitioning to in-service, followed by the active state of the outside pools changing from down to standby.

```
*A:CGN-2# configure card 1 mda 2 no shutdown

9 2016/10/13 15:55:10.39 UTC MINOR: NAT #2024 Base NAT
"The state of NAT group 1 changed to in-service."

10 2016/10/13 15:55:10.39 UTC MINOR: NAT #2025 Base NAT
"The NAT group 1 is not degraded."

17 2016/10/13 15:55:10.40 UTC WARNING: NAT #2017 vprn1000 NAT redundancy
"The Large Scale NAT activity changed to false for pool "4-to-4" - state changed
to "Standby".

18 2016/10/13 15:55:10.39 UTC WARNING: NAT #2017 vprn1000 NAT redundancy
"The Large Scale NAT activity changed to false for pool "6-to-4" - state changed
to "Standby".

19 2016/10/13 15:55:10.39 UTC MINOR: NAT #2020 Base NAT
"The NAT MDA 1/2 is now active in group 1."
```

Conclusion

NAT stateless dual-homing provides a compromise between a lack of redundancy and the protocol and state synchronization requirements for stateful NAT redundancy. This is particularly true when CGN nodes provide a gateway to the Internet where Service Level Agreements (SLAs) are often difficult to guarantee.

This chapter provides an example of how NAT stateless dual-homing is configured and describes how SR OS provides the redundancy mechanism for NAT44 and NAT64. The example in this chapter does not represent the only way that NAT stateless dual-homing can be delivered. It uses VPRNs in both the inside and outside routing contexts, but the GRT is also an option for either. It uses IP filtering for NAT diversion for both NAT44 and NAT64, but a routing-based approach using destination-prefix is also option for NAT44. It uses BGP in the NAT outside routing context and BGP-VPN in the NAT inside routing context to advertise redundancy routes externally, but any routing protocol that can be accessed through the route-policy framework is applicable.

Triple Play Service Delivery Architecture

This section provides TPSDA configuration information for the following topics:

- [ARP Hosts](#)
- [Bridged CO](#)
- [DHCP Server Failover States](#)
- [DHCPv4 Server Basics](#)
- [Diameter Application NASREQ](#)
- [Diameter Base Protocol: Establishing a Diameter Peer Connection](#)
- [ESM 128-bit Mode for DHCPv6 IA_NA WAN Hosts](#)
- [ESM Basics](#)
- [ESM IPv4: Multicast in a Wholesale/Retail Scenario](#)
- [ESM IPv4: Multicast with Redirection](#)
- [ESM IPv4: Multicast with SRRP](#)
- [ESM SLAAC Prefix Assignment via Local Address Server](#)
- [ESMv4: PPPoE Hosts](#)
- [ESMv6: IPoE Dual Stack Hosts](#)
- [ESMv6: PPPoE Dual Stack Hosts](#)
- [Flexible Authentication Model in ESM](#)
- [GTP Access](#)
- [High Scale QoS IOM in ESM Context: Expanded SLA Mode](#)
- [High Scale QoS IOM in ESM Context: Single SLA Mode](#)
- [Ingress Multicast Path Management](#)
- [IPoE Sessions](#)
- [IPv4 DHCP Hosts](#)
- [L2TP for Subscriber Access — LAC](#)
- [Local User Database Basics](#)
- [Local User Database for DHCPv4 Server](#)
- [Local User Database for Enhanced Subscriber Management](#)
- [Managed SAPs with Routed CO](#)
- [Multi-Chassis Ring Layer 2 with Enhanced Subscriber Management](#)
- [Python Cache Support for ESM Applications](#)
- [RADIUS-Triggered Dynamic Data Service Provisioning](#)
- [Raw Formatting of DHCPv4/v6 Options in ESM](#)

- [Routed CO](#)
- [Subscriber Redundancy for Routed CO](#)
- [Virtual Residential Gateway Authentication Scenarios](#)
- [Virtual Residential Gateway Home LAN Extension](#)
- [Virtual Residential Gateway Home Pool Management](#)
- [WiFi Aggregation and Offload — Basic Open SSID](#)
- [WiFi Aggregation and Offload — Basic Secure SSID with Distributed RADIUS Proxy](#)
- [WiFi Aggregation and Offload — IPv4/v6 Dual-Stack UEs](#)
- [WiFi Aggregation and Offload — Migrant User Support](#)
- [WiFi Aggregation and Offload — Open SSID with DSM and Lawful Intercept](#)

ARP Hosts

This chapter describes advanced ARP host configurations.

Topics in this chapter include:

- [Applicability](#)
- [Summary](#)
- [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.

Summary

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 section 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 document.



WARNING:

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

Overview

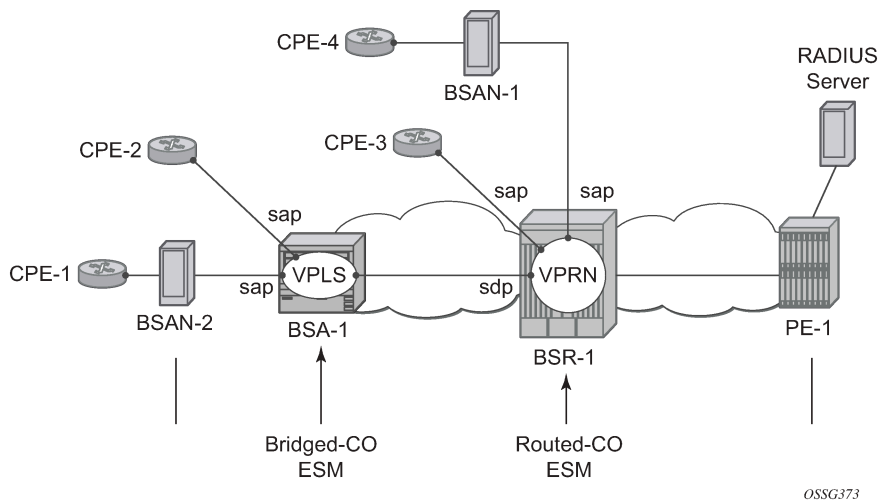
ARP host is supported in the bridged CO (VPLS) and the routed CO (Subscriber Interface) model. It is triggered by the first ARP packet received from the host. ARP host is also supported in a wholesale/retail context and on managed SAPs (MSAP).

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.

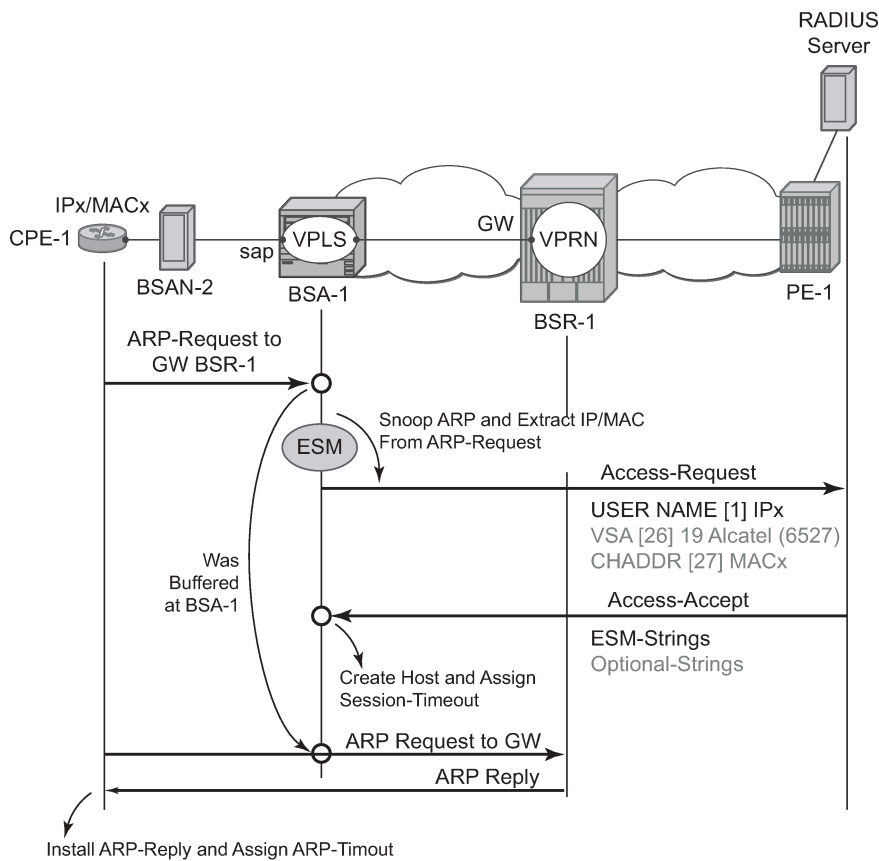
Figure 83: Bridged CO and Routed CO Example



Configuration

ARP Hosts in a Bridged CO Environment

Figure 84: 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
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.

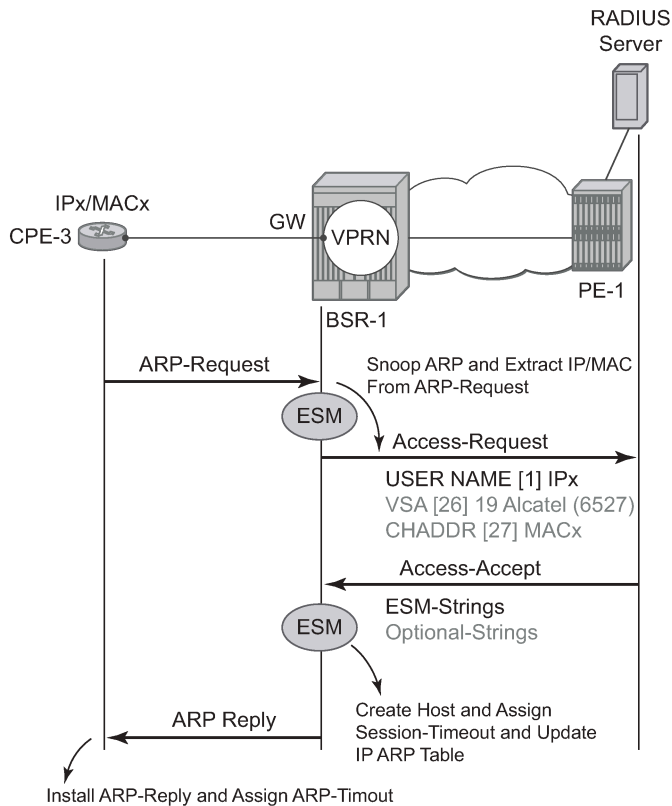
```

# 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 85: 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
    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 IP1 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 gy3yhtT5dF9YYilHtiinnk
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: SVCAGR #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
                00:00:0a:01:00:01 1/1/1:1         Time           Stdbby
-----
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
                  Time                               Stdbby
-----
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 12: 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**.

```

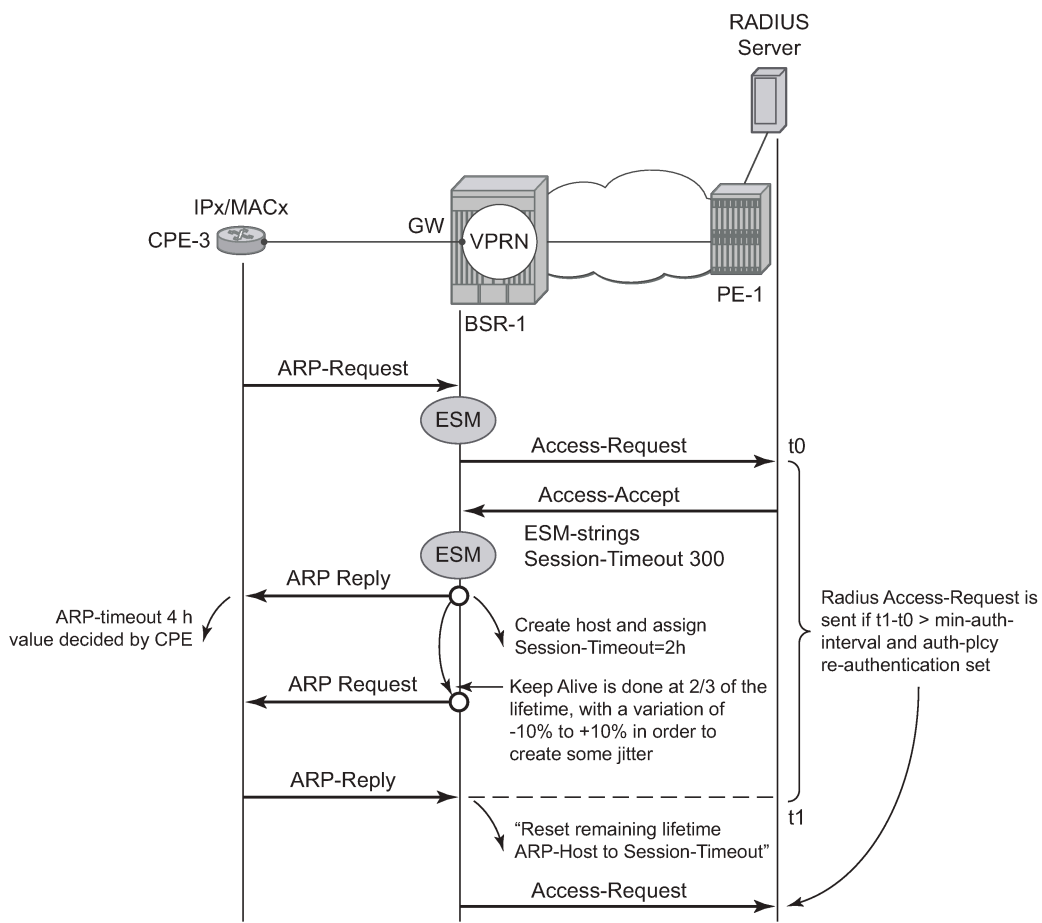
A:BSR-1# show service id 1 arp 10.1.0.1
=====
ARP Table
=====
IP Address      MAC Address      Type    Expiry    Interface    SAP
-----
10.1.0.1       00:00:0a:01:00:01 Managed 00h00m00s group-int-1  1/1/1:1
=====
A:BSR-1#
    
```

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 86: 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 86: 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
"

```

```

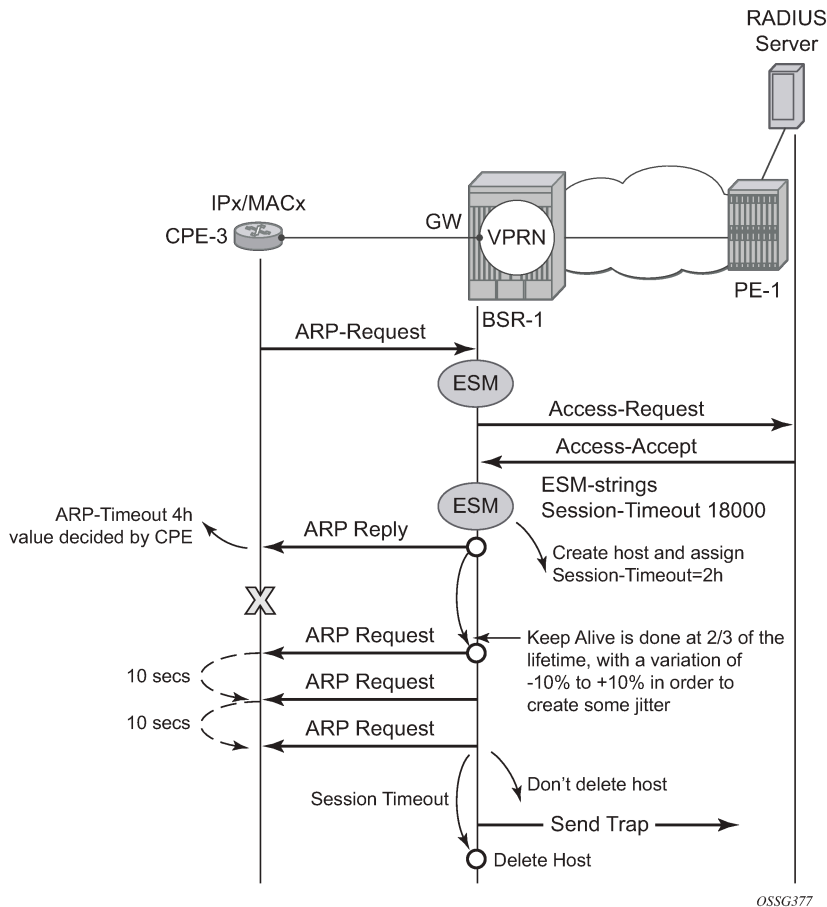
A:BSR-1# show service id 1 arp-host statistics
=====
ARP host statistics
=====
Num Active Hosts          : 1
Received Triggers        : 3
Ignored Triggers         : 0
Ignored Triggers (overload) : 0
SHCV Checks Forced       : 1
Hosts Created            : 1
Hosts Updated            : 2
Hosts Deleted            : 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 87: Trap Generation Example](#).

Figure 87: Trap Generation Example



```

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

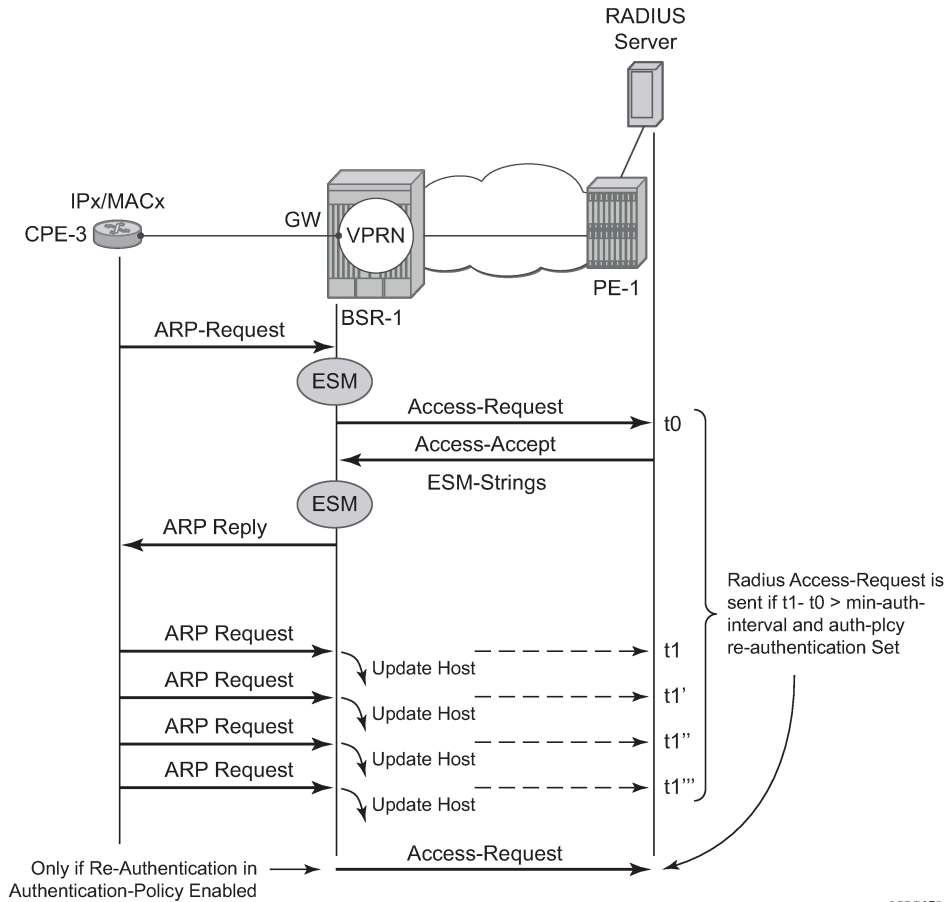
 IP: 10.1.0.1
 MAC: 00:00:0a:01:00:01
"
```

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 88: Throttling Toward RADIUS Example



OSSG378

```
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
```

```

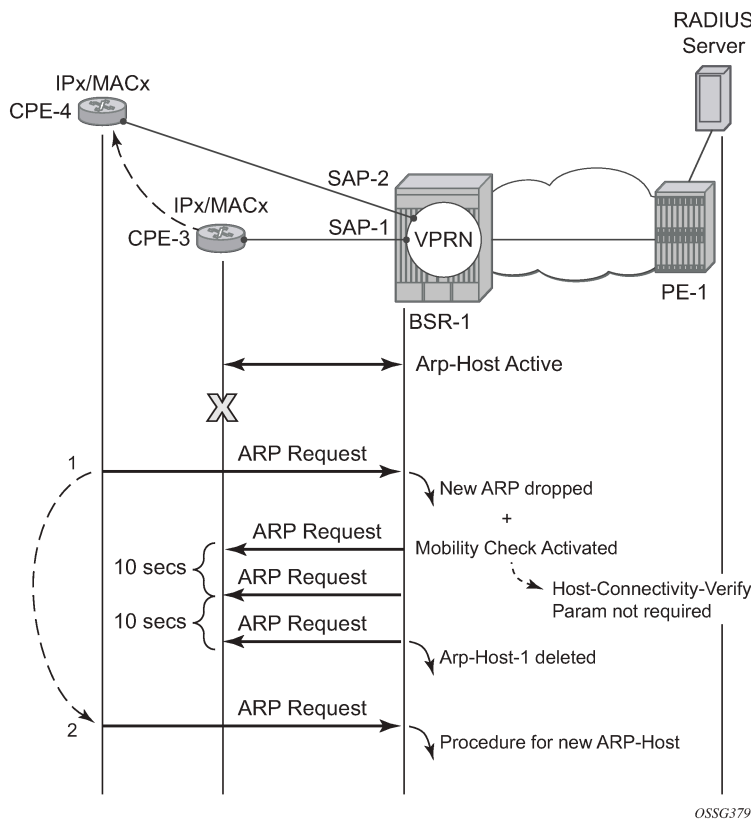
Last Auth       : 06/22/2015 15:59:32
Last Refresh   : 06/22/2015 16:00:33
Persistence Key : N/A
-----
Number of ARP hosts : 1
=====
A:BSR-1#
    
```

ARP Host Mobility

In order for ARP host mobility to function, host-connectivity-verification must **not** be enabled. This is different when compared with DHCP host mobility.

The implementation for routed CO is displayed in [Figure 89: ARP Host Mobility Example](#) and works the same for bridged CO. The **mac-pinning** command in routed CO context has no influence on this behavior.

Figure 89: ARP Host Mobility Example



OSSG379

ARP Host Persistency

ARP hosts can be made persistent across reboots and do not differ with other host types such as DHCP hosts.

```

configure
system
    
```

```

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 : 02DAFF0000008558814C4
Multi-Sess-Id : 02DAFF0000009558814C4
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: SVCNMR #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.

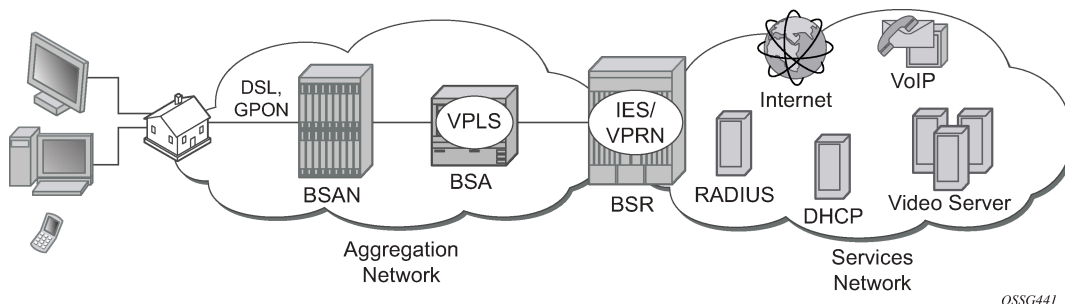
Knowledge of the Triple Play Service Delivery Architecture (TPSDA) concepts is assumed throughout this chapter.

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 90: Bridged CO Network Topology](#).

Figure 90: Bridged CO Network Topology



Following types of nodes are defined in Bridged CO model:

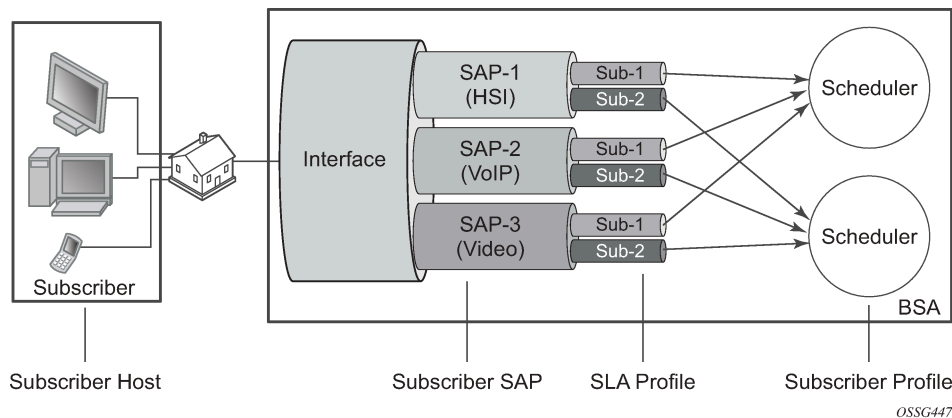
- Broadband Service Access Node (BSAN) — Access node connected to Layer 2 domain to aggregate all traffic from subscribers (IP DSLAM, ethernet switch).

- 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 91: 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 91: 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"
```

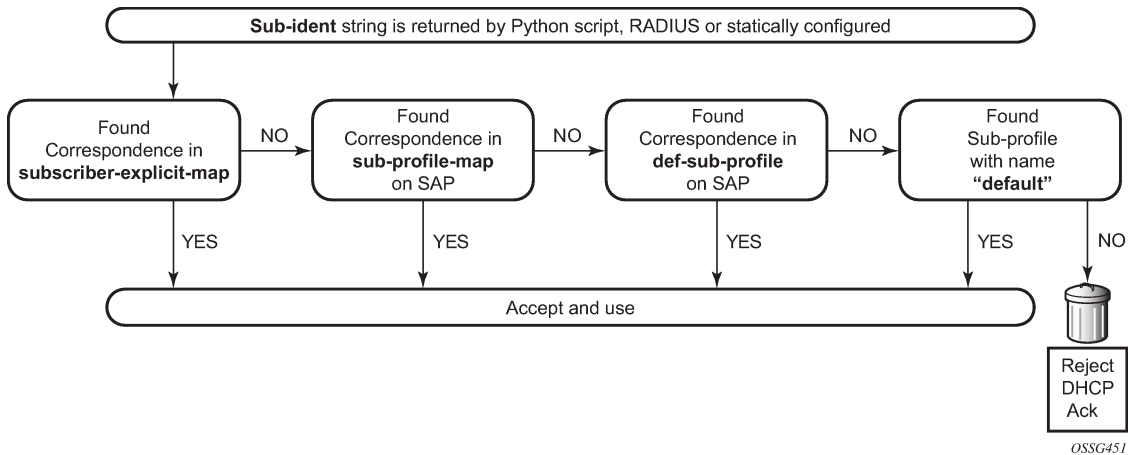
- c. If provisioned, the sub-profile-name is taken from the def-sub-profile attribute on the SAP. If no entry was found go to [1.d](#).

```
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 92: 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"
  
```

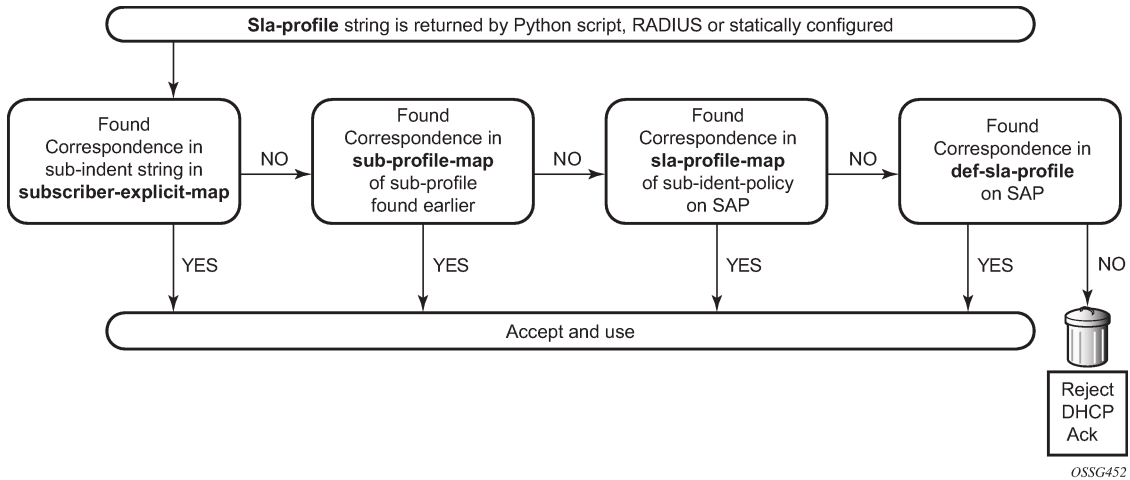
- d. The sla-profile-name is taken from the def-sla-profile attribute on the SAP. If no entry was found DHCP Ack is dropped.

```

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

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

Figure 93: 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 13: Correlation of Hosts and BSA/BSR Services](#) .

Table 13: 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 94: Sample Topology](#).


Figure 94: Sample Topology



OSSG446

Bridged CO model requires certain techniques and features to be used on different nodes. Major methods are presented in [Figure 95: Functionality of Each Node](#).

Figure 95: 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 14: 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 15: 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 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 |

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 16: 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
```


Figure 96: 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 SVCNMR
"SVCMGR: 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 97: Host-2 Setup Process](#).

Figure 97: 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 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: Discover
---snip---
"
23 2009/12/15 13:00:36.35 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: Offer
---snip---
"
24 2009/12/15 13:00:36.34 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: Offer
---snip---
"
25 2009/12/15 13:00:36.46 UTC MINOR: DEBUG #2001 Base SVCNMR
"SVCNMR: 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 98: Host-3 Setup Process](#).

Figure 98: 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 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: Discover
  ---snip---
  "
34 2009/12/15 13:02:34.38 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: Discover
  ---snip---
  "
35 2009/12/15 13:02:34.40 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: Offer
---snip---
"
36 2009/12/15 13:02:34.40 UTC MINOR: DEBUG #2001 Base SVCMMGR
"SVCMMGR: 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 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: Request
---snip---
"
38 2009/12/15 13:02:34.52 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: Request
---snip---
"
39 2009/12/15 13:02:34.53 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: Ack
---snip---
"
40 2009/12/15 13:02:34.53 UTC MINOR: DEBUG #2001 Base SVCMMGR
"SVCMMGR: 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  SapId/      DestIp      Timestamp      Time since  Oper
Id   SdpId       Address     last-reply/conn-lost  Reply/Lost  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 99: 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 99: General Redundancy Model




26352

Figure 100: 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 99: General Redundancy Model), but can also run over backbone links if no direct link is present (scenario 2 in Figure 99: 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 100: 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 101: 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 101: 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

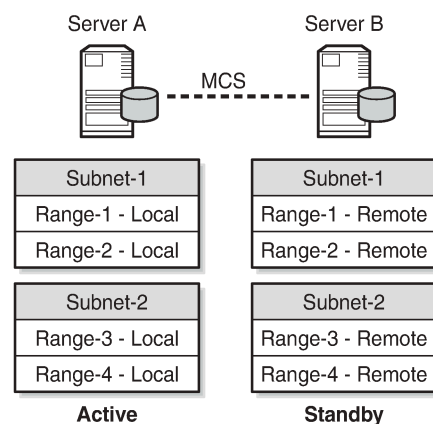
The first example of the local-remote model is an active-standby configuration; see [Figure 102: Local-Remote Model – Active-Standby](#). All the ranges in all pools are declared local on DHCP server A (the active server), and remote on DHCP server B (the standby server).

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 102: 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 103: 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 103: Local-Remote Model – Subnet-Based Load Sharing



26356

A variant of this model is shown in [Figure 104: 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 104: 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 104: 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 105: 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 105: 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 106: 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 106: 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
```

```
# 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     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
  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 107: 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 107: 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.

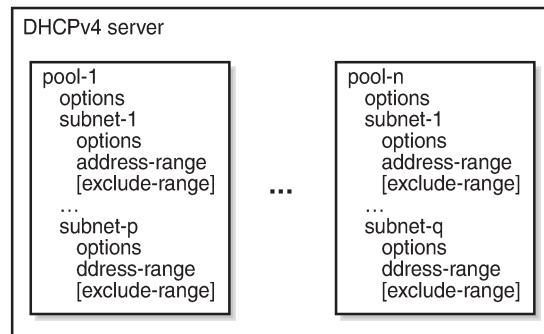
```
tools perform router <router-id> dhcp local-dhcp-server <server-name>
pool <pool-name> create-sticky-lease <hostname>
    [mac <ieee-address>]
    [circuit-id <circuit-id>]
    [client-id <client-id>]
    [requested-ip-address <ip-address>]
    [circuit-id-hex <circuit-id-hex-string>]
    [client-id-hex <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 108: Addresses, Subnets, and Pools in a DHCPv4 Server](#).

Figure 108: 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.

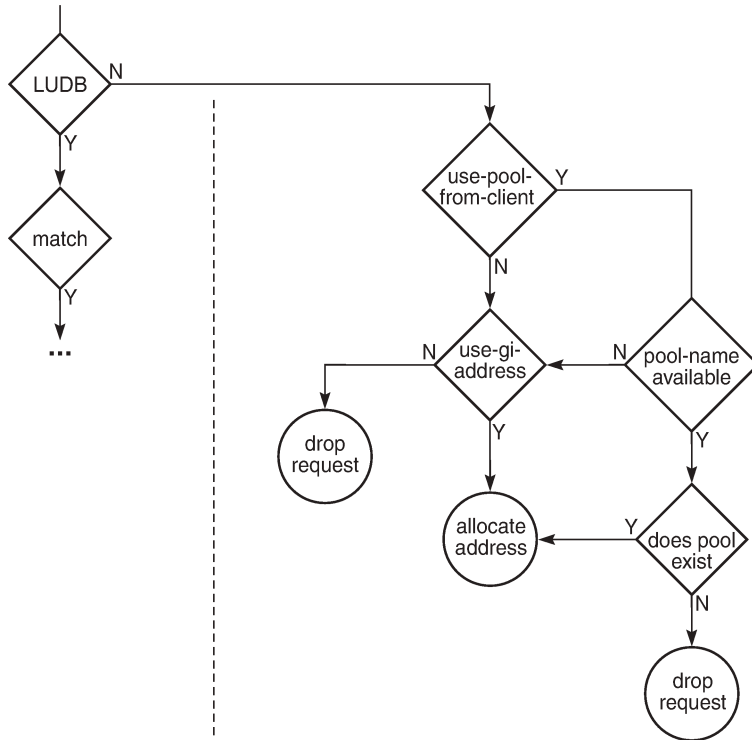
In terms of address assignment, an LUDB attached to a DHCP server can return:

- 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 109: General Address Allocation for DHCP](#). The [Local User Database for DHCPv4 Server](#) chapter applies when an LUDB is attached to the DHCP server.

Figure 109: 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.

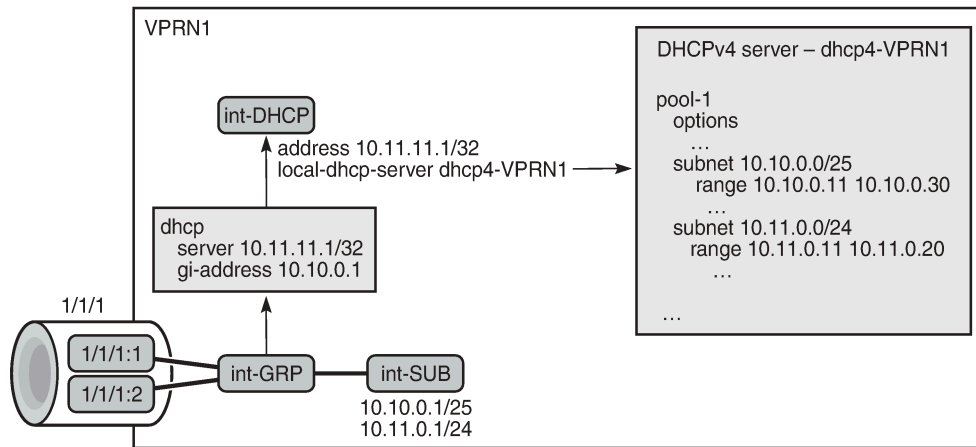
Configuration

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 110: Baseline Service Configuration](#).

Figure 110: 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 110: 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-VPRN1 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 109: 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-VRN1"
            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

```

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-VPRN1router 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     : cf1:\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-VPRN1
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-VPRN1" 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-VPRN1" 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-VPRN1          inService
=====
*A:PE1#
```

The following command shows all leases currently allocated by DHCP server dhcp4-VPRN1 in VPRN-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-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
  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
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-VPRN1" leases 10.11.0.0/24

=====
```

```

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
-----
10.11.0.14      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-VPRN1" leases 10.11.0.18/32 detail
=====
Lease for DHCP server dhcp4-VPRN1 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-VPRN1-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-VPRN1" server-stats
=====
Statistics for DHCP Server dhcp4-VPRN1 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-VPRN1" pool-ext-stats

=====
Extended pool statistics for server "dhcp4-VPRN1"
=====

```

| | 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 IpoE 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 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 0a
---snip---
31 2f 31 3a 31 ff 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-VPRN1
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-VPRN1
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
"

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
```

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      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 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-VPRN1" 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
    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

---snip---

```

```
=====
*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-VPRN1
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-VPRN1
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-VPRN1
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-VPRN1
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-VPRN1
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-VPRN1
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-VPRN1
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-VPRN1
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-VPRN1
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-VPRN1
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-VPRN1" summary
=====
DHCP server dhcp4-VPRN1  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---

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

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:

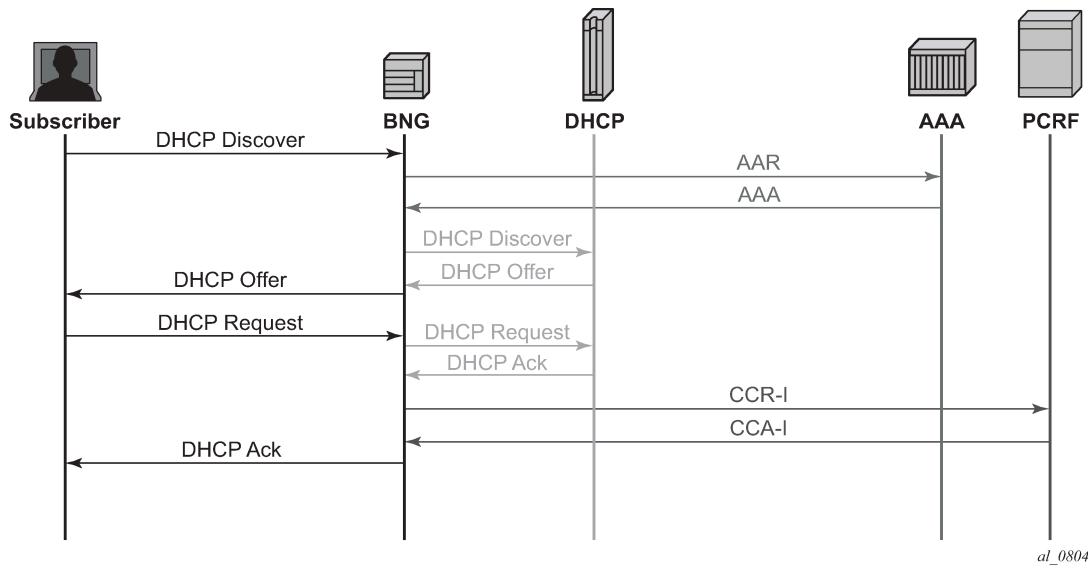
- Gx provides advanced authorization capabilities for subscribers and usage monitoring, and interfaces with a Policy and Charging Rules Function (PCRF).
- Gy or Diameter Credit Control Application (DCCA) provides on-line charging functionality, and interfaces with an On-line Charging Server (OCS).
- 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 111: 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 111: NASREQ trigger



As shown in [Figure 111: 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 *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 *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
  diamater
    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 failur* 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\\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:

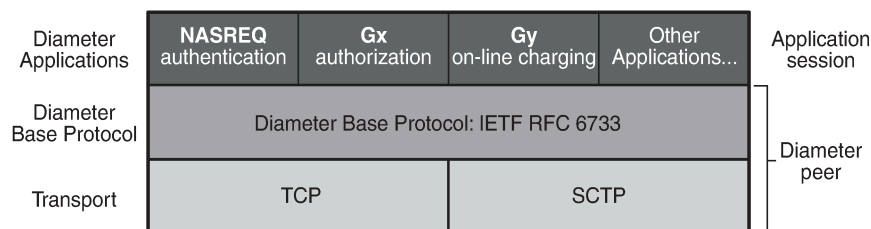
This chapter covers the Diameter base protocol implementation that is available from SR OS Release 16.0.R4 onward (configured in the **aaa** CLI context as **diameter node**). The legacy Diameter base implementation (configured in the **aaa** CLI context as **diameter-peer-policy**) is supported in maintenance 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 112: Diameter protocol stack](#).

Figure 112: Diameter protocol stack



35618


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 113: 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 113: Diameter network topology



35619

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 113: 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-gy.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-0S
  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
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](#)
- [Summary](#)
- [Overview](#)
- [Configuration](#)
- [Conclusion](#)

Applicability

The information and configuration in this chapter are based on SR OS Release 11.0.R4.

Summary

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.

Overview

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, smart-phones, 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, etc.

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 114: 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 114: Bridged RGW Scenario

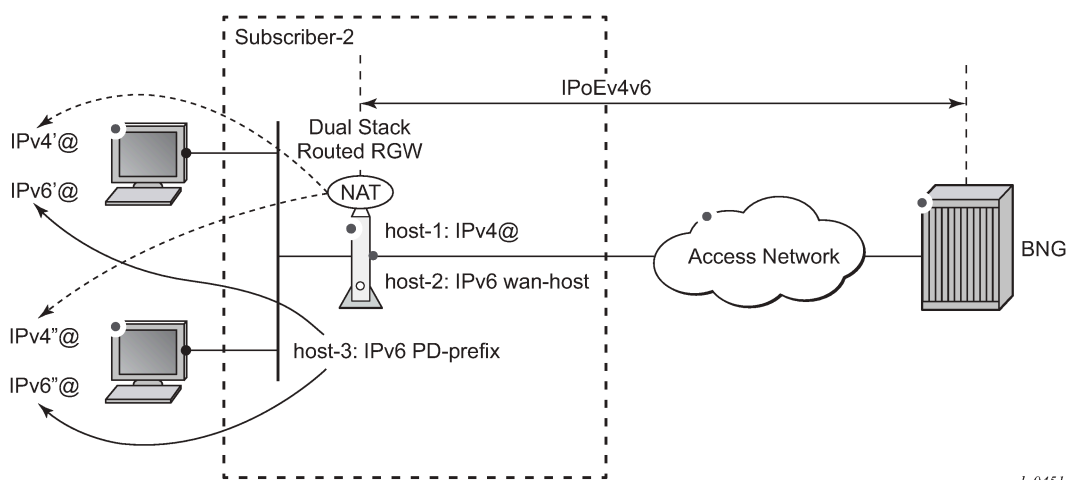


al_0450

Figure 115: 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 115: Routed RGW Scenario



al_0451

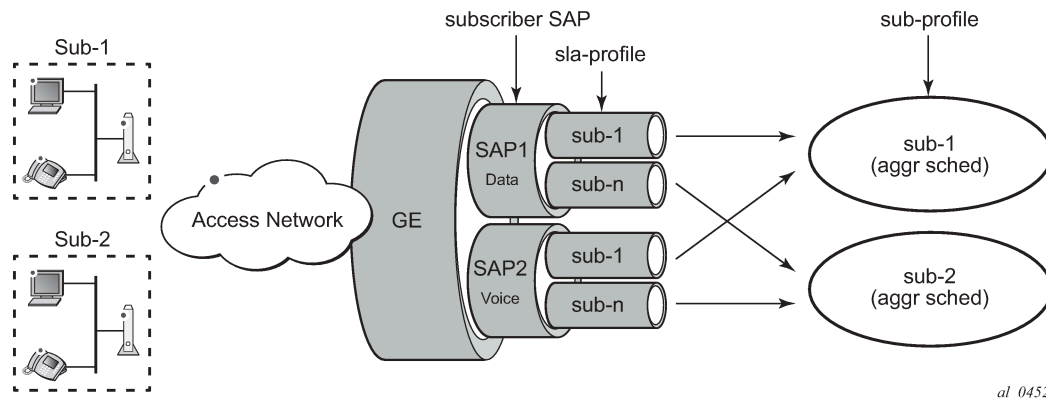
Detailed information on both scenarios can be found in [ESMv6: IPoE Dual Stack Hosts](#) and the [ESMv6: PPPoE Dual Stack Hosts](#) respectively.

Configuration

Figure 116: SLA-Profile and Sub-Profile shows two objects closely related to subscribers and subscriber hosts:

- SLA profile
- Subscriber profile

Figure 116: 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
      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 117: Subscriber Host Identification and Instantiation Process](#) and encompasses fetching the ESM data for a connecting device.

Figure 117: Subscriber Host Identification and Instantiation Process



al_0453

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 117: 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 117: 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
  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
  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 117: 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 118: 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 118: 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.

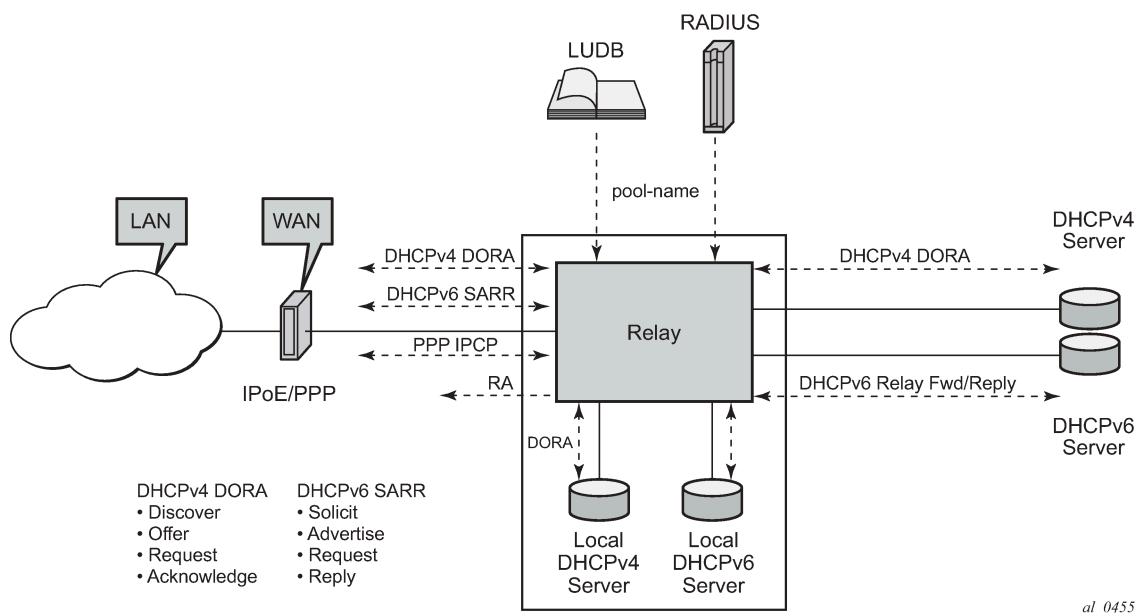
Indirect Address Assignment using a DHCP server

In the example shown in [Figure 119: 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 119: 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.

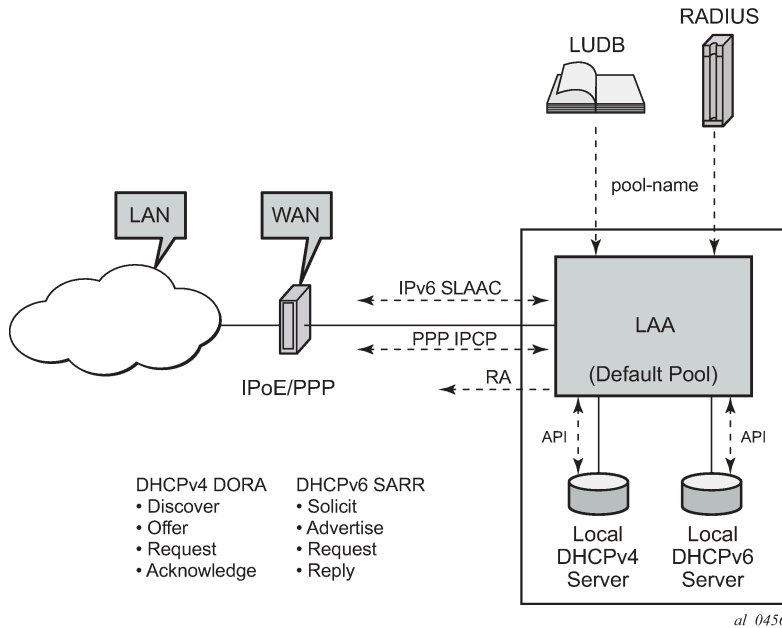
```
configure
  service
    ies 1
      subscriber-interface "sub-int-1"
        ---snip---
        group-interface "grp-int-1-1"
          ipv6
            ---snip---
            dhcp6
              relay
                server 2001:DB8::11
                client-applications dhcp ppp
                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 120: 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 120: 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
    
```

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:BN# #
```

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:BN# 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:BN# #
```

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
    
```


- 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
      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
    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
```

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
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
      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
    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
```

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 17: 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 addresses to DHCP IA_NA requests. | Automatically assigns incremental /128 addresses to DHCP IA_NA requests. | Requires no configuration. |

| | 64-bit WAN mode | 128-bit WAN mode | Comments |
|--|--|---|---|
| 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 mode128 create
```


- 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 mode128 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
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
    |
```

```

+-- IPOE-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
    |
    +-- IPOE-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 121: 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 122: 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 121: Wholesale/Retail Model 1

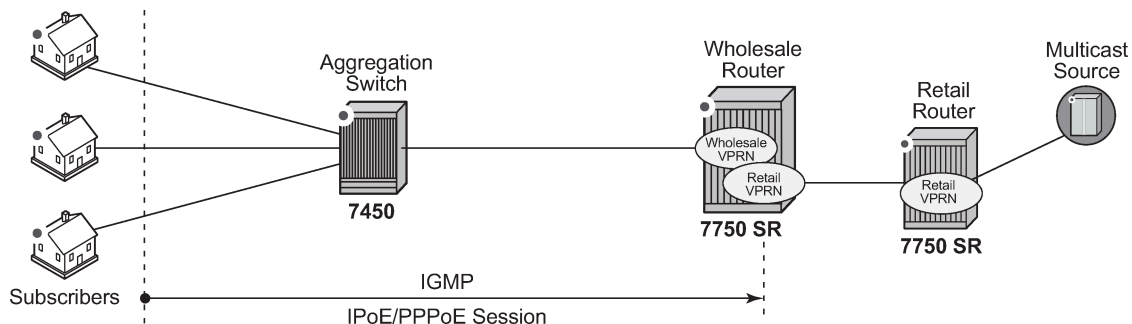



Figure 122: Wholesale/Retail Model 2



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

Figure 29: Local URL-list filtering setup 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 123: 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
    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

```

```
        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                               IGMP-Policy
HostAddr      GrpAddr      SrcAddr      GrpItf      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      0d 00:02:16      1
                                           Dynamic      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 124: L2TP Wholesale-Retail Multicast](#). LAC-1 contains the wholesale configuration while LNS-1 contains the retail configuration.

Figure 124: 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
  ---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
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
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-VPRN-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 | 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

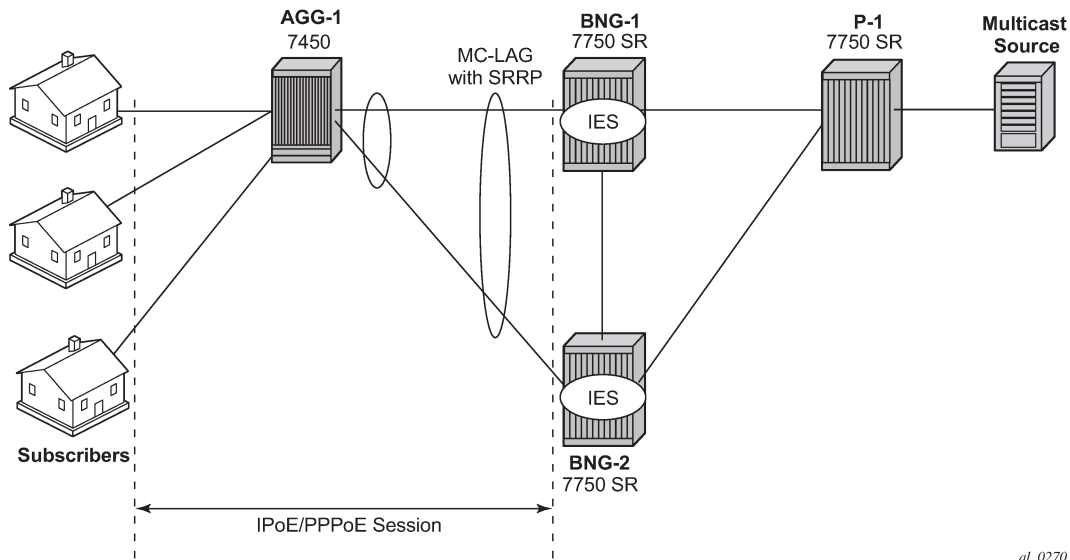
Triple Play Service Delivery Architecture (TPSDA) allows operators to integrate High Speed Internet (HSI), voice, and video services within a single network. The goal of this chapter is to walk through a TPSDA multicast architecture with redirection. The topics are divided into the following sections:

- 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 125: 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 125: 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 126: 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 126: 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
        exit
        no shutdown
    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

```

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
                   Up   Up    192.168.10.1    3/3     1        none
-----
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 127: 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 127: 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
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

```

```

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
no shutdown
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

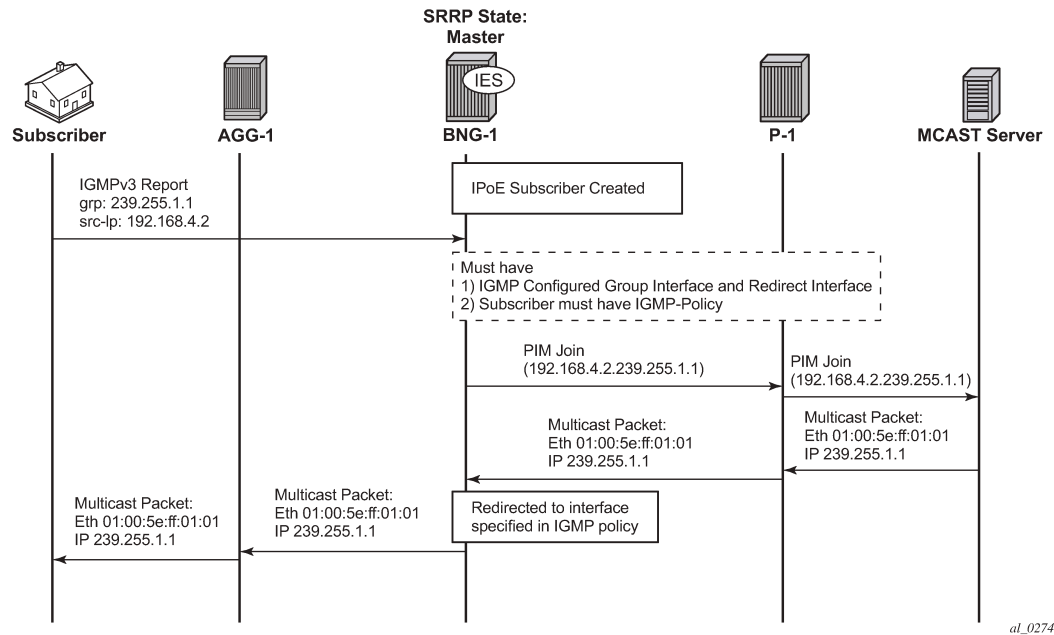
```

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 128: 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
                   Up   Up    192.168.10.253   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
-----
HSMDA-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                               IGMP-Policy
HostAddr  GrpAddr  SrcAddr  GrpItf  Type  Up-Time  NumGroups
Mode
Blk/Fwd
-----
ipoe-sub-1                               igmp-policy-1
10.0.0.11  grp-int-1  Dynamic  0d 00:03:30  1
239.255.1.1  Dynamic  Include
192.168.4.2  Dynamic  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:BNB-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.

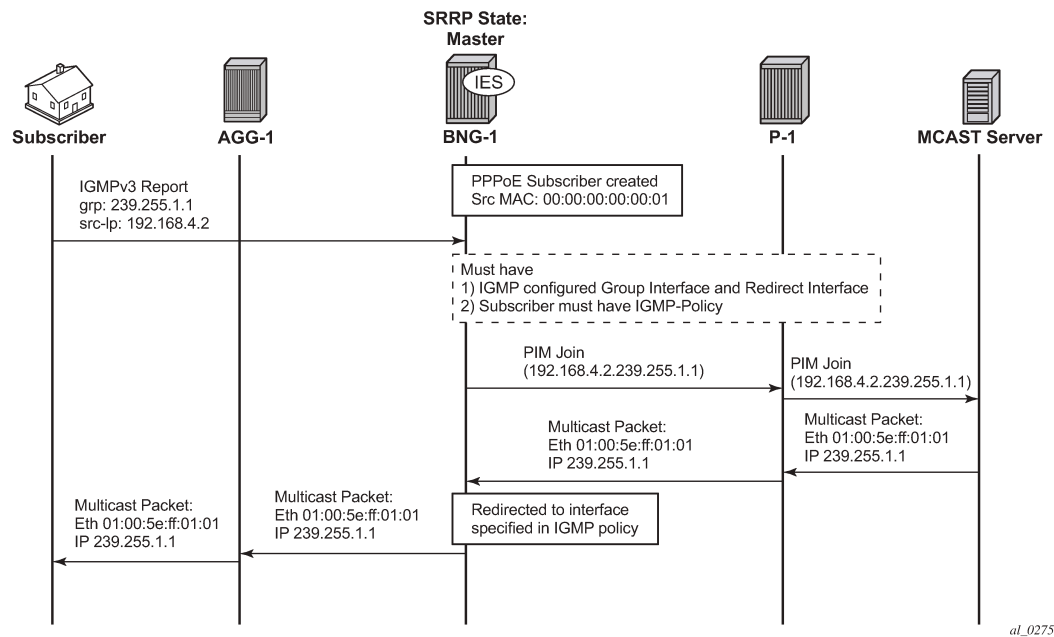
```

*A:BNB-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:BNB-1#
    
```

ESM IGMP PPPoE Walkthrough

The same baseline configuration is used for PPPoE subscriber. [Figure 129: PPPoE Multicast Flow](#) shows the message flow for delivery of multicast streams to PPPoE subscribers.

Figure 129: 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 IPEoE 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
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 Avail Mand BW : unlimited
MCAC In use Mand BW: 0
MCAC Avail Opnl 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               : lag-1:4
Admin/Oper version: 3/3
Num Groups        : 0
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
-----
```

```
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                               IGMP-Policy
HostAddr      GrpItf                        NumGroups
  GrpAddr      Type              Up-Time      Mode
  SrcAddr      Type                               Blk/Fwd
-----
pppoe-sub-1      igmp-policy-1
  10.0.0.12      grp-int-1                1
  239.255.1.1    Dynamic                  0d 00:04:42  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 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
Omcrc Alarmed Entries: 0
Omcrc Standby Entries: 0
Associated Shared Records (ALL): 0
Associated Shared Records (LD): 0
*A:BNB-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
```

```
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 (GlblDel)
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 130: 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 130: Network Topology Overview



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 131: 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 131: 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
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
    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
  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
  no shutdown
  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
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

```

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
mcaac               + 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
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.

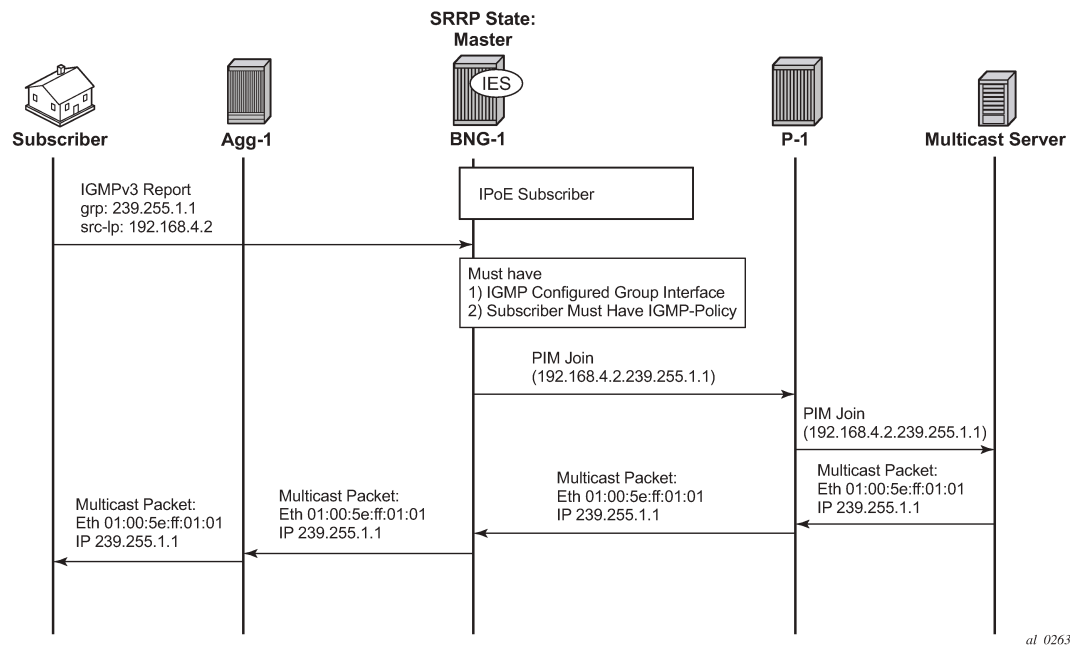
ESM IGMP IPoE Walkthrough

With the baseline configuration applied, the BNG is ready to process IGMP messages and deliver multicast. [Figure 132: 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 132: 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. Policer Ctrl. : N/A
I. vport-hashing : Disabled
E. Agg Rate Limit: Max
```

```

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
-----
HSMDA-2
-----
I. Qos Policy    : 1                    E. Qos Policy    : 1
                                           E. Agg Rate Limit: Max
                                           Pkt Byte Offset  : add 0*
E. WRR Policy    : N/A
-----
Last Mgmt Change : 06/21/2017 11:58:50
=====
* 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.

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: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  : 10.0.0.11                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:BNB-1#

```

For more IGMP details on the group interface, including maximum multicast groups and bandwidth management, use the following command:

```

*A:BNB-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:BNB-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        NumGroups
SrcAddr                   Type           Mode           Blk/Fwd
-----
subscr-1                  igmp-policy-1
10.0.0.11                 grp-int-1      0d 00:08:50    1
239.255.1.1               Dynamic
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:BNB-1#

To check for an individual subscriber and its requested (S,G)s, the following command can be used.

```
*A:BNB-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:BNB-1#

ESM IGMP PPPoE Walkthrough

IGMP message processing and delivery of multicast streams to PPPoE subscribers is considered next. [Figure 133: PPPoE Multicast Flow](#) shows the message flow for multicast stream delivery to PPPoE subscribers.

Figure 133: 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:BNB-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:BNB-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:BNB-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:BNB-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:BNB-1#

```

To view each individual subscriber and their respective (S,G)s, use the following command.

```

*A:BNB-1# show service active-subscribers igmp subscriber "ppoe-sub-1" detail
=====
Active Subscribers Detail
=====
Subscriber          IGMP-Policy
HostAddr            GrpItf
GrpAddr             Type          Up-Time      NumGroups
SrcAddr             Type          Mode
Blk/Fwd
-----
ppoe-sub-1          igmp-policy-1
10.0.0.15           grp-int-1
239.255.1.1         Dynamic       0d 00:07:31 1
192.168.4.2         Dynamic       Include
Fwd
-----
Number of Subscribers : 1
=====
*A:BNB-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
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:BNB-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:BNB-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
-----
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
Omcrc Alarmed Entries: 0
Omcrc 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, G1b"
```

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
  commit
exit
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
  default-action accept
  exit
  exit
  commit
exit
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
```

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
```

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

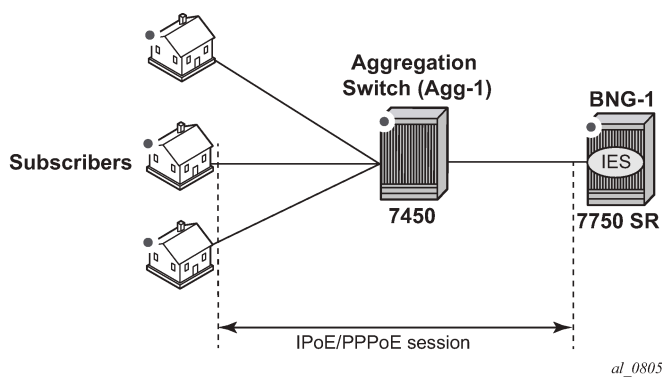
The information and configuration in this chapter was based on SR OS Release 13.0.R1. The CLI in the current edition is SR OS Release 16.0.R7 based. Both Internet Protocol over Ethernet (IPoE) and Point-to-Point Protocol over Ethernet (PPPoE) are supported.

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 134: 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 134: 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
      group-interface "group-int-1" create
        ipv6
          router-advertisements
            prefix-options
              autonomous
            exit
          no shutdown
        exit
      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
-----
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 : "0217FF00000000C5CB82FA7"
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
  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

```

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

```

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
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 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
exit

```

For operators that do not enable IPoE 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
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:BNB-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:BNB-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](#)
- [Summary](#)
- [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.

Summary

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.

Overview

PPPoE Hosts in a Routed CO Environment

The network topology for a Routed CO environment is displayed in [Figure 135: Routed CO Network Topology](#).

Figure 135: 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 18: Reserved PPPoE Tags](#). All other tag values between 0 and 65535 are to be assigned by IANA.

Table 18: Reserved PPPoE Tags

| Tag Value | Tag Name |
|-----------|-------------|
| 0 0x0000 | End-Of-List |

| Tag Value | Tag Name |
|------------|--------------------|
| 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 136: 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 to 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:

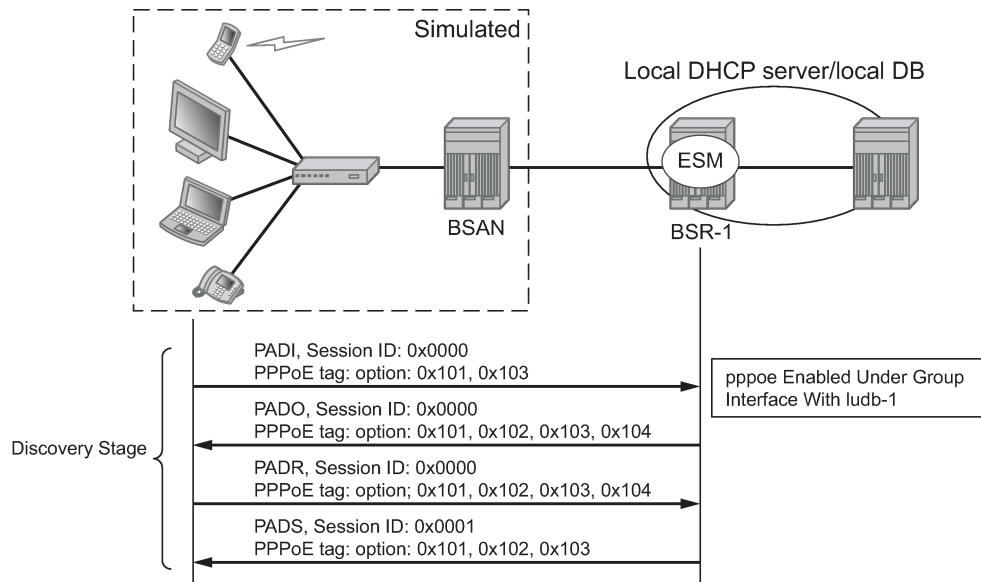
- Peer supports the option and its content.
 - Peer will agree and send config-ack.
- Peer does not support an option
 - Peer will send configure-reject with the option that is not supported.
 - Resend of configure-request without that option.
 - Peer agrees and sends config-ack.
- Peer does support 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 19: 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 |
| 13 | Time-Remaining |
| 14 | Reset-request CCP |

| Code | Packet Type |
|------|---------------|
| 15 | Reset-Ack CCP |

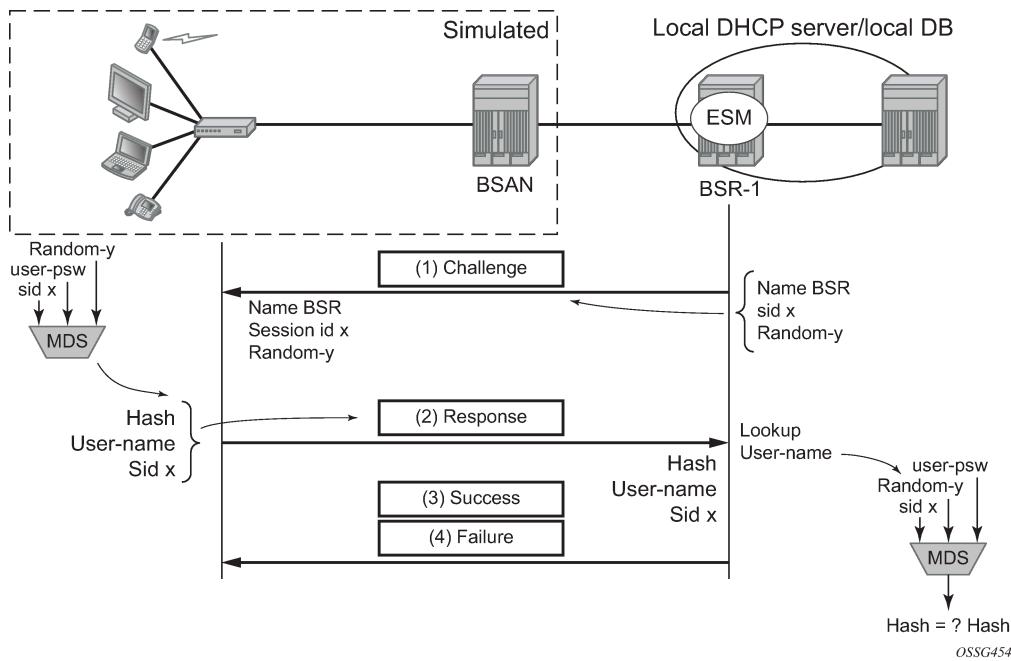
Figure 137: LCP Phase Messages



Authentication Phase

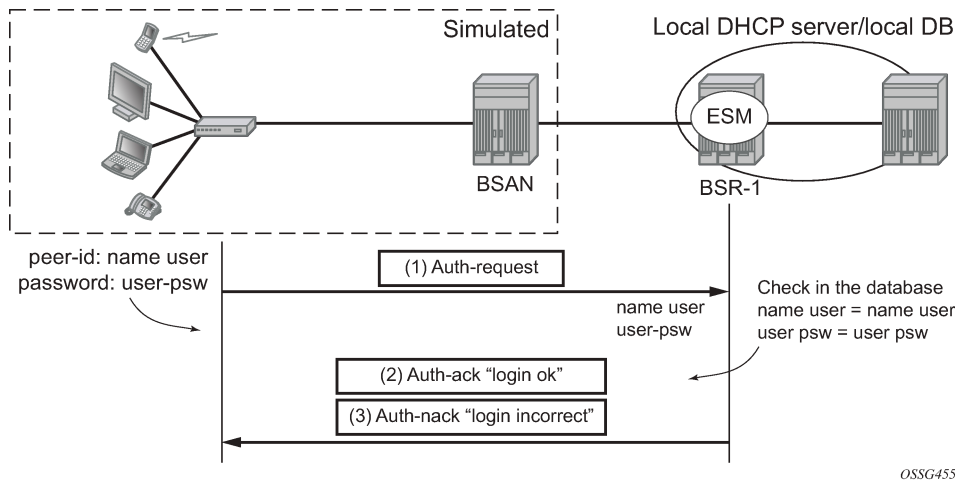
The client authenticates itself through PAP (PPP Password Authentication Protocol) or CHAP (Challenge Handshake Authentication Protocol) to check for access permission. For the CHAP authentication, the BSR initiates the authentication as shown in [Figure 138: CHAP Handshaking Overview Process](#). The password is hashed on the link and plain text in the RADIUS Access-Request message.

Figure 138: CHAP Handshaking Overview Process



For PAP, the client initiates the authentication as shown in [Figure 139: PAP Overview Process](#). The password is sent as plain text on the link and hashed in a RADIUS Access-Request message.


Figure 139: PAP Overview Process



Network-Layer Protocol Phase (PPP IPCP Opening Phase)

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 140: 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 141: Keepalive Messages

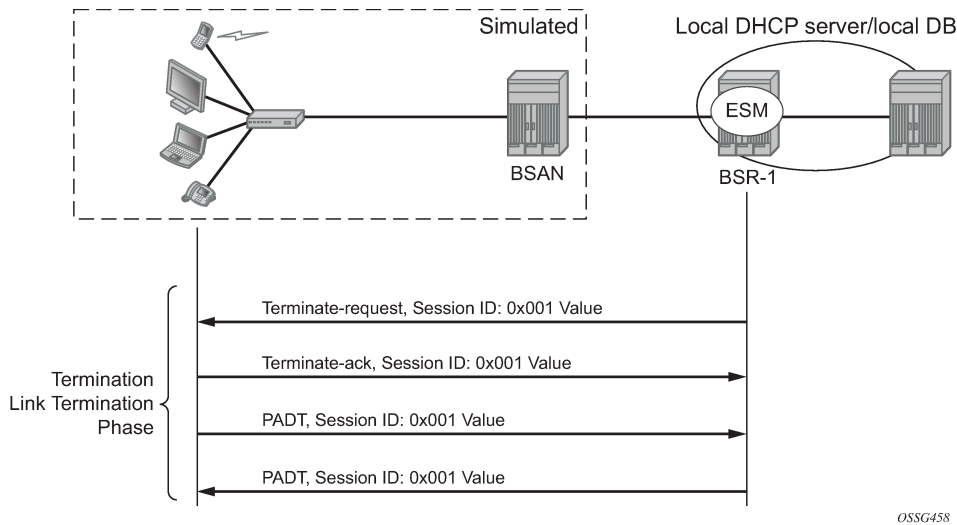


Termination

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 142: Link Termination Phase



Configuration

Session Set-Up, 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
      exit
    pppoe
      policy "ppp-policy-1"
      user-db "ludb-1"
      no shutdown
```

```

        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
```


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 <deci-seconds>

<deci-seconds>      : [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..0xFFFFFFFF...(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 143: Authentication Flow Chart](#) and starts with the request to the middle left.

Figure 143: 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
  ---snip---
  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
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
  ---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
PADO 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 LUBD, 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         00:00:67:14:01:02 1    0d 00:14:05     local
  10.2.0.4

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: SVCMgr #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: "cicuit11"
        [0x02] Agent-Remote-Id: "remotel1"
"
```

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/VPRN 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

```

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: "cicuit11"
          [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---
```


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
  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 144: Pado-Delay Scenario



OSSG459

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

Applicability

This chapter describes ESMv6: IPoE dual stack hosts and is based on SR OS 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/VP RN service) with Enhanced Subscriber Management (ESM)
- Routed Gateway (RG) in the home

Summary

In this chapter, the configuration, operation, and troubleshooting of IPoE dual stack hosts in a routed home gateway environment is described. 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.

Overview


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 145: Stateful IPoE Dual Stack Subscriber Hosts

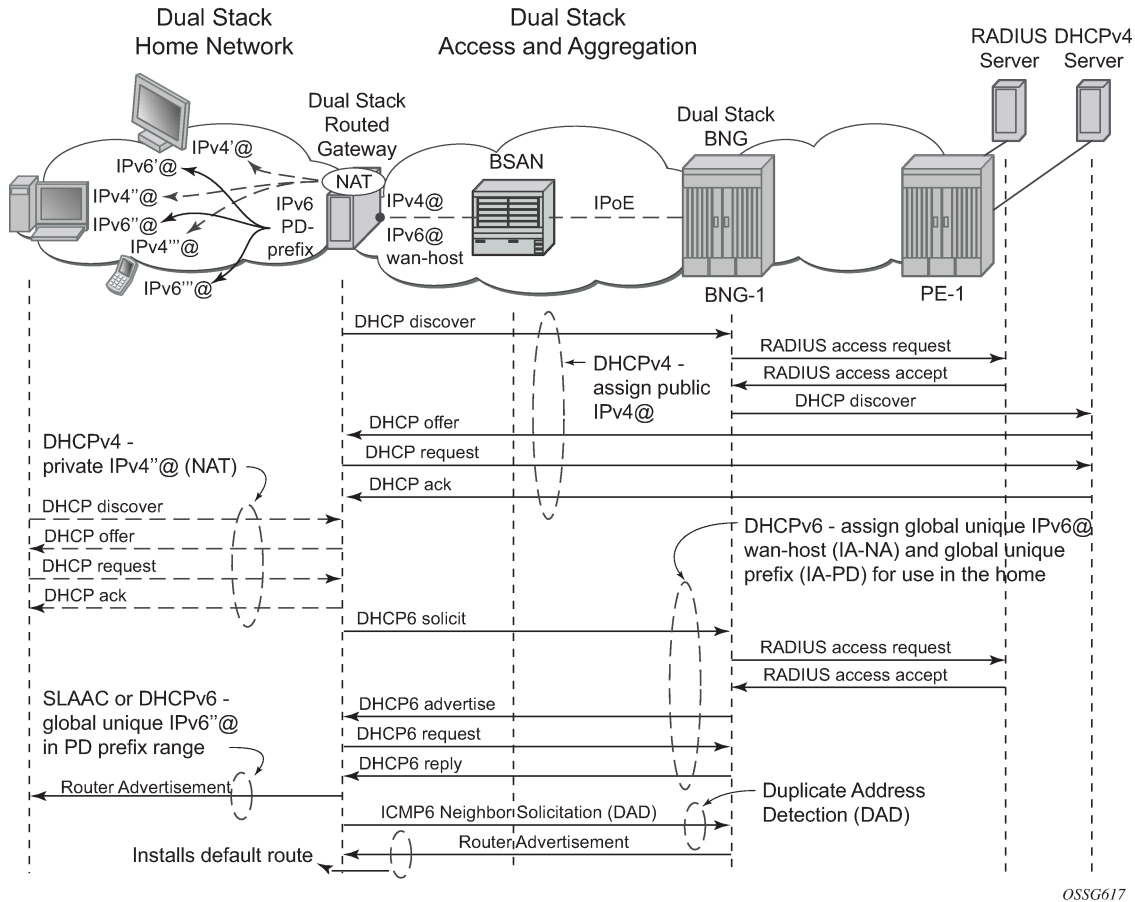


Dual Stack IPoE Routed Gateway

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 146: Dual Stack IPoE Routed Gateway Service



OSSG617

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
```

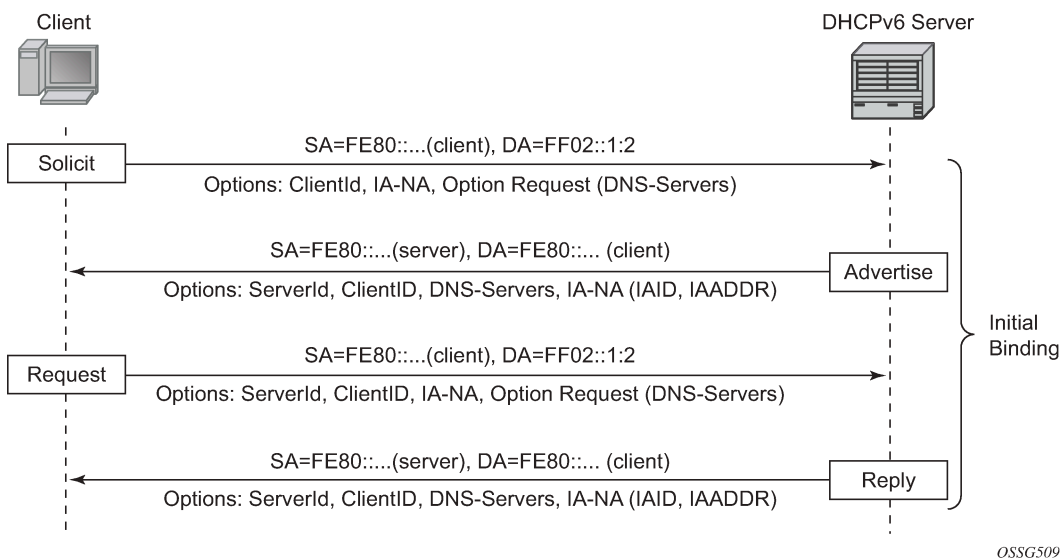
- Identity Association for Prefix Delegation (IA-PD), defined in RFC 3633. Used for prefix delegation assignment (for an explanation on prefix delegation, see [Prefix Delegation](#))

```

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 147: DHCPv6 Lease Process \(Part A\)](#) and [Figure 148: DHCPv6 Lease Process \(Part B\)](#).

Figure 147: 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 148: 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 149: 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 149: 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 149: 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
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
        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
```

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:ddd:1::1 2001:db8:ddd:2::1
      exit
      no shutdown
    exit
    ---snip---
  exit
  no shutdown
exit
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
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 20: Applicable Subscriber-Prefix Parameters](#) provides an overview of the subscriber-prefix parameters that apply:

Table 20: 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

```

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 150: IPv6 Address/Prefix Timers



When violating the following rule, the default timers will be used:

Table 21: 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 |
| 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
```

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
          exit
        exit
      exit
    exit
  exit
exit
```

```

reachable-time 0
retransmit-time 0
router-lifetime 4500
exit
exit
exit
exit
exit
exit
exit
exit

```

Table 22: Router Advertisements Parameters

| Parameter | Description (RFC 4861, Neighbor Discovery for IP version 6 (IPv6)) | 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) |
| 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) |

| Parameter | Description (RFC 4861, <i>Neighbor Discovery for IP version 6 (IPv6)</i>) | Value Range (default) |
|---|--|-----------------------|
| 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
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 23: RADIUS AVPs](#).

Table 23: 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:ddd:1::1 2001:db8:ddd:2::1
            exit
            no shutdown
          exit
        exit
      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
                    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 24: 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 25: 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. |
| 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. |

| Parameter | Comment |
|-----------------------|--|
| 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. |
| 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. |

| Parameter | Comment |
|------------------|---|
| 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
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:BNB# 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::1/128          Remote Sub Mgmt 01h00m03s 0
      [grp-int-1]                    0
2001:db8:b001:103::3/128          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:BN#

```

RADIUS statistics can be shown as follow:

```

*A:BN# 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:BN#

```

```

*A:BN# 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 Alc-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
                   submgt
                   14.0 (current) - cf1:\submgmt.012
                   13.0           - cf1:\submgmt.011
                   12.0           - cf1:\submgmt.010
                   11.0           - cf1:\submgmt.009
                   10.0           - cf1:\submgmt.008
                   9.0            - cf1:\submgmt.007
                   8.0            - cf1:\submgmt.006
                   7.0            - cf1:\submgmt.005
                   6.0            - cf1:\submgmt.004
                   5.0            - cf1:\submgmt.003
                   4.0            - cf1:\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 151: 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

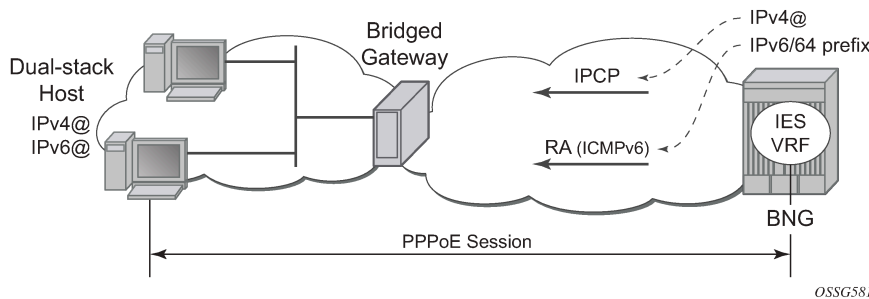
In the dual stack PPPoE host service, the PPPoE session is initiated directly from a dual stack device in the home network. PPPoE is used to carry IPv6 and (optionally) IPv4 traffic from the device to the broadband remote access server (BRAS), also called broadband network gateway (BNG).

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 152: 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 153: 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
```

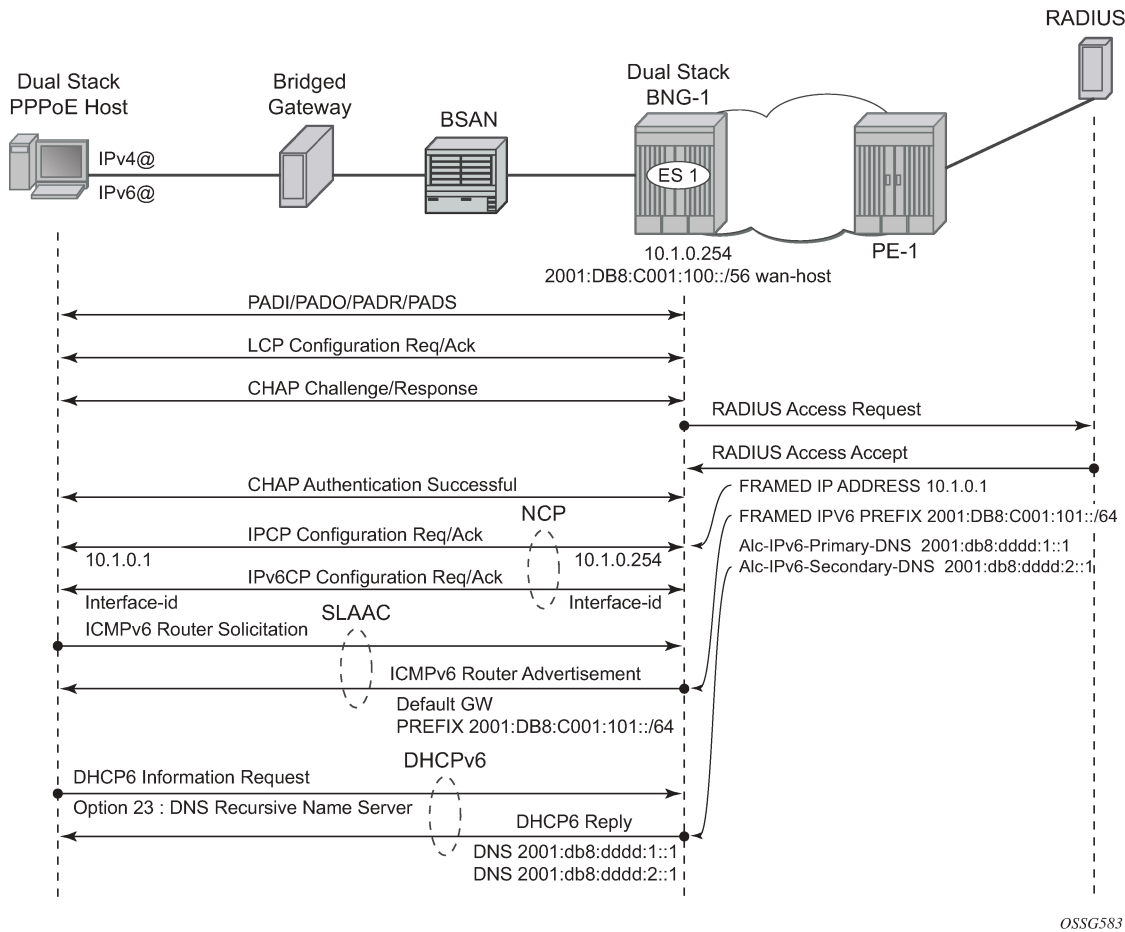
```
configure
  subscriber-mgmt
    sla-profile "sla-profile-1" create
  exit
  sub-profile "sub-profile-1" create
  exit
  sub-ident-policy "sub-ident-1" create
    sub-profile-map
      use-direct-map-as-default
  exit
  sla-profile-map
    use-direct-map-as-default
  exit
  exit
exit
exit
```

Service

Dual Stack PPPoE for Bridged Gateway

[Figure 154: 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 154: 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
  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 Ipv6 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 Ipv6 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 26: Subscriber Prefix Parameters](#) and [Table 27: 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 26: 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 |

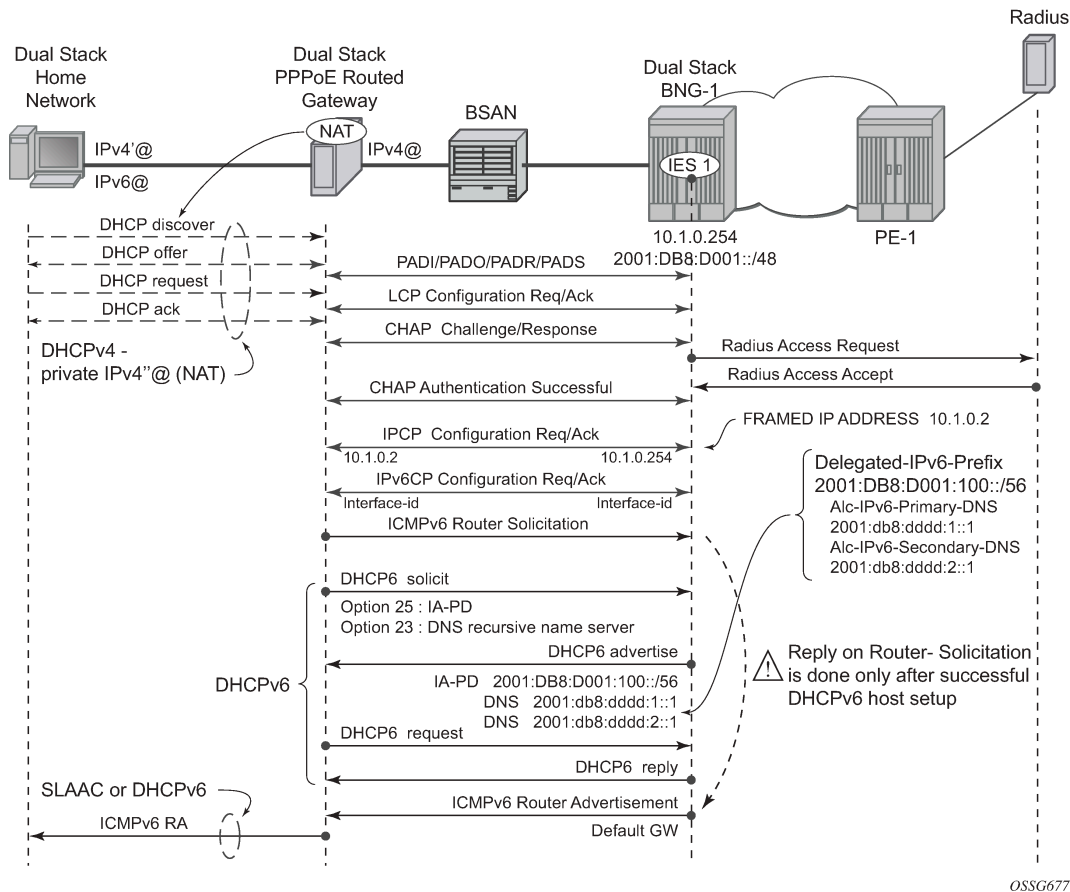
Table 27: Subscriber Prefix Subnetting for SLAAC

| 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 155: 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 155: Dual Stack PPPoE for Routed Gateway



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
  no shutdown
  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 28: Subscriber-Prefix Parameters](#) and [Table 29: 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 28: 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 29: 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::/48 | 2001:db8:d002:100::/56 |
| 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 30: RADIUS AVPs](#).

Table 30: 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
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
```

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
```

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 31: SLAAC-Related Parameters](#) .

Table 31: 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 Ipv6 are listed and explained in [ESMv6: Ipv6 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 Ipv6 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

```

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 156: DHCPv6 Renewals](#).

Figure 156: 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-Stdby
-----
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: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: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:BN#
```

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 : ORO (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:BN# 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:BN#

```

To clear the statistics use following command:

```
*A:BN# 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:BN# 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:BN#

```

```

*A:BN# 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:BN#

```

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: BNG#

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 : ORO (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:ddd:1::1
    Server : 2001:db8:ddd: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_SRVR (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_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
"

```

Use the following command to display the DHCPv6 statistics.

```

*A:BN# 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:BN#
```

Use the following command to clear the DHCPv6 statistics.

```
A:BN# 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:BN# 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
    exit
  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 32: 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 32: 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 157: Topology](#).

Figure 157: 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 IPoE 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
```

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
        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
    
```

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
                subnet 10.10.0.0/24 create
                    options
                        subnet-mask 255.255.255.0
                    exit
            exit
        exit
    
```

```

        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
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

```

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
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:BNB-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:BNB-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:BNB-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:BNB-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 7x50 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
          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
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
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
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
-----

```

| 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 LUBD.

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
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
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
```



```

-----
(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:BNB-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:BNB-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 LUBB + RADIUS Authentication

The IP address is assigned via LUBB.

- RADIUS provides sub/sla-profile strings and a framed IPv4 route.
- LUBB 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 LUBB.

```
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
exit
```

The LUBB 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 LUBB 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-ppoe-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

```

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 7x50 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
          dhcp
            proxy-server
              emulated-server 10.12.0.1
              no shutdown
            exit
          user-db "ludb-4"
        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
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          Subscriber
IP Address
MAC Address  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 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
    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
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 158: GTP Access to the BNG](#).

Figure 158: 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 159: 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 159: EPS Bearer Across the Different Interfaces



27983

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 160: GTP-C and GTP-U Encapsulation](#). UDP ports 2123 and 2152 are used for GTP-C and GTP-U, respectively.

Figure 160: GTP-C and GTP-U Encapsulation

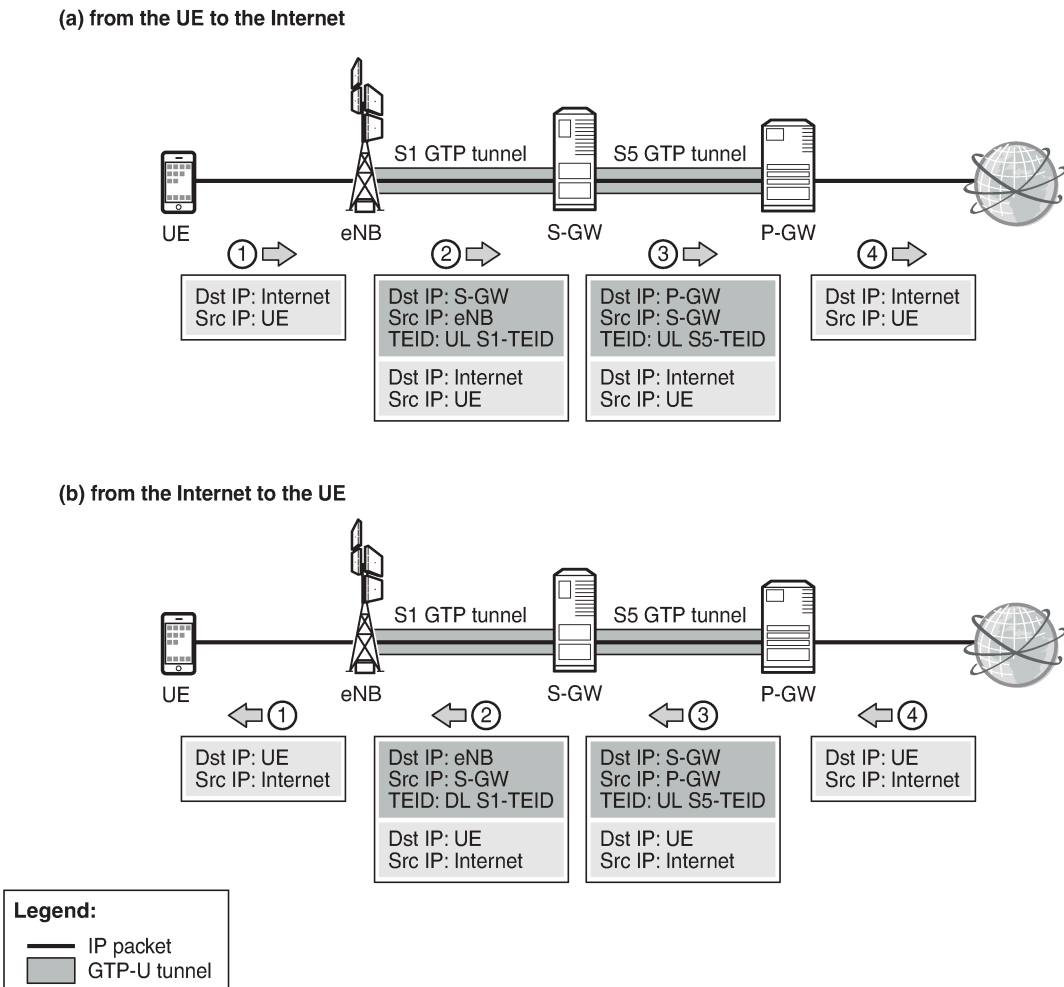


27984

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 161: 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 161: 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.

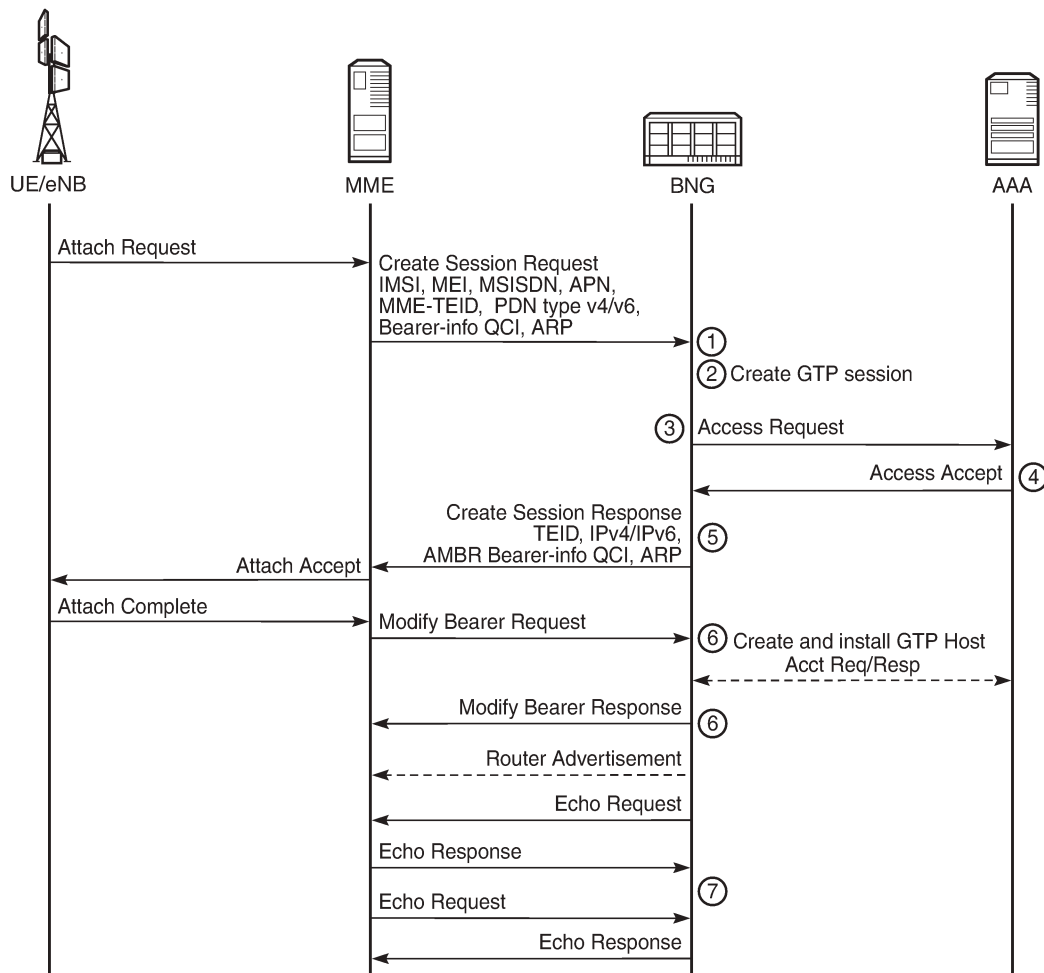
ESMoGTP requires PXC FPE in order to support L3 re-routing in the access network and in-line encapsulation/decapsulation of the GTP protocol.

GTP-C Control - GTP Host Creation and Host Deletion

Figure 162: 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 162: GTP-C Control - GTP Host Creation




27986

Figure 163: 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 163: 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 163: GTP-C Control - GTP Host Deletion




27987

Configuration

Figure 164: GTP Access Topology shows the topology used for demonstrating GTP Access.

Figure 164: GTP Access Topology



27988

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 7450 ESS, 7750 SR, and 7950 XRS Advanced Configuration Guide - Part I, the Port Cross-Connect (PXC) chapter for more information. The PXC port 1/2/4 additionally requires subscriber management extensions to be enabled in the forwarding plane (FPE), see 7450 ESS, 7750 SR, and 7950 XRS Advanced Configuration Guide - Part II, the VXLAN Forwarding Path Extension chapter, 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   Admin   Oper    Port    Description
Id    State   State   Id
-----
1     Up      Up      1/2/4   (Not Specified)
-----
No. of PXC: 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   : N/A
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
  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
```

Two services are used to provide GTP access to the BNG, as shown in [Figure 165: 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 165: GTP Tunnel and Subscriber Termination Configuration Logic



GTP Tunnel Terminating VPRN

In [Figure 165: 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
  exit
  static-route-entry 192.168.100.0/24
  next-hop 192.168.0.2
  no shutdown
  exit
  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

```

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

```

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 165: 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 "001020000000111"
        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 164: 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
  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.  Ing.  Egr.  Egr.  Adm  Opr
                   QoS       Fltr  QoS   Fltr
-----
[pxc-1.b:1.8]         11         1    none  1     none  Up   Up
-----
Number of SAPs : 1

```

```
=====
* A: BNG-1#
```

When the UE with IMSI 00102000000111 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: BNG-1# show subscriber-mgmt gtp s11 session imsi 00102000000111
```

```
=====
GTP S11 sessions
=====
```

```
IMSI                : 00102000000111
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: BNG-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: BNG-1# show service active-subscribers hierarchy
```

```
=====
Active Subscribers Hierarchy
=====
```

```
-- 00102000000111 (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] 192.168.50.12    001020000000111
00:03:ff:e0:01:20  N/A      GTP      Fwding
[pxc-1.b:1.8] 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]
Next Hop[Interface Name]
-----
2001:db8:101::/48
int-SUSBC
Type Proto Age Pref
Local Local 21h23m02s 0
0
2001:db8:101:1::/64
[int-GRP]
Type Proto Age Pref
Remote Sub Mgmt 00h29m21s 0
0
2001:db8:101:2::/64
[int-GRP]
Type Proto Age Pref
Remote Sub Mgmt 00h01m42s 0
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
exit
subscriber-mgmt
  local-user-db "ludb-GTP"
  detail all
exit
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 162: 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

"

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

"

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 vprn1 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 vprn1 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 vprn1 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=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

"

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=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=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 163: 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

"

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

"

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 **hosts\sessions** 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 166: 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 166: Test Environment Example



26876

The QoS hierarchy associated with the subscribers is shown in [Figure 167: 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.



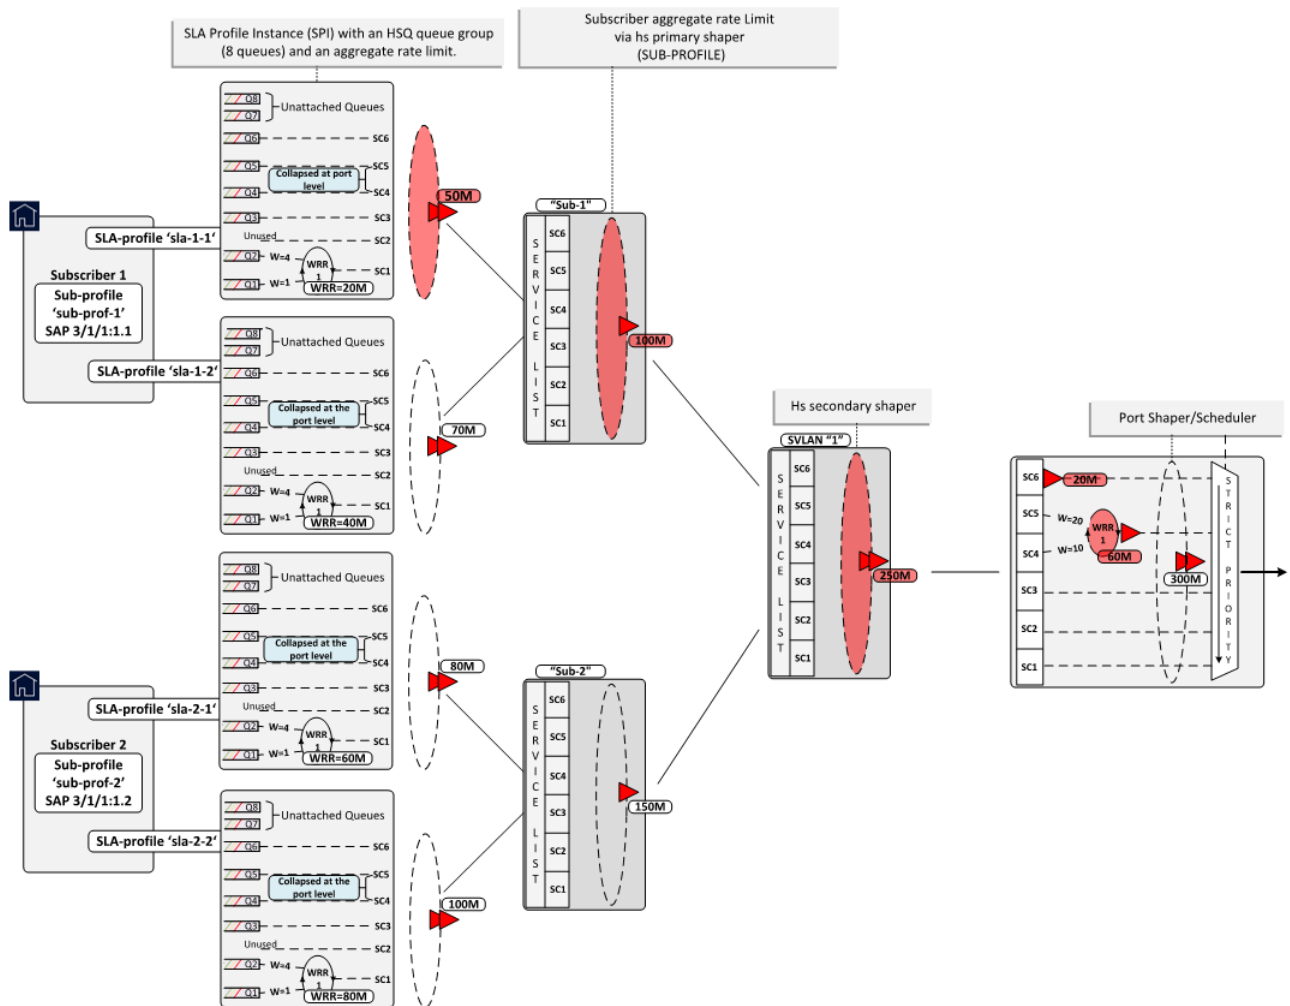
Note:

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 167: 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-agg-rate-limit 70000
  exit
exit
sla-profile "sla-2-1" create
  egress
    qos 30
  exit
  hs-agg-rate-limit 80000
  exit
exit
sla-profile "sla-2-2" create
  egress
    qos 40
  exit
  hs-agg-rate-limit 100000
  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
```

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
      encap-type qinq
      egress
        hs-scheduler-policy "hs1"
        hs-secondary-shaper "1" create
          aggregate
            rate 250000
            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 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"
Icmp 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
-----


| 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

                Packets      Octets
CPM Ingress      : 5          1077

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

-----
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 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 |
Network Egress Classification - | | 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 168: 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 168: 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 168: 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 168: 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 168: Managing Congestion on HSQ in Expanded SLA Mode

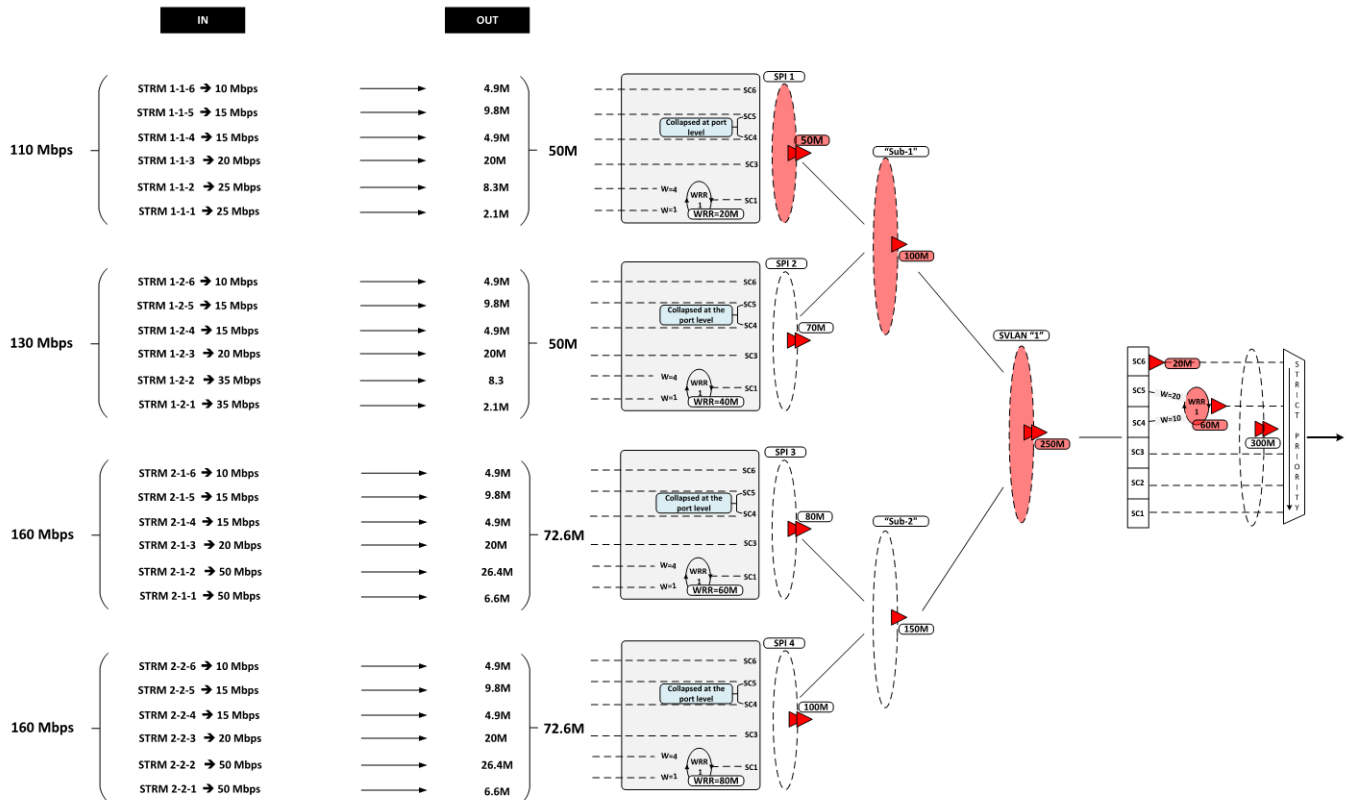


Table 33: Input/Output Rates lists the input/output rates throughout the subscriber QoS hierarchy. Red shaded table cells represent congested objects. The numbers in blue represent the L2 output rates measured on the traffic generator. The numbers in red (in parentheses) are configured aggregate rates 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 in 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 33: 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 (20) | - | |
| 1-1-4 | 15 | | 4.9 | - | | | | - | - | |

| 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 | | | | | SC 5 | | 60 | |
| 1-1-1 | 25 | | 2.1 | (4:1) | | | | | 40 | | (60) (2:1) | |
| 1-2-6 | 10 | 130 | 4.9 | - | 50 (70) | | | | | | | |
| 1-2-5 | 15 | | 9.8 | - | | | | | | | | |
| 1-2-4 | 15 | | 4.9 | - | | | | | | | SC 4 | |
| 1-2-3 | 20 | | 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) | | | | | | |
| 2-1-5 | 15 | | 9.8 | - | | | | | | | SC 3 | - |
| 2-1-4 | 15 | | 4.9 | - | | | | | | | 81 | - |
| 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) | | | | | | | |
| 2-2-5 | 15 | | 9.8 | - | | | | | | | - | |
| 2-2-4 | 15 | | 4.9 | - | | | | | | | SC 1 | - |
| 2-2-3 | 20 | | 20 | - | | | | | | | 88 | - |
| 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 33: 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 33: 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 (39.6Mb/s x 2) is below the configured limit of 100Mb/s for sub-1.
- 79.2Mb/s (39.6Mb/s x 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 (39.6Mb/s x 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 33: 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 169: 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 169: Test environment example




26891

The QoS hierarchy associated with the subscribers is shown in [Figure 170: 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 170: 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 queuing on ingress because ingress shared queuing is not supported on HSQ. With shared queuing disabled and service queuing 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 queuing.

Removing the SAP queues on ingress becomes even more relevant on HSQ, which does not support shared queuing on ingress. Non-shared queuing 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 "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 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
-----


| 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
-----
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 170: 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 |
| 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 171: 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 171: 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 171: Managing Congestion on HSQ in Single SLA Mode

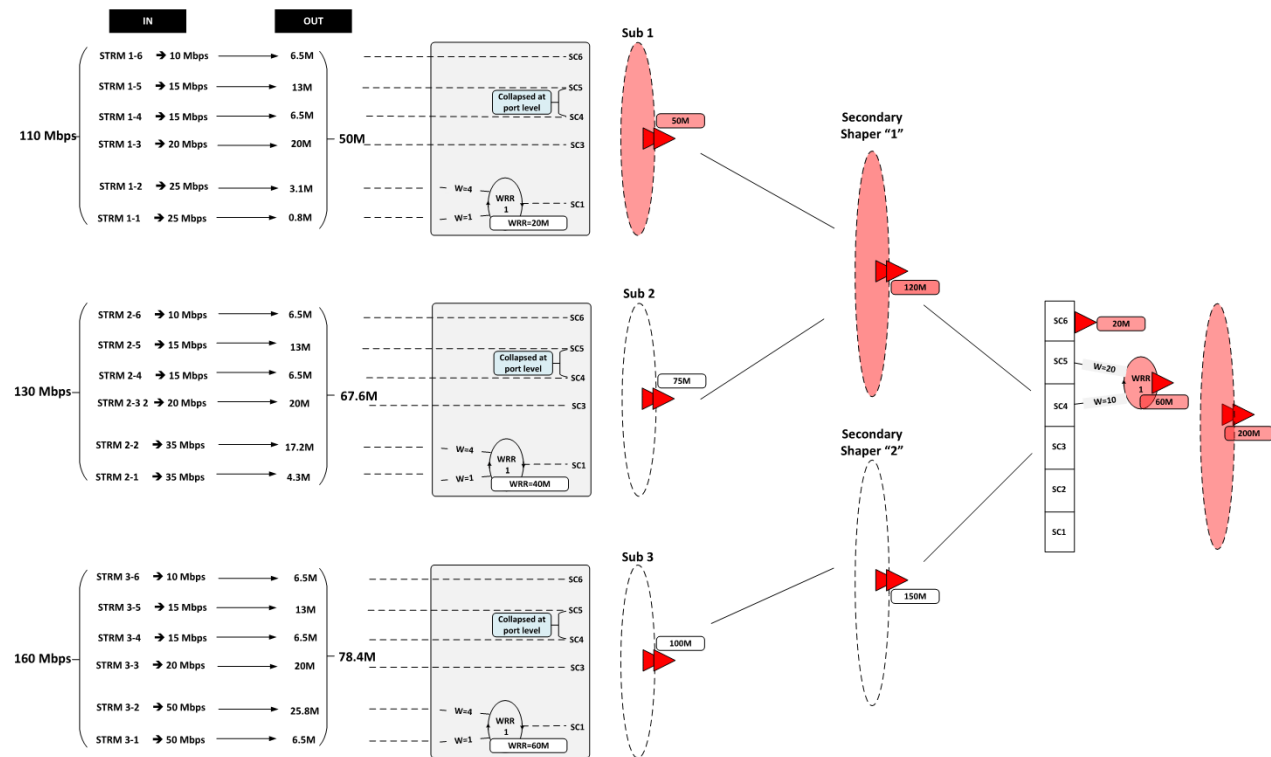


Table 34: Input and output rates throughout the subscriber QoS hierarchy lists the input/output rates throughout the subscriber QoS hierarchy. Red shaded table cells represent congested objects. The numbers on the blue background represent the output rates measured on the traffic generator. The numbers 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 in 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 34: 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 | - | | | SC 5 | 60 | |
| | 1-3 | 20 | | 20 | - | | | | | |
| | 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 34: 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 34: 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 34: 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 34: 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 34: 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 34: 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
```

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](#)
- [Summary](#)
- [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.

Summary

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.

Overview

IMPM introduces the concept of paths on a line card (IOM/IMM) which connect to planes on the chassis switch fabric ([Figure 172: IOM/IMM Paths Connecting to Switch Fabric Planes](#)).

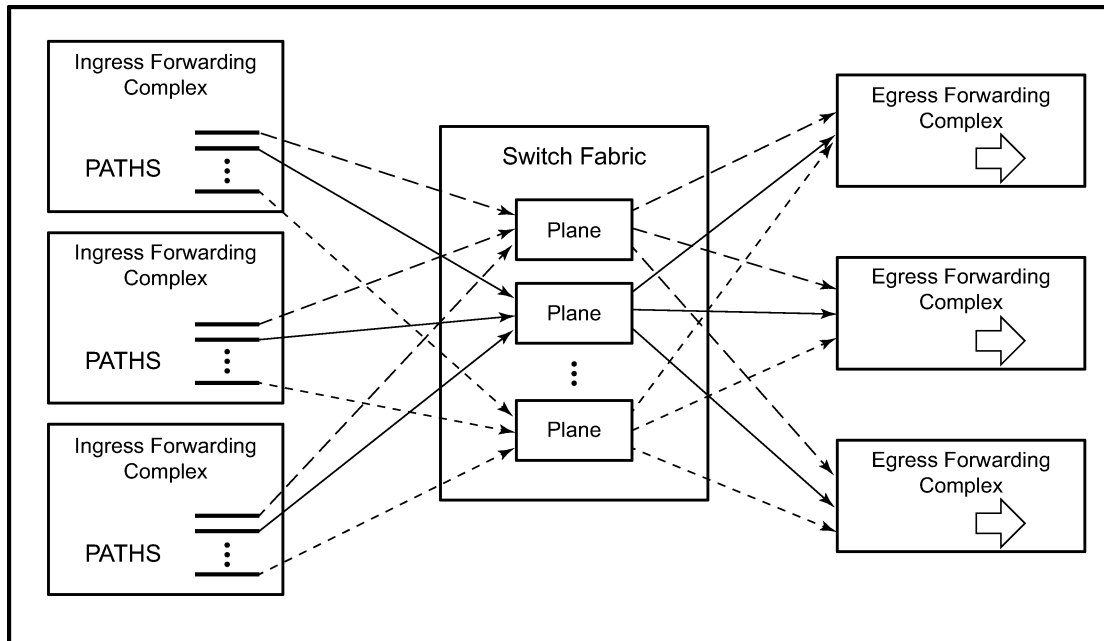
IMPM monitors the ingress rate of IP multicast channels (S,G multicast streams) on line card paths and optimizes the use of the capacity of each switch fabric plane. Its goal is to forward as many channels as possible through the system in order to make maximum use of the switch fabric planes without 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 172: 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 given0 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      Equipped      Admin   Operational      Comments
Card-type  Card-type
-----
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
=====
Cap:      Planes:
Slot/Mda  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      Equipped      Admin   Operational      Comments
Card-type   Card-type
-----
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
=====
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#
```

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
  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 dedicate 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 173: 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 173: Dynamic Bandwidth Rate Management




OSSG726

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 174: 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 174: 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 175: 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 175: 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  PLANE:
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  PLANE:
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      Equipped      Admin   Operational      Comments
Card-type  Card-type
-----
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      Equipped      Admin   Operational      Comments
Card-type  Card-type
-----
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      PoolSize
Admin ResvCBS
-----
MDA       6/1     Acc-Ing  MC Path Mgmt      18816      37632
50%
=====
*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
P     1         0        -        -    0    1         0    2000000  2000000
S     1         0        -        -    1    1         0    1800000  1800000
B     0         0        -        -    -    -         -     -        -
    
```

```

*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
    
```

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 cplx: 0 path: 0
Unmanaged traffic  0      3     primary    1     8
slot: 6 cplx: 0 path: 1
Unmanaged traffic  0      5     primary    2     8
slot: 6 cplx: 0 path: 2
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:
      PLANE:
Type SGs      InUseBW  AvailBW  TotalBw  ID  SGs      InUseBW  AvailBW  TotalBw
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:
      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      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
      default      Ancillary    0 Kbps      down
8/2   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:
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  PLANE:
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#

```

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         -         -
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
-   -         -         -         -
-   -         -         -         -
*A:PE-1#
    
```

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 cplx: 0 path: 0  0       0    none     0
Unmanaged traffic      0       0    secondary 1  8
slot: 5 cplx: 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 cplx: 1 path: 0  0       0    none     0
Unmanaged traffic      0       3    secondary 1  8
slot: 5 cplx: 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 176: IPoE Session](#).

In this way, IPoE sessions provide behavior similar to PPPoE sessions for authentication, mid-session changes, and accounting.

Figure 176: 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 177: 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 177: 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 178: 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.

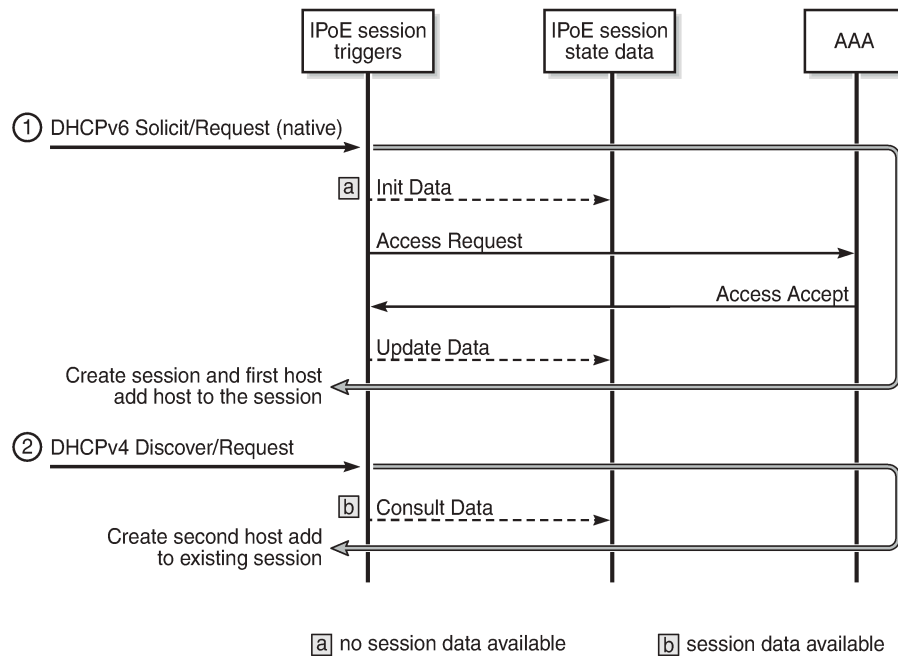
Figure 178: IPoE Session Creation Flow



25643

Figure 179: 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 179: 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 178: 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
```

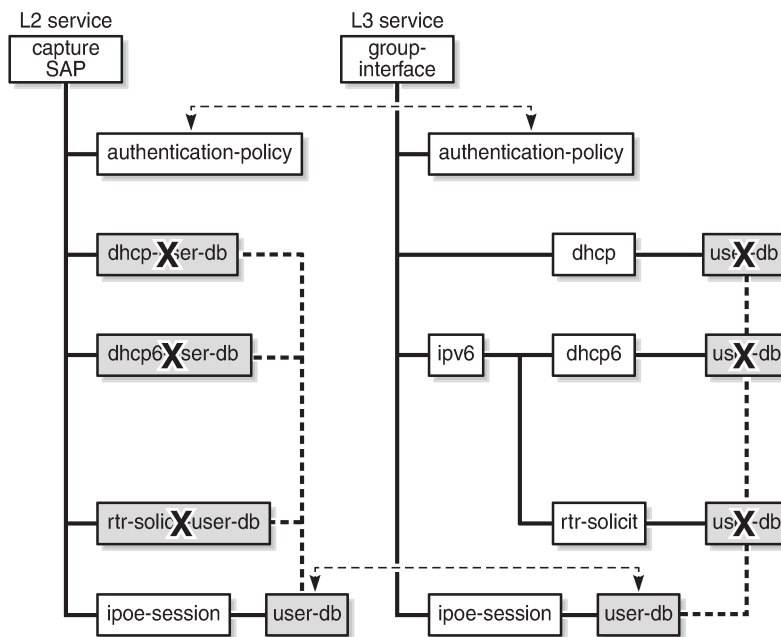
In IPoE, private-retail-subnets only apply to IPoEv6 in single-homed scenarios (no MCS). Therefore, the use cases for private-retail-subnets in combination with IPoE sessions are limited.

Practical Considerations

The rules for configuring authentication for regular, capture, and managed SAPs, also apply to IPoE sessions; see [Figure 180: 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 180: 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 181: 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 181: Baseline configuration



25645

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
  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-
inter-dest-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-
inter-dest-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
  subscriber-mgmt
    local-user-db "ludb-1"
    detail all
  exit
exit
exit

```

IPoE session failure events are also issued to log-id 99:

```

*A:BNB-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 luidb-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 : ORO (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      Stdbby
-----
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      Stdbby
-----
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 : ORO (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_SRV (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 : 5f746d6e785f696e7465726e616c5f6336c65616e7570 (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 182: Bridged CO Network Topology](#) and for a Routed CO environment in [Figure 183: Routed CO Network Topology](#).

Figure 182: Bridged CO Network Topology

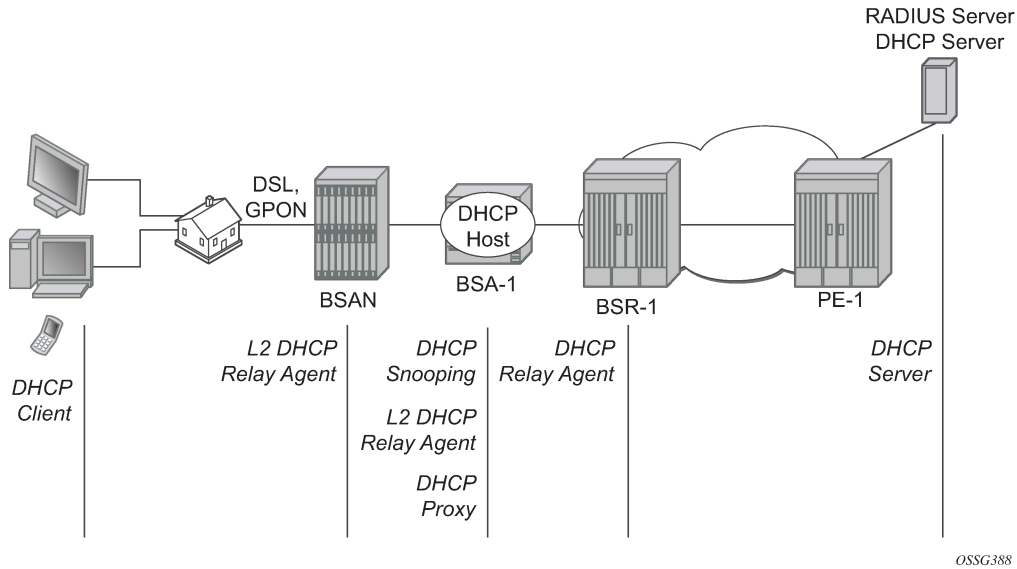



Figure 183: Routed CO Network Topology



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 184: DHCP Lease Process](#).

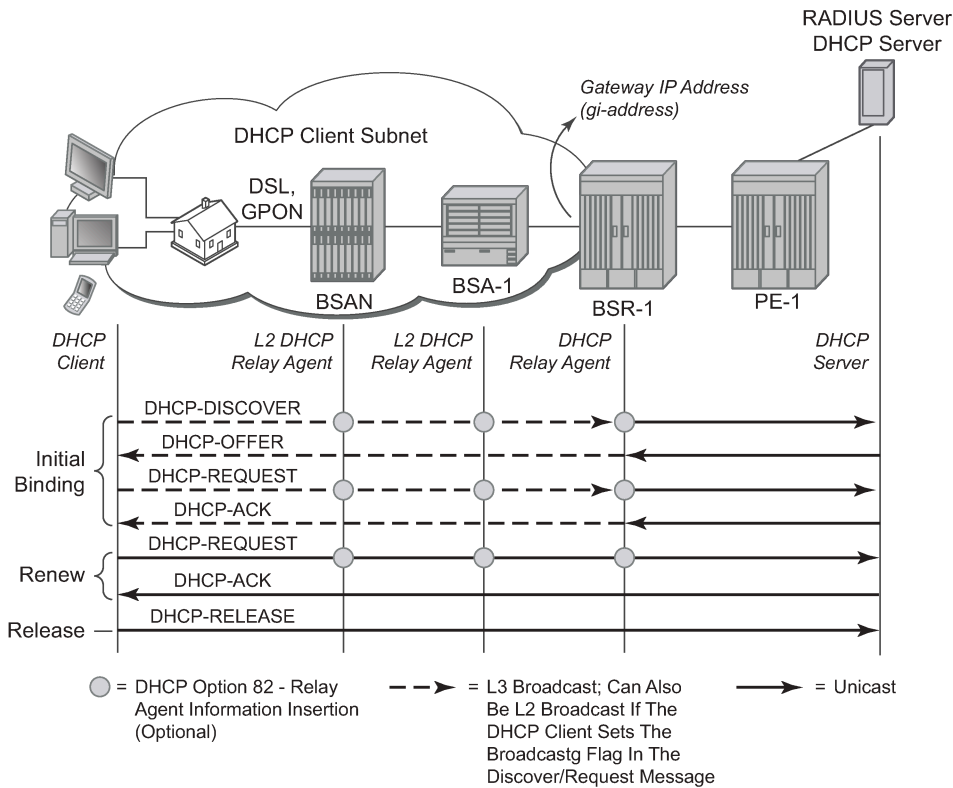
When a DHCP client boots, a DHCP discover message is broadcast on the local subnet (dest-ip = 255.255.255.255).

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 184: 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
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---
      exit
    exit
  exit
exit
```


Table 35: 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" |
| | 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
            remote-id string "Opt-82 [2] - Remote ID"
            vendor-specific-option
              system-id
              client-mac-address
              service-id
              sap-id
              string "Opt-82 [9][5] - Vendor ID"
            exit
          exit
        no shutdown
      exit
    exit
  exit
exit
```

```
# 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
            option
              action replace
              circuit-id
              remote-id string "Opt-82 [2] Remote-ID"
              vendor-specific-option
                system-id
                client-mac-address
                pool-name
                service-id
                sap-id
                string "Opt-82 [9][5] string"
            exit
          exit
        server 172.16.0.1
        trusted
      ---snip---
      gi-address 10.2.0.254
      no shutdown
```



```

        exit
    exit
exit
exit
exit
exit

```

Drop

Drop all incoming DHCP discover/request with Option 82 information present.

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 the configured options.

```

# 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 drop
exit
exit
exit
exit
exit
exit

```

```

# 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
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 CO @ 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 CO @ 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
exit
```

```
# 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
            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 36: 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 36: 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 (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 |
| 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 |

| 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 | The DHCPv4 is deleted when its uptime reaches this value |
| 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 C0"
  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

```

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
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 SVCMMGR
"SVCMMGR: 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 "DCHP 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:\submgmt.012
                    13.0           - cf3:\submgmt.011
                    12.0           - cf3:\submgmt.010
                    11.0           - cf3:\submgmt.009
                    10.0           - cf3:\submgmt.008
                    9.0            - cf3:\submgmt.007
                    8.0            - cf3:\submgmt.006
                    7.0            - cf3:\submgmt.005
                    6.0            - cf3:\submgmt.004
                    5.0            - cf3:\submgmt.003
                    4.0            - cf3:\submgmt.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 SVCNMR
"SVCNMR: 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 notstored 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
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.

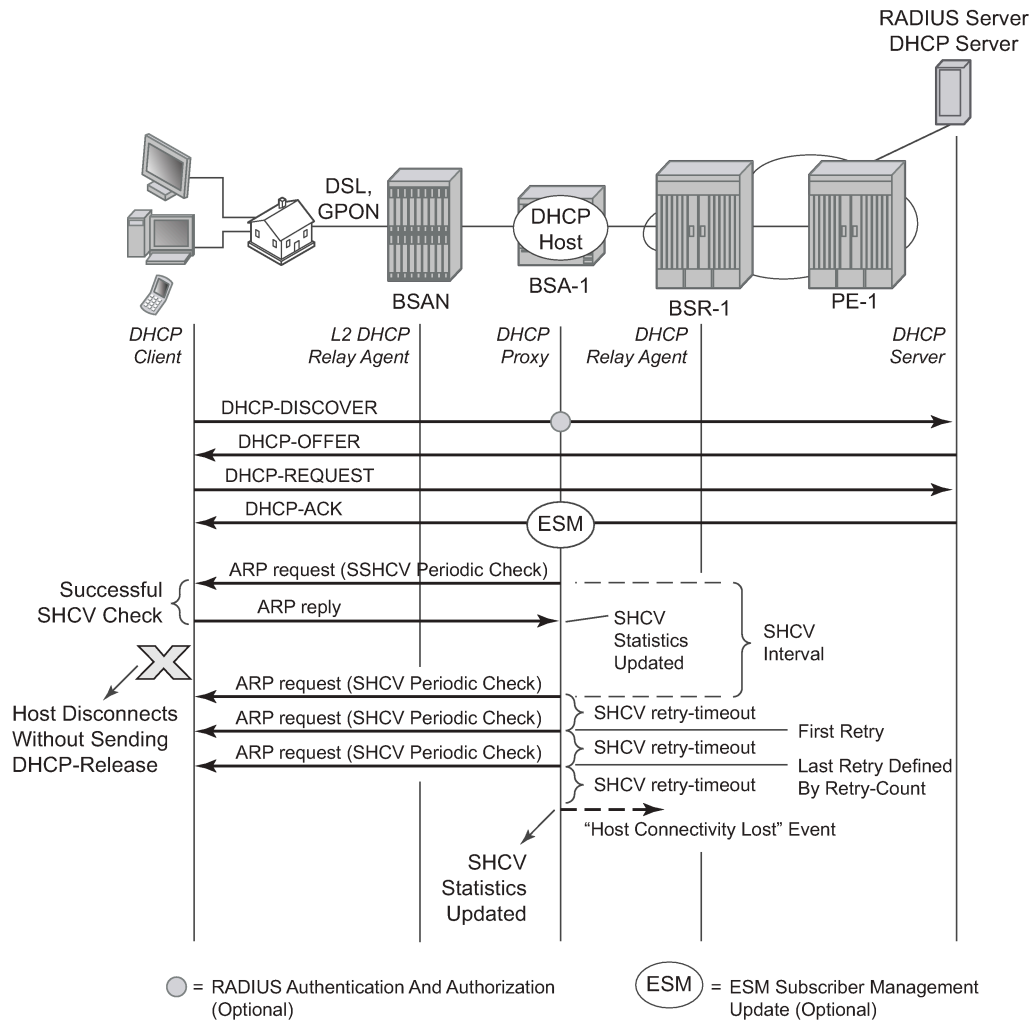
The maximum time for DHCP host connectivity loss detection in this case is:

$$((\text{host-connectivity-verify interval}) + ((\text{retry-count}) * (\text{retry-timeout})))$$

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 185: 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
    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: SVCMGR #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
  DestIp
  Oper      Last-reply/Conn-lost          MAC
-----
1       1/1/3:2              10.1.0.101
  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
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 186: 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
```



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

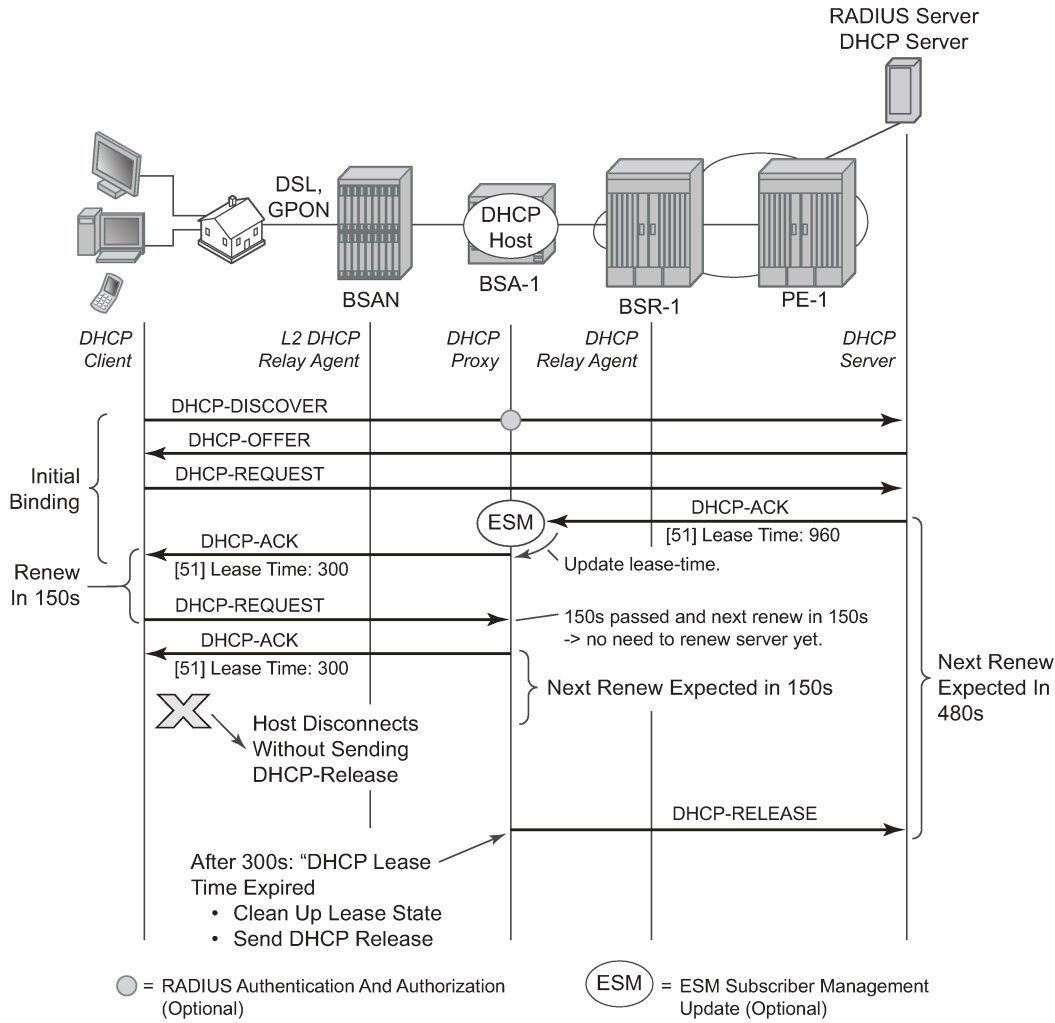
=====
*A:BSA-1#

Figure 186: 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 187: DHCP Proxy Server: Lease Split Operation, DHCP Client Disconnected

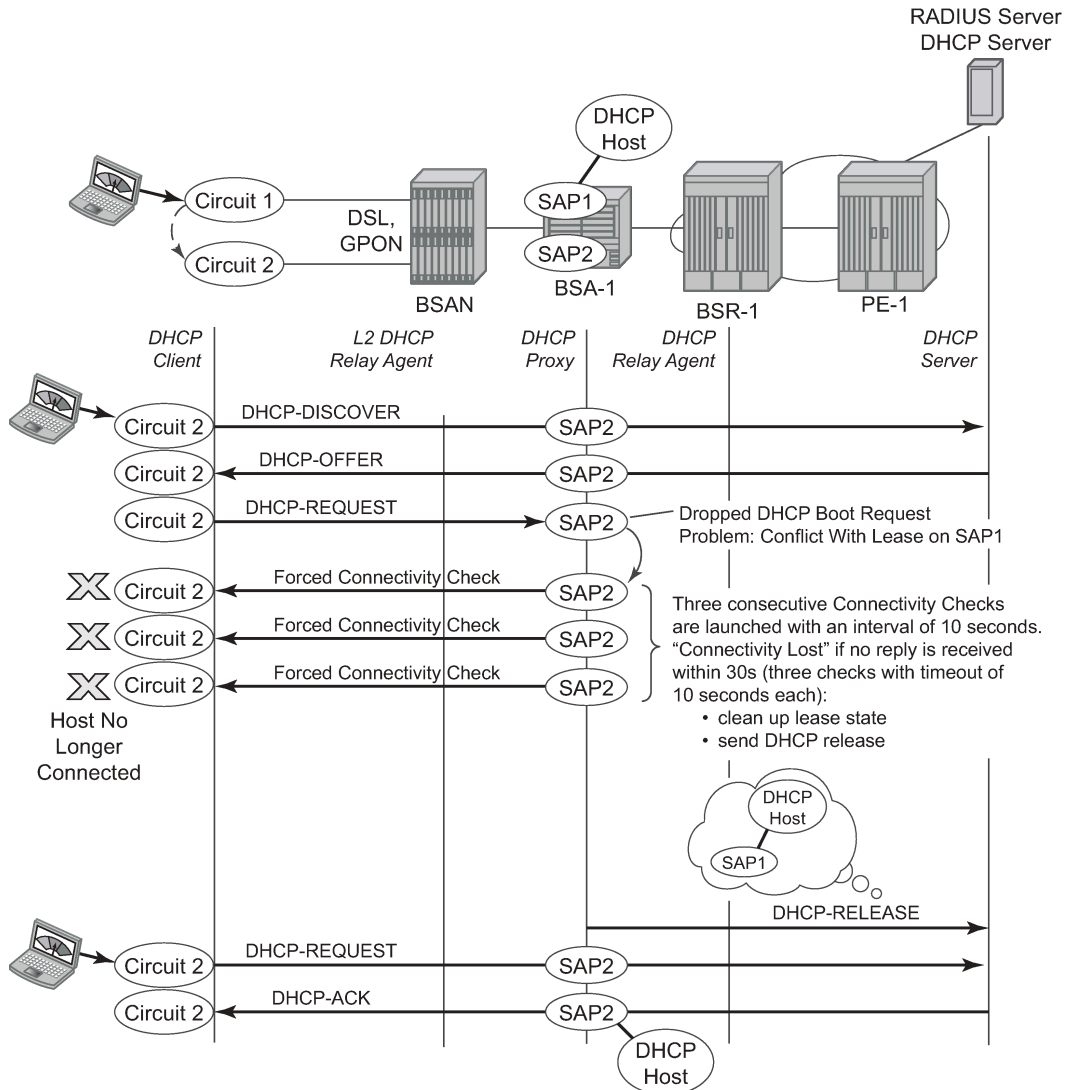


DHCP Host Mobility

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 188: DHCP Host Mobility



OSSG393

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 189: 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 189: 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 189: 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 189: 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 190: Supported L2TP reachability options](#).

Figure 190: Supported L2TP reachability options



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):




al_0513A

Table 37: 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):



al_0513B

Table 38: AVP header fields and descriptions

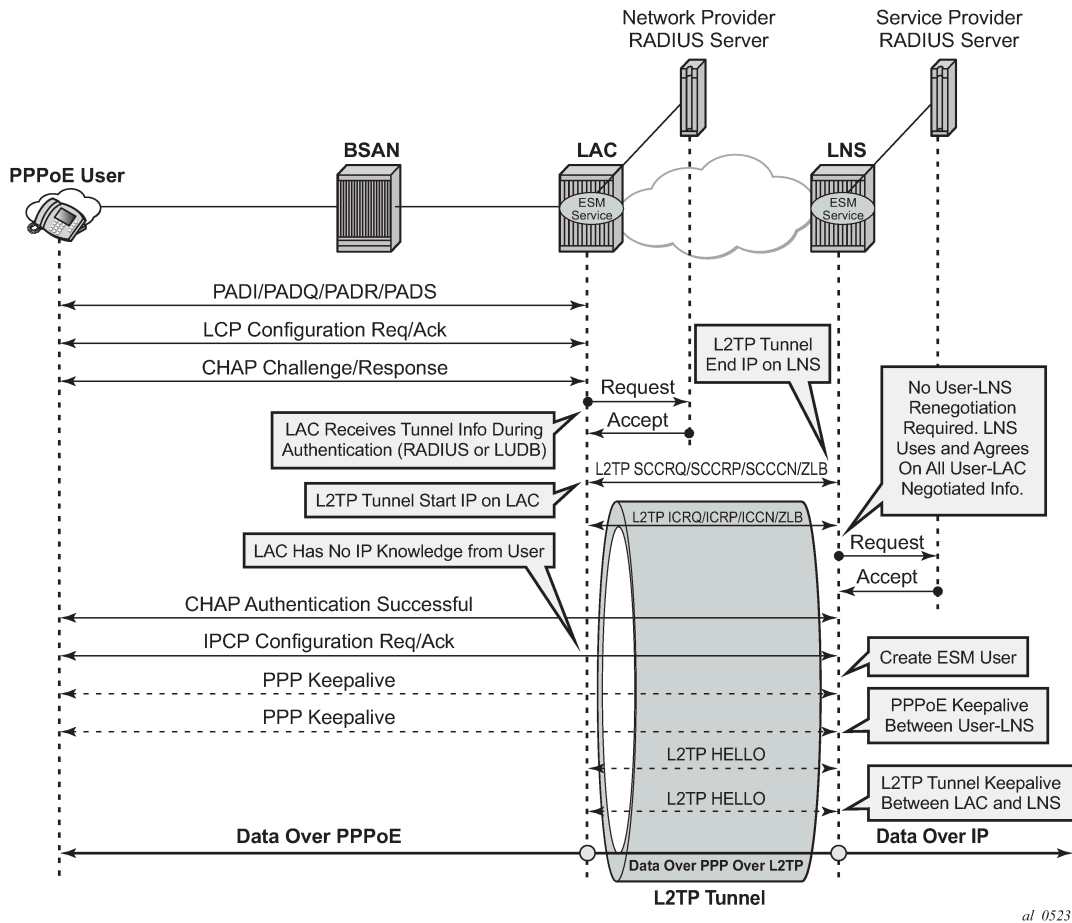
| Field | Description |
|-------|---|
| M | Mandatory bit — If the M bit is set on an unrecognized AVP within a message associated with a particular session, the session associated with this message MUST be terminated (1 bit). |
| H | Hidden bit — Identifies the hiding of data in the Attribute-Value 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 191: 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.

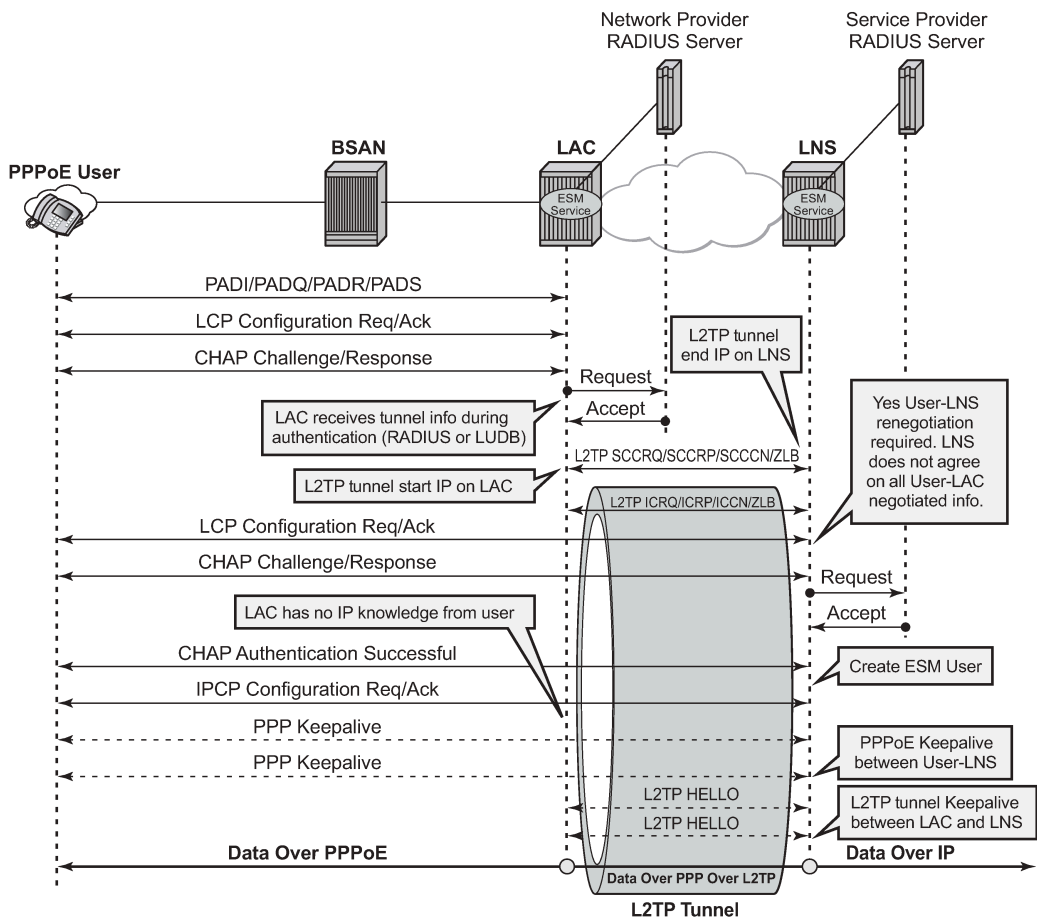
Figure 191: RADIUS-triggered tunnel/session setup without LNS renegotiation



RADIUS-triggered tunnel/session setup with LNS renegotiation

Figure 192: 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 191: RADIUS-triggered tunnel/session setup without LNS renegotiation.

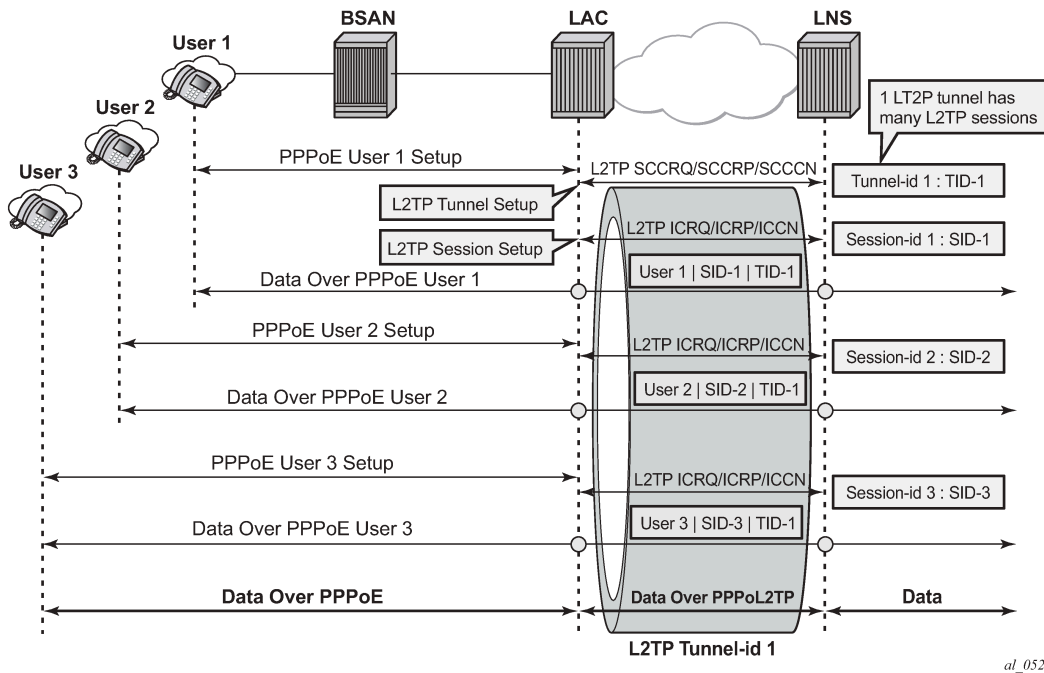
Figure 192: RADIUS-triggered tunnel/session setup with LNS renegotiation



Running multiple PPP sessions over a single L2TP tunnel

Figure 193: 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).

Figure 193: Running multiple PPP sessions over a single L2TP tunnel




al_0522

PPP user-initiated release/terminate

Figure 194: PPP user-initiated release/terminate shows the user initiated terminate_request tunneled by the LAC followed by the user initiated PADT terminated on the LAC. The LAC informs the LNS about the termination of the session via the L2TP CDN message. The L2TP tunnel can be optionally (idle-timeout) terminated via the L2TP StopCCN message.

Figure 194: PPP user-initiated release/terminate



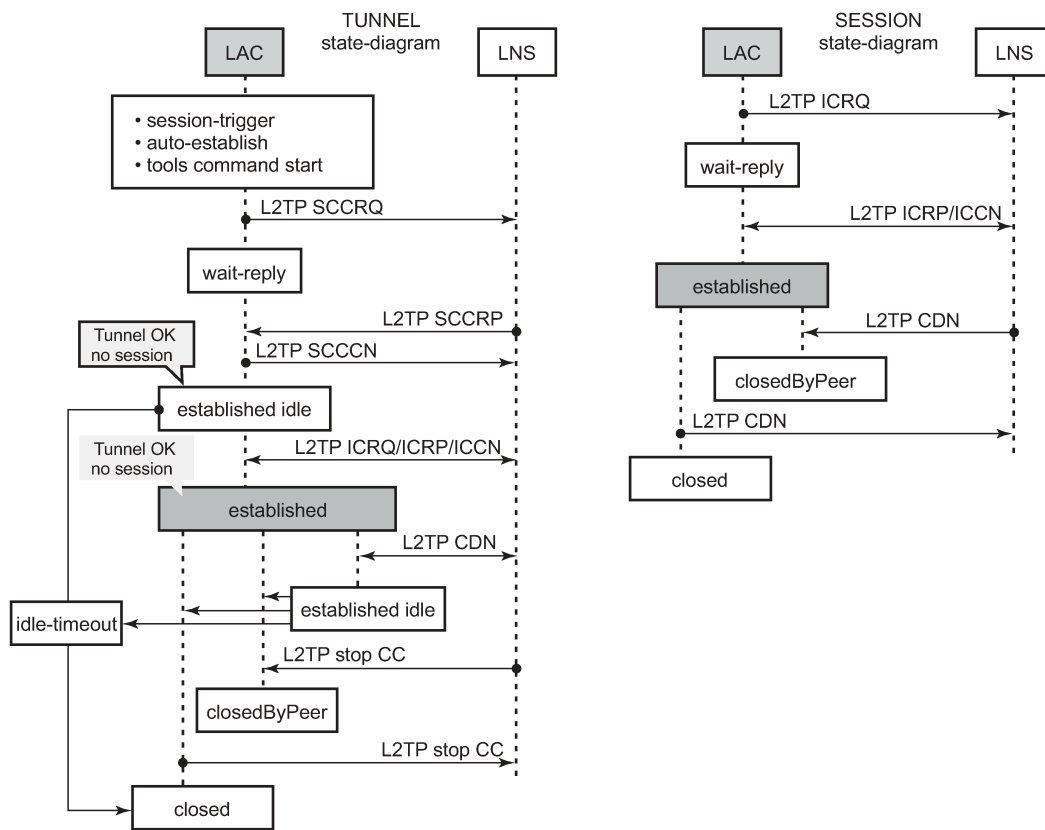
al_0525

L2TP tunnel/session state diagram

Figure 195: 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 195: 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 39: 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 40: 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. 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 196: 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
```

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.

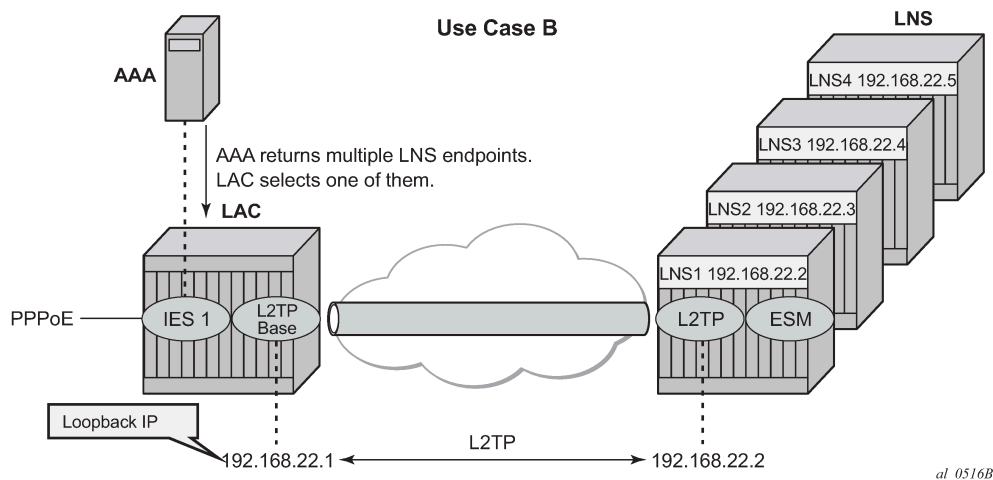
Missing L2TP parameters are taken from defaults defined in the router l2tp context.

```
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 197: 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",

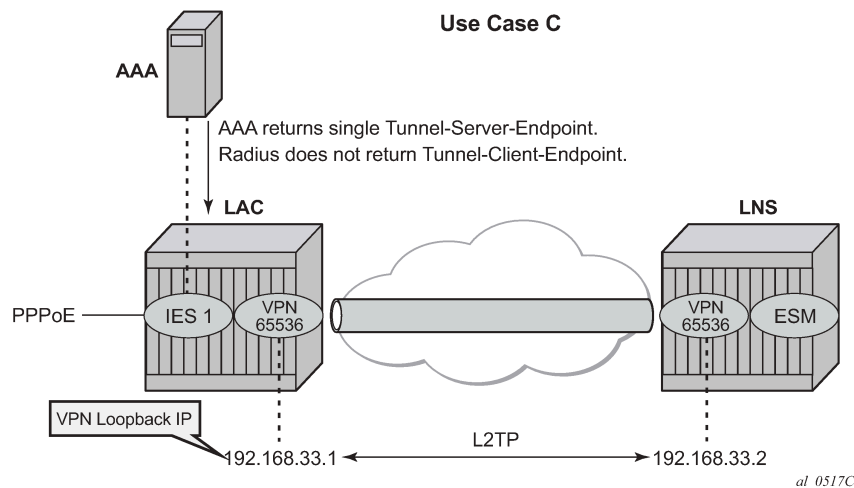
---snip---
```

This scenario shows the PPPoE session termination (base IES service 1) and the L2TP tunnel setup in the base router instance.

VRF hosted LAC

Figure 198: 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 198: 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

In use case D, the L2TP tunnel-group-name is assigned during RADIUS authentication.

Figure 199: 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
      exit
    exit
  exit
  no shutdown
exit
no shutdown
exit
no shutdown
exit
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

In use case E, the L2TP tunnel-group-name is assigned during LUDB authentication, so this essentially is a RADIUS-less scenario.

Figure 200: LUDB returns L2TP tunnel group



The PPPoE user enters on an IES service 1, sap 1/1/3:100, and is authenticated via the LUDB which provides L2TP wholesale/retail and ESM information. The PPPoE context refers to a local-user database *l2tp* to provide the subscriber authentication and the tunnel setup parameters, so no RADIUS is needed.

```
configure
service
  ies 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
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 unnamed | 1 | |
| ----- | | |
| 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     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 Tun | Ses Total | Active Ses | Total |
|--------------|-------|--------------|------|-----|------------|-----------|------------|-------|
| 192.168.22.2 | | | 1701 | 0 | 0 | 0 | 0 | 0 |
| 192.168.22.3 | | | 1701 | 1 | 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 |
| | PPP Sessions | - PPPoEoA | 0 0 |
| | PPP Sessions | - PPPoA | 0 0 |
| | 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
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
PADO Delay            : 0msec
Diameter app policy   : (Not Specified)
Diameter auth policy  : (Not Specified)
Force IPv6CP          : Disabled
Ignore DF Bit         : Disabled

DHCPv6 lease times
  Renew timer         : > 9999 days
  Rebind timer        : > 9999 days
  Preferred lifetime  : 0d 00:00:00
  Valid lifetime      : 0d 00:00:00

Identification Strings (option 254)
  Subscriber Id       : user@wholesale.com
  SLA Profile String  : sla-profile-1
  Sub Profile String  : sub-profile-1

L2TP
  Service             : 65536
  Tunnel Group        : wholesale.com

MSAP defaults

Filter Overrules

Access loop info
=====
*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-Stdby
-----
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                               1
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
no shutdown
exit
no shutdown
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
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 TO (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 TO (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 201: 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 | unnamed | | | | 1 |
| No. of tunnels: 1 | | | | | |

Figure 201: L2TP keepalive mechanism

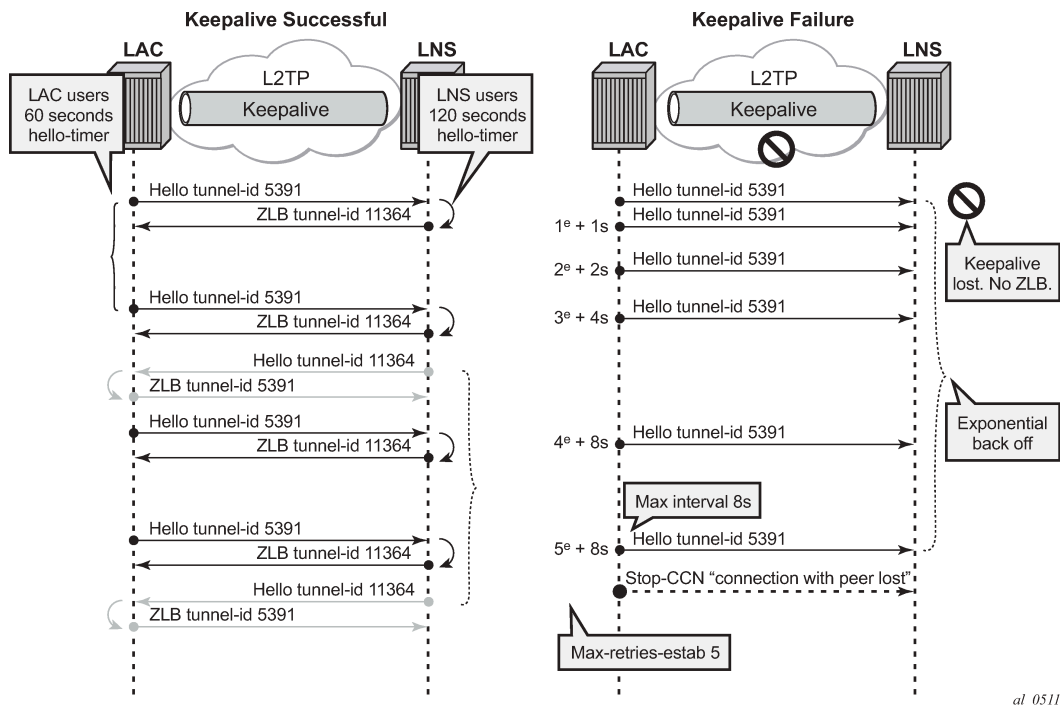


Figure 201: L2TP keepalive mechanism shows the tunnel being closed after 5 unsuccessful Hello deliveries with error-message *connection with peer lost*.

```
*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.


Figure 202: Floating peers accept



al_0512B

- Ignore — Specifies that this system ignores any source IP address change for received L2TP control messages, does not learn any new peer IP address and does not change the destination address in subsequent L2TP messages.


Figure 203: Floating peers ignore



al_0512C

- Reject — Specifies that this system rejects any source IP address change for received L2TP control messages and drops those messages.

Figure 204: 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 SCCRQ
Peer UDP   : 1701
Tx dst-IP  : 192.0.2.12 # (3) peer address used in SCCCN
Tx dst-UDP : 1701
Rx src-IP  : 192.0.2.12 # (2) SCCRP different IP received
Rx src-UDP : 1701
---snip---
```

Tx/Rx connect speed - AVP 24/38

The connect speed (TX AVP 24 and RX AVP 38) is passed in the ICCN messages sent from the LAC to the LNS. The L2TP AVP 24 defines the (Tx) connect speed in bps from the perspective of traffic flowing from the LAC towards the subscriber (BNG downstream rate). The L2TP AVP 38 defines the (Rx) connect speed in bps from the perspective of traffic flowing from the subscriber 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     Assignment
-----
1023868928 15623     9957     established     not-blacklisted  1
wholesale.com
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     Assignment
-----
1023868928 15623     9957     draining        not-blacklisted  1
wholesale.com
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     Assignment
-----
1023868928 15623     9957     drained         not-blacklisted  0
wholesale.com
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.

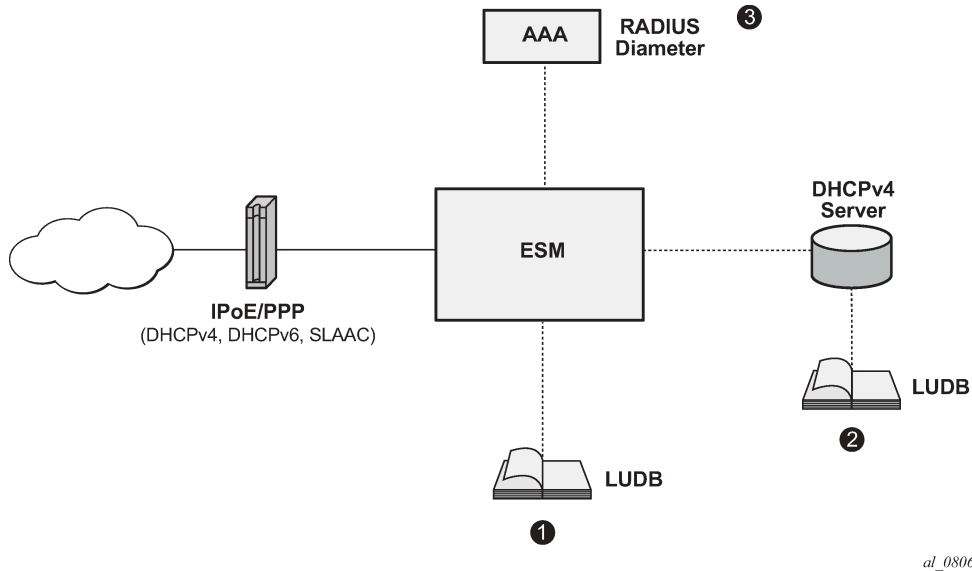
Overview

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 205: 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 205: LUDB applications



The LUDB lookup process is common to the applications shown in [Figure 205: 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.

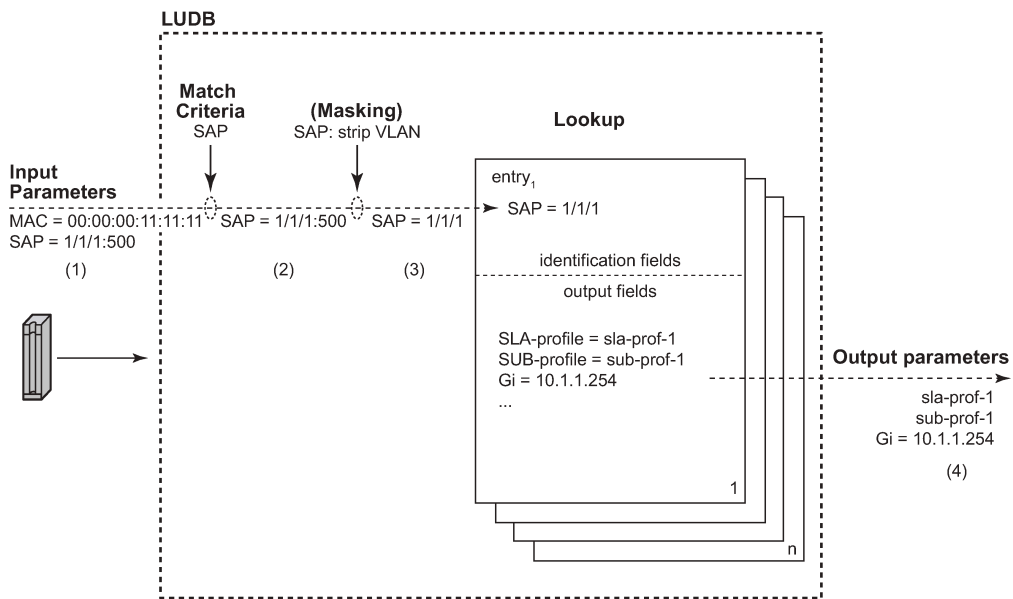
The LUDB lookup process translates a set of input parameters (the host identification fields) into a set of output parameters; see [Figure 206: Processing an LUDB lookup request](#) for the following example:

- 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 206: Processing an LUDB lookup request



al_0807

Configuration

Creating LUDBs

An LUDB is identified by a name of 32 characters maximum ([Figure 207: 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.

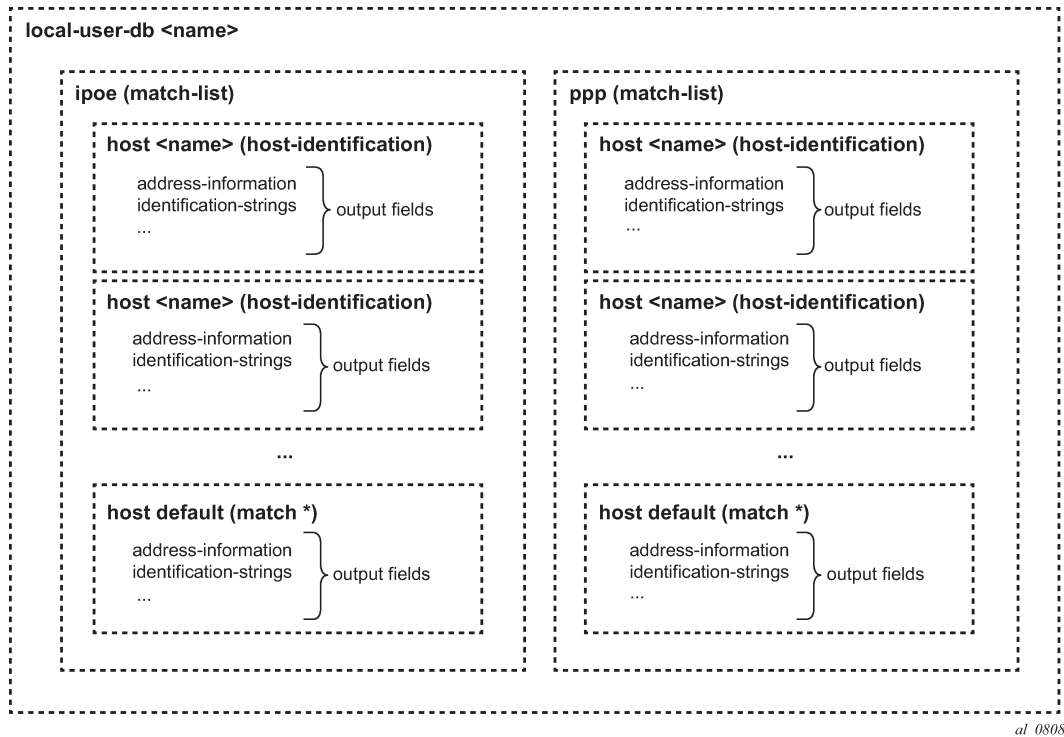
An LUDB can provide the data for IPoE (DHCPv4, DHCPv6, and SLAAC) as well as for PPP users.

```
*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 207: Creating LUDBs and LUDB entries



al_0808

Creating host entries

A host entry is identified by name of 32 characters maximum ([Figure 207: 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: DHCPDS #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: DHCPDS #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: DHCPDS #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 41: Masking examples](#) illustrate masking. For the third example, a combination of both prefix and suffix matching is used.

Table 41: 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.

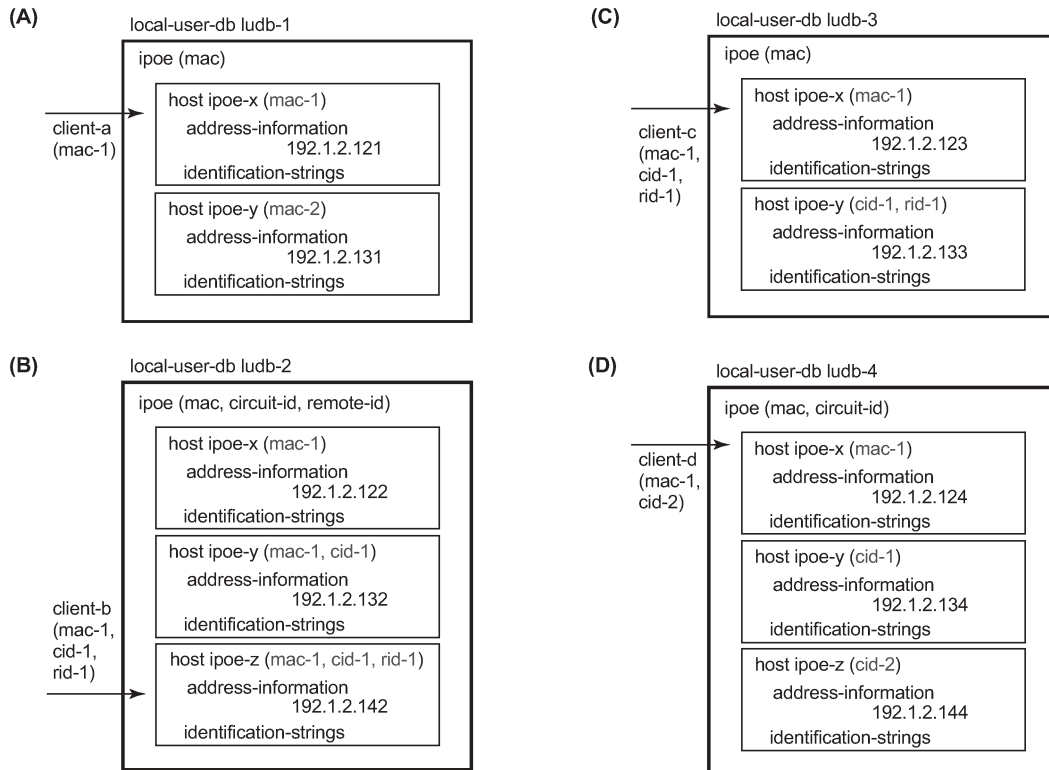
3. As many of the host-identification fields as possible must be matched, while still obeying rule 1.

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 208: Host matching examples

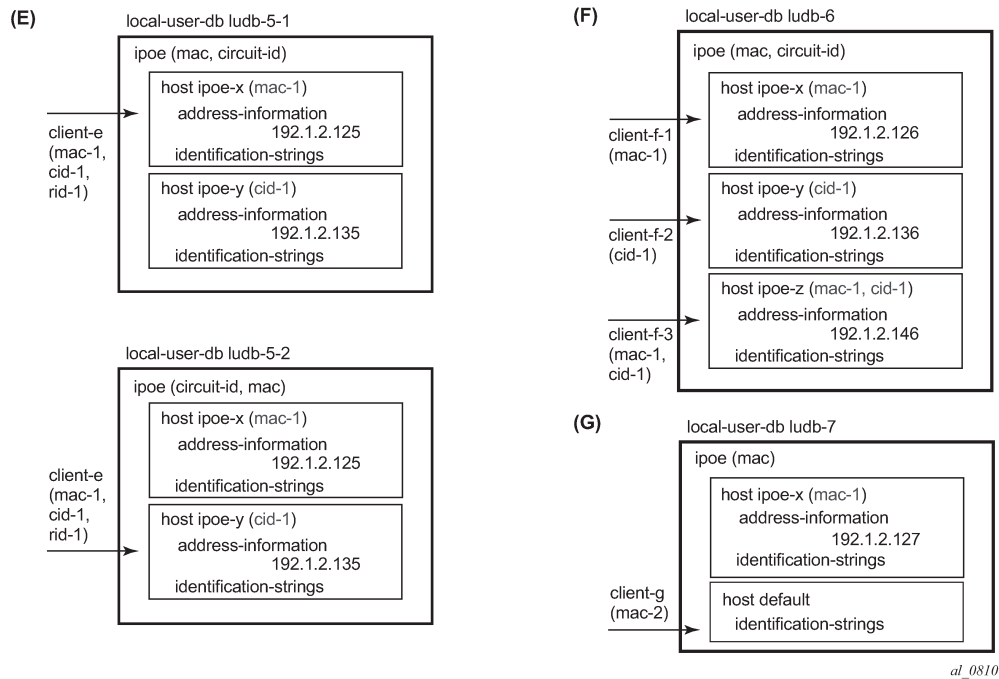


al_0809

The examples in [Figure 208: Host matching examples](#) and [Figure 209: 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 209: 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 LADB
"LADB: 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 LADB
"LABD: 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 labd-4"
```

Example 5: VLAN range

LABD 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: DHCPS #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: DHCPS #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: DHCPS #1103 User data base still referenced

*A:BNG-1>config>subscr-mgmt# no local-user-db "ludb-11"
MINOR: DHCPS #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 State   Matched objects
-----
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 State   Matched objects
-----
entry-55               Up           userName
---snip---
-----
Number of PPP Hosts : 5
=====
*A:BNB-1#
    
```

To find the places where a specific LUDB is applied, use the following command.

```
*A:BNB-1# show subscriber-mgmt local-user-db "ludb-1" association
```

```
=====
```

```
DHCP Server associations for ludb-1
```

```
=====
```

| Server-Name | Router-Name |
|-------------|-------------|
| 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 |

```
-----
```

```
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 |

```
-----
```

```
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 |

```
-----
```

```
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 OVERRULES
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:BNG-1#
    
```

```

*A:BNG-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:BNG-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
    
```

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 ladb-2"
```

```
37 2018/11/20 12:18:59.141 CET MINOR: DEBUG #2001 Base LADB  
"LADB: User lookup failed  
  Problem: user db is shutdown"
```

Conclusion

In this chapter general LADB concepts are explained. LADBs 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

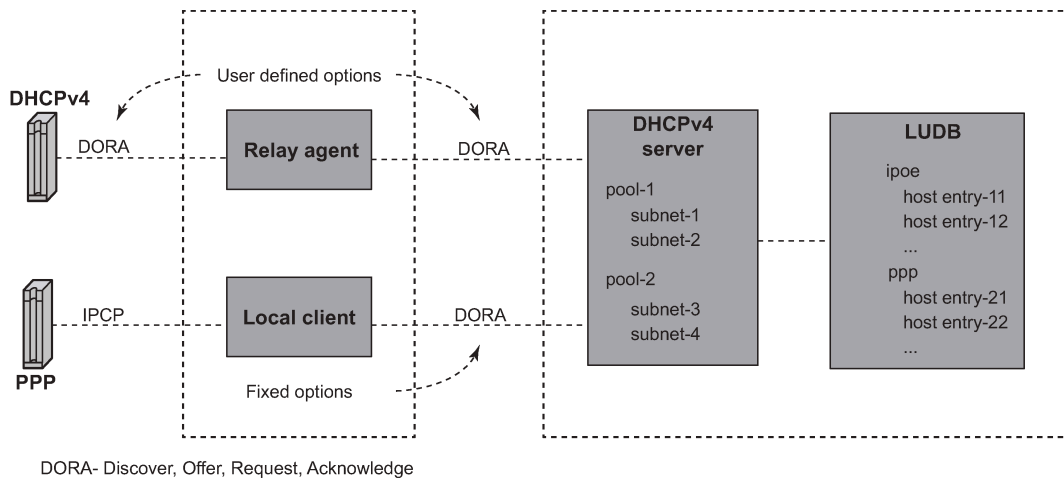
A local DHCPv4 server can be used for IPoE users as well as for PPP users (see [Figure 210: LUDB access via a DHCPv4 server](#)).

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 210: 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/VPRN 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/VPRN service ID)
 - [4] SAP ID (not included for retail VPRN 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.

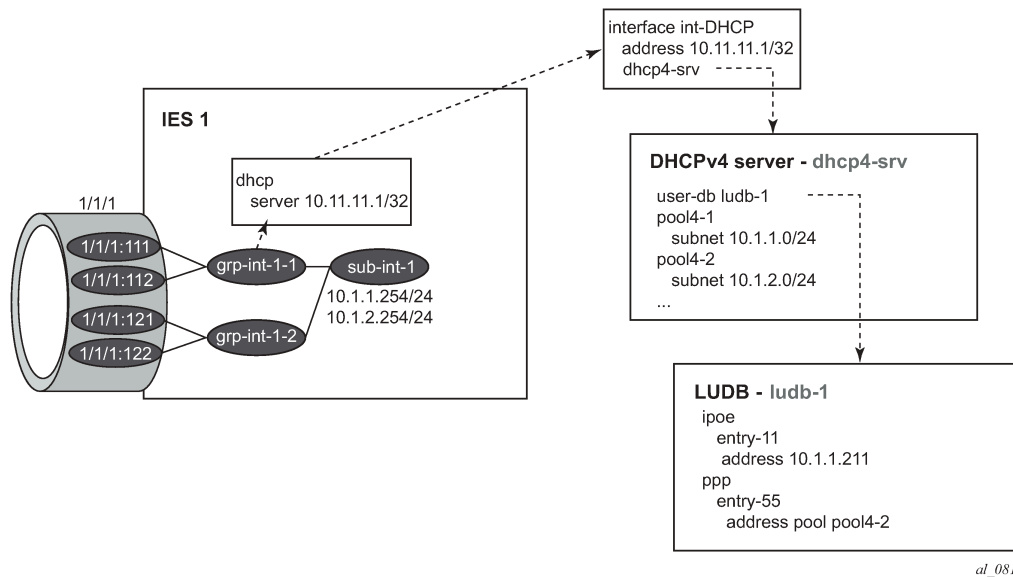
Other parameters

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 211: Example configuration shows the example configuration used in this chapter.

Figure 211: 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
        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
            exit
        exit
    exit

```

```

        no shutdown
    exit
exit
---snip---
pppoe
    session-limit 100
    sap-session-limit 100
    no shutdown
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 212: 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 ludb-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

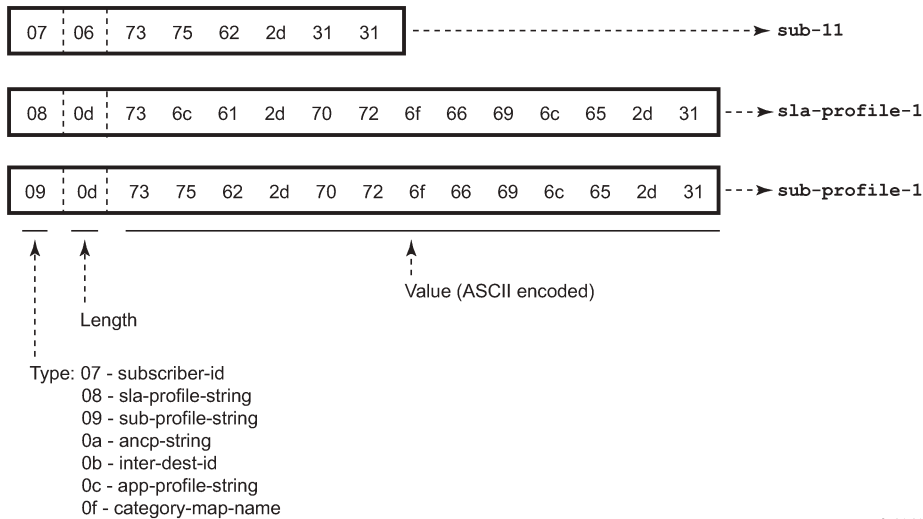
```

```

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 212: Decoding the ESM user option



al_0813

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: ALU7XXXSBM
[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 via Local Address Server of the Advanced Configuration Guide 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 213: LUDB authentication](#):

- DHCPv4 Discover
- DHCPv6 Solicit
- PPPoE PADI
- PPPoA Conf Req
- Router Solicit [RS]

Figure 213: LUDB authentication



(*) – indirect trigger only, via RADIUS fallback

25536

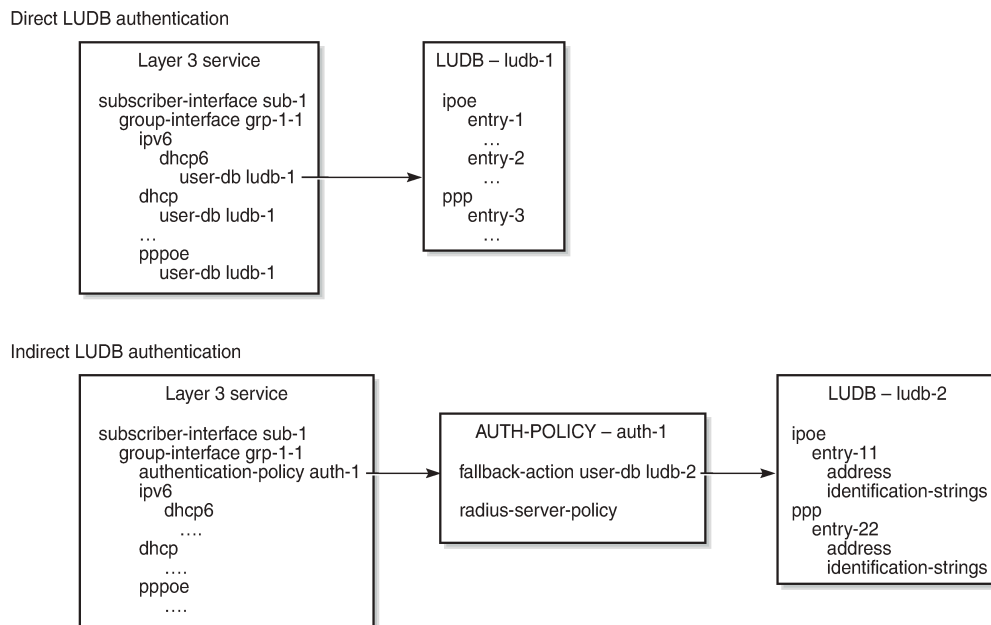
When triggered, ESM can directly access an LUDB because the LUDB is applied to the service directly (through one of its sub-contexts), or indirectly as a fallback action for RADIUS (through the authentication policy); see [Figure 214: Direct and indirect LUDB authentication](#). ARP requests can only trigger LUDB authentication indirectly.



Note:

An authentication policy can be referenced from a group interface and a capture SAP, or from an LUDB.

Figure 214: 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 215: 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 216: 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 215: LUDB parameters for IPoE](#) and [Figure 216: LUDB parameters for PPPoE](#) for the full list of input and output parameters for IPoE and PPPoE, respectively.

Figure 215: LUDB parameters for IPoE

| Parameter | | 1 | 2 | 3 | 4 |
|----------------------------|------------------------|---|---|---|---|
| <u>host-identification</u> | <u>circuit-id</u> | | | | |
| | <u>derived-id</u> | | | | |
| | <u>encap-tag-range</u> | | | | |
| | <u>ip-prefix</u> | | | | |
| | <u>mac</u> | | | | |
| | <u>option60</u> | | | | |
| | <u>remote-id</u> | | | | |
| | <u>sap-id</u> | | | | |
| | <u>service-id</u> | | | | |
| | <u>string</u> | | | | |
| | <u>system-id</u> | | | | |

| Parameter | | 1 | 2 | 3 | 4 |
|-------------------------------|-----------------------------|---|---|---|---|
| <u>identification-strings</u> | <u>ancp-string</u> | | | | |
| | <u>app-profile-string</u> | | | | |
| | <u>category-map-name</u> | | | | |
| | <u>inter-dest-id</u> | | | | |
| | <u>sla-profile-string</u> | | | | |
| | <u>spi-sharing-group-id</u> | | | | |
| | <u>sub-profile-string</u> | | | | |
| | <u>subscriber-id</u> | | | | |

| | |
|--|--|
| | 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 |

| Parameter | | 1 | 2 | 3 | 4 |
|------------------------------------|-----------------------------------|---|---|---|---|
| <u>acct-policy</u> | | | | | |
| <u>address</u> | <u>ip-address</u> | | | | |
| | <u>gi-address</u> | | | | |
| | <u>pool</u> | | | | |
| | <u>use-from-pool-client</u> | | | | |
| - | <u>auth-domain-name</u> | | | | |
| - | <u>auth-policy</u> | | | | |
| - | <u>ipv6-address</u> | | | | |
| - | <u>ipv6-delegated-prefix</u> | | | | |
| - | <u>ipv6-delegated-prefix-len</u> | | | | |
| - | <u>ipv6-delegated-prefix-pool</u> | | | | |
| - | <u>ipv6-slaac-prefix</u> | | | | |
| - | <u>ipv6-wan-address-pool</u> | | | | |
| <u>diameter-application-policy</u> | | | | | |
| <u>diameter-auth-policy</u> | | | | | |
| <u>gi-address</u> | | | | | |
| <u>link-address</u> | | | | | |
| <u>ipv6-lease-times</u> | <u>preferred-lifetime</u> | | | | |
| | <u>rebind-timer</u> | | | | |
| | <u>renew-timer</u> | | | | |
| | <u>valid-lifetime</u> | | | | |
| <u>ipv6-slaac-prefix-pool</u> | | | | | |
| <u>match-radius-proxy-cache</u> | | | | | |
| <u>msap-defaults</u> | <u>group-interface</u> | | | | |
| <u>msap-defaults</u> | <u>policy</u> | | | | |
| <u>msap-defaults</u> | <u>service</u> | | | | |
| <u>options</u> | <u>custom-options</u> | | | | |
| | <u>default-router</u> | | | | |
| | <u>dns-server</u> | | | | |
| | <u>domain-name</u> | | | | |
| | <u>lease-rebind-time</u> | | | | |
| | <u>lease-renew-time</u> | | | | |
| | <u>lease-time</u> | | | | |
| | <u>netbios-name-type</u> | | | | |
| | <u>netbios-node-type</u> | | | | |
| <u>subnet-mask</u> | | | | | |
| <u>options6</u> | <u>dns-server</u> | | | | |
| - | <u>retail-service-id</u> | | | | |
| <u>rip-policy</u> | <u>server</u> | | | | |
| <u>server6</u> | | | | | |
| <u>to-client-options</u> | <u>ipv4</u> | | | | |
| | <u>ipv6</u> | | | | |

Figure 216: LUDB parameters for PPPoE

| Parameter | | 1 | 2 | 3 |
|----------------------------|-----------------|---|---|---|
| host-identification | circuit-id | | | |
| | derived-id | | | |
| | encap-tag-range | | | |
| | mac | | | |
| | remote-id | | | |
| | sap-id | | | |
| | username | | | |

| Parameter | | 1 | 2 | 3 |
|-------------------------------|----------------------|---|---|---|
| identification-strings | ancc-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 | | | |
| | pl-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 | | | |
| options6 | netbios-name-server | | | |
| | 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 217: 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 217: 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 218: LUDB authentication for capture and managed SAPs](#).

Figure 218: 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 215: LUDB parameters for IPoE](#) and [Figure 216: 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 215: LUDB parameters for IPoE](#) and [Figure 216: 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.

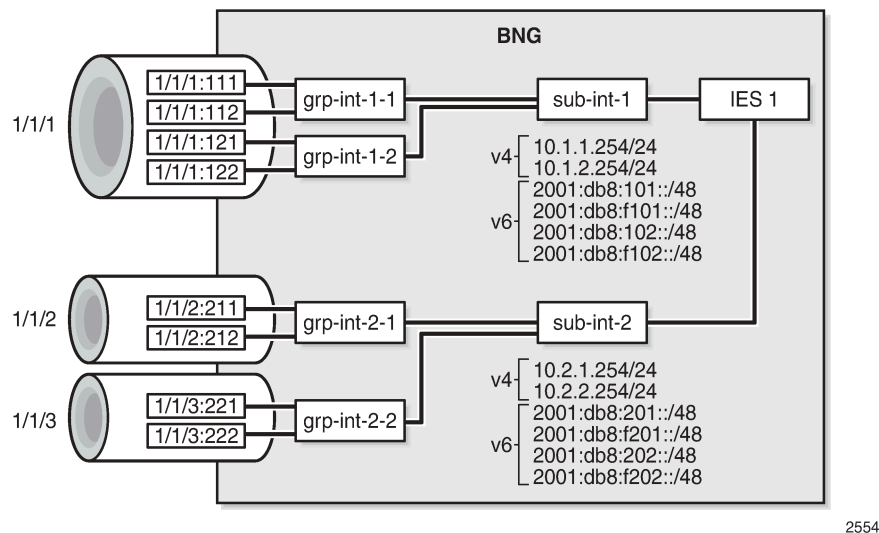
- If an LUDB is applied in the ppp or pppoe context of a group interface or capture SAP, PAP or CHAP authentication is based on the password configured in the entry. If no password is required, the password parameter in the LUDB entry must be explicitly set to ignore.

A password verification failure leads to a setup failure.

Configuration

Figure 219: 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 219: 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
```

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      Fwding
1/1/1:111    sub-11
2001:db8:102:11::11/128
00:00:00:11:11:11  N/A      IPoE-DHCP6  Fwding
1/1/1:111    sub-11
2001:db8:f102:1100::/56
00:00:00:11:11:11  N/A      IPoE-DHCP6  Fwding
-----
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
    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 LADB
"LABD: 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 LADB
"LABD: 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 LADB is accessed for every incoming message. In a proxy case, the LADB 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 LADB is accessed four times per host.

The command to list the active subscriber hosts is the same as for the IPE session model, and is not repeated here. The same applies to the other commands providing origin codes.

PPPoE authentication

In this example, the LADB *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
    
```

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
  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 42 4e 47 2d 31"
    
```

```
38 2018/11/22 12:52:59.420 CET MINOR: DEBUG #2001 Base PPPoE
"PPPoE: RX Packet
  IES 1, SAP 1/1/1:111

  DMAC: 02:17:01:01:00:01
  SMAC: 00:00:00:55:55:55
  Ether Type: 0x8864 (Session)

  PPPoE Header:
  Version: 1                Type      : 1
  Code   : 0x00             Session-Id: 0x0001 (1)
  Length : 36

  PPP:
  Protocol : 0xc223 (CHAP)
  Code     : 2 (Response)
  Identifier: 1             Length    : 34

  Value-Size: 16
  Value      : c5 02 13 0e 6c bf f4 58 61 51 e8 92 91 7c 53 94
  Name       : "sub55@domain1"

  Hex Packet Dump:
  11 00 00 01 00 24 c2 23 02 01 00 22 10 c5 02 13 0e 6c bf f4 58 61 51 e8 92
  91 7c 53 94 73 75 62 35 35 40 64 6f 6d 61 69 6e 31 00 00 00 00"

39 2018/11/22 12:52:59.420 CET MINOR: DEBUG #2001 Base LADB
"LABD: User lookup success - host found
  user-name:
    original: sub55@domain1
    masked:   sub55@domain1

  Host entry-55 found in user data base ladb-rsap"

40 2018/11/22 12:52:59.420 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 : 33

  PPP:
  Protocol : 0xc223 (CHAP)
  Code     : 3 (Success)
  Identifier: 1             Length    : 31

  Message: "CHAP authentication success"

  Hex Packet Dump:
  11 00 00 01 00 21 c2 23 03 01 00 1f 43 48 41 50 20 61 75 74 68 65 6e 74 69
  63 61 74 69 6f 6e 20 73 75 63 63 65 73 73"

---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        PPPoE-SID Origin   Fwding State
  MAC Address
-----
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 Lease      MC
                :                :                : LeaseTime  Origin  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 218: 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 217: LUDB authentication for regular SAPs](#) and [Figure 218: 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
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
service
id 3
dhcp
mode egr-ingr-and-dropped
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 IPoE 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 215: LUDB parameters for IPoE](#) and [Figure 216: 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 DHCP application (DHCP #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 220: Network Topology](#). This chapter uses the RCO model with PPPoE, IPv4, and RADIUS authentication for demonstrating MSAPs.

Figure 220: Network Topology



OSSG412

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
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
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
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-ISPI" 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
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
    exit
no shutdown

```

```

exit
exit
exit
    
```

Verify the details of capture SAP:

```

*A:BNB-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      : 474539          Packets          : 474539          Octets           : 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.

VPRN 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
      no shutdown
    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
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-Stdby
-----
[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-Stdby
-----
[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
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
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: ""
Icmp 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 IP-Address | Adm | Opr(v4/v6) | Mode | Port/SapId PfxState |
|-------------------------------|-----|-----------------|---------|------------------------|
| grp-int-1 | Up | Up /Down | VPRN G* | 1/2/2 |
| sub-int-1 10.255.255.254/8 | Up | Up/Down | VPRN S* | subscriber 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-ISPI
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:BNB-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
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
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
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: SVCNMR #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: SVCNMR #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
    id 2
      ppp
        packet
          mode egr-ingr-and-dropped
          detail-level medium
          discovery
          ppp
          dhcp-client
        exit
      exit
    exit
  exit
exit

```

```

configure
  log
    log-id 1

```

```

        from debug-trace
        to session
    exit
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:BNB-1#

```

The MSAP policy associations can be checked as follows:

```

*A:BNB-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:BNB-1#

```

All MSAPs created and associations with the services can be checked as follows:

```

*A:BNB-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:BNB-1#

```

It is possible to use a **tools** command to update an existing MSAP when a specific msap-policy has changed.

```

*A:BNB-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](#)
- [Summary](#)
- [Overview](#)
- [Configuration](#)
- [Conclusion](#)

Applicability

The information and configuration in this chapter is based on SR OS Release 7.0.R5.

Summary

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 Triple Play Guide.

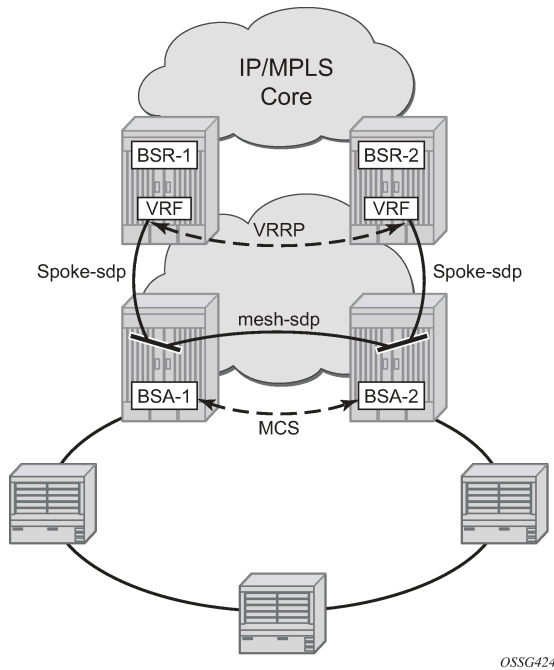
Overview

MC-Ring

[Figure 221: 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 Distribution 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 221: MC-Ring Layer 2 CO Dual Homing



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 221: 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/VPRN 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 subdivide 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 222: 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 222: Dual homing Under Steady-State Condition



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 222: 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 223: 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 223: Broken Ring State



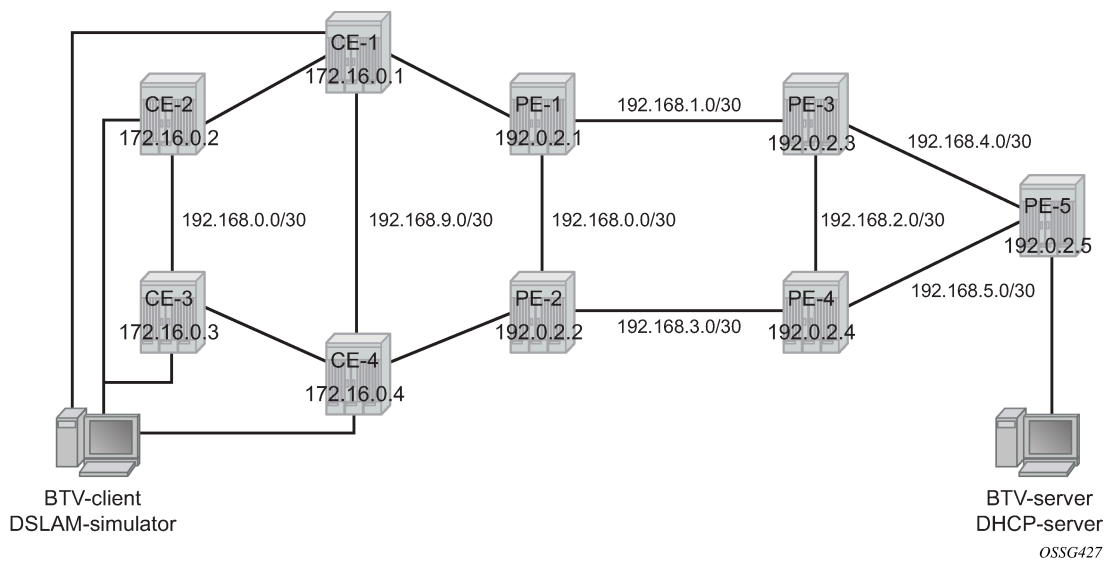
Transition from Broken to Closed Ring State

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 224: Network Topology



The network topology is displayed in [Figure 224: 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 225: 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 225: 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 appropriate platform OS Triple Play 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  : IGMPsnooping 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          Stdby
-----
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-Stdby 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.

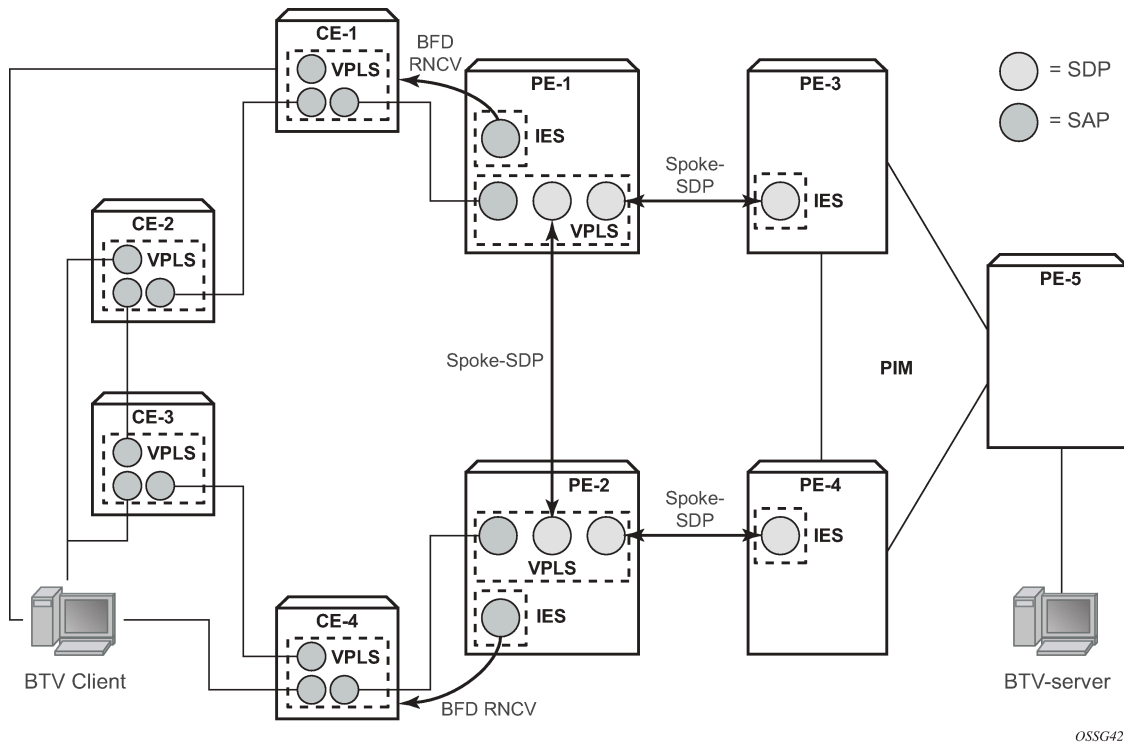
Following the ring failure, the BSA nodes send a MAC address withdrawal message to all the SDP peers configured in the VPLS.

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 226: Multicast Service — Logical Setup](#) shows the logical setup for the multicast service with a MC-ring.

Figure 226: 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
  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 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 227: Test topology](#).

Figure 227: 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
    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
#-----

```

```

echo "AAA Configuration"
#-----
aaa
  radius-server-policy "radius-svr-policy-1" create
    servers
      router "management"
      server 1 name "radius-svr-1"
    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 0x000100000000000000019
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 0x000100000000000000019
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 0x000100000000000001a36
  VSA [26] 12 Alcatel(6527)
    OUTPUT_OUTPROF_OCTETS_64 [22] 10 0x00010000000000000000
  VSA [26] 12 Alcatel(6527)
    OUTPUT_INPROF_PACKETS_64 [25] 10 0x000100000000000000037
  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

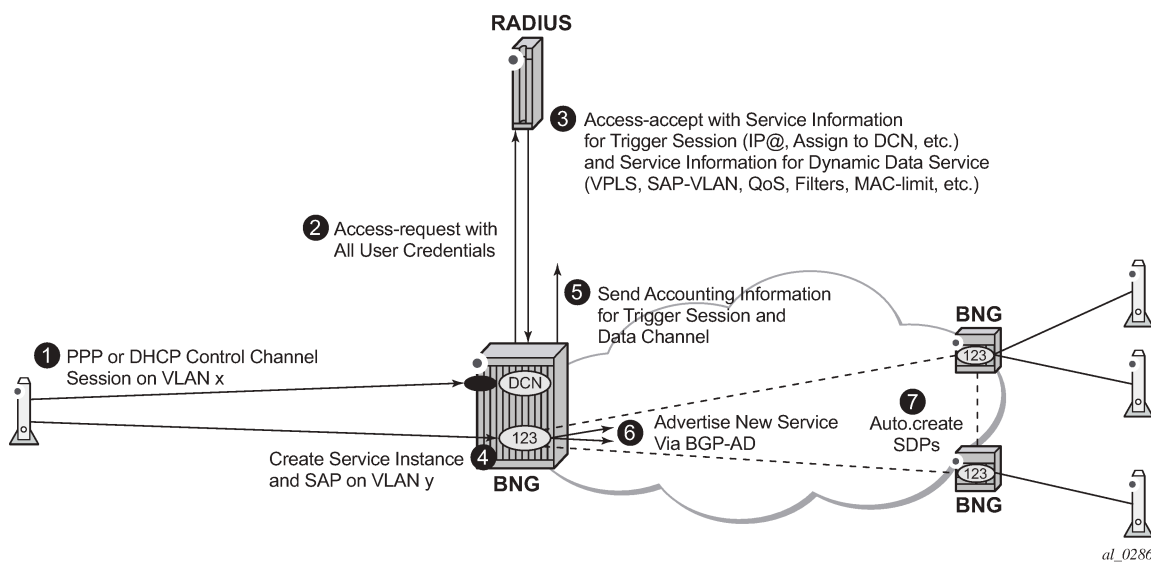
include local Epipe VLL services, Epipe VLL services with dynamic Multi-Segment PseudoWires (MS-PWs) (FEC129) or spoke SDP, VPLS services with BGP-AD PWs or mesh/spoke SDPs, IES, and VPRN services.

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 228: Principle Model of Dynamic Data Services](#) with the example of a VPLS service.

Figure 228: 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 229: Test Topology](#).

Figure 229: Test Topology

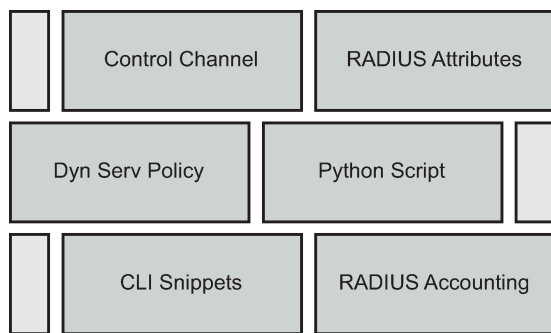


al_0287

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 to 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 230: Building Blocks of Dynamic Data Services



al_0288

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 42: 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 42: 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 43: 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 43: 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 44: Function and Dictionary Relationship](#) :

Table 44: 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 231: 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 42: 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
exit
""" % d)
    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
exit
""" % d)
    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
exit
""" % d)

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
exit
"" % d)
    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
exit
"" % d)
    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
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',

```



```

router
""" % d)
    for route in d['routes']:
        dyn.add_cli("""
            static-route %s next-hop %s metric %s tag 80
            """)
        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
            """) (route["to"], route["n-h"])
        dyn.add_cli("""
            exit
        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
        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

```

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
        sai-type2 64496:192.0.2.1:3333
        tai-type2 64496:192.0.2.5:3333
        no shutdown
    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 also 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 45: 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] = 172.22.250.250 172.22.250/251 | Domain search list [24] = 'Nokia.com' 'test.com' |
| Domain name [15] = 'alcatel.com' | Vendor specific-option [17] = |

| RADIUS | |
|--|----------------------------------|
| DHCPv4 ToClient-Dhcp-Options | DHCPv6 ToClient-Dhcp6-Options |
| | '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 46: 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 47: CLI Inserted DHCP Options

| CLI DHCP Server Pool/prefix Options | |
|---|---|
| DHCPv4 | DHCPv6 |
| DNS server [6] = 172.22.250.253 | DNS server [23] = 2001:db8:1:1:1:1:3 |
| Custom option [231] = 'dhcp injected custom option 231' | Custom option [232]= 'v6 custom option 232' |
| | IA-NA Preferred-Lifetime = 20min (1200sec) |
| | IA-NA Valid-Lifetime = 20min (1200sec) |

| CLI DHCP Server Pool/prefix Options | |
|-------------------------------------|--|
| DHCPv4 | DHCPv6 |
| | 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) |

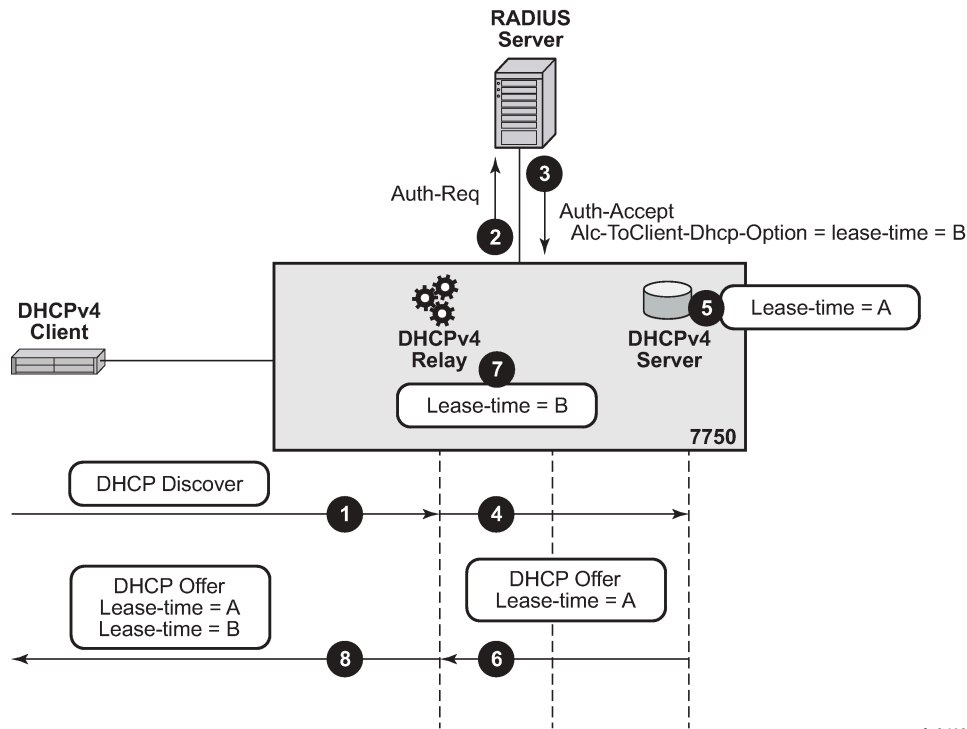
The RADIUS injected raw options are prepended by the DHCP relay agent in the SR OS node to any existing DHCP options already present in the DHCP message before being sent to the client. The existing options could be generated by the DHCP server (internal or external) or by the LUDB. No check is performed 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 232: 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 232: 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 232: DHCPv4 Lease-Time Inserted by RADIUS and DHCPv4 Server



al_0613

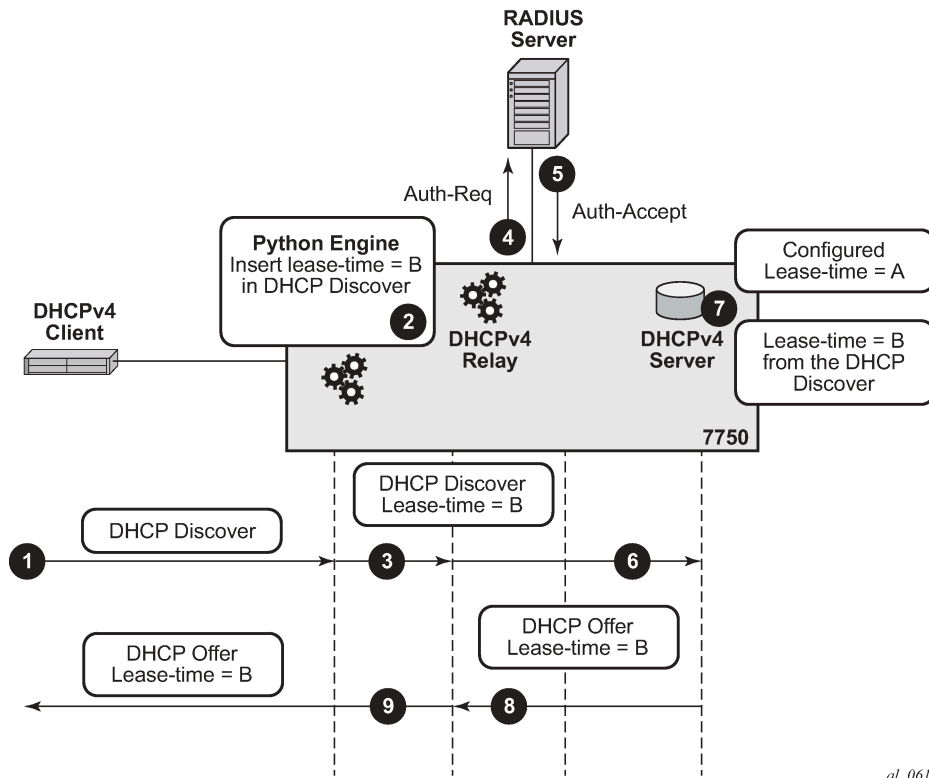
To ensure DHCPv4 lease time unambiguity, the lease-time should be supplied by a single source, in this case by the DHCPv4 server.

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 233: 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 233: Python Injected Hint for Lease-Time

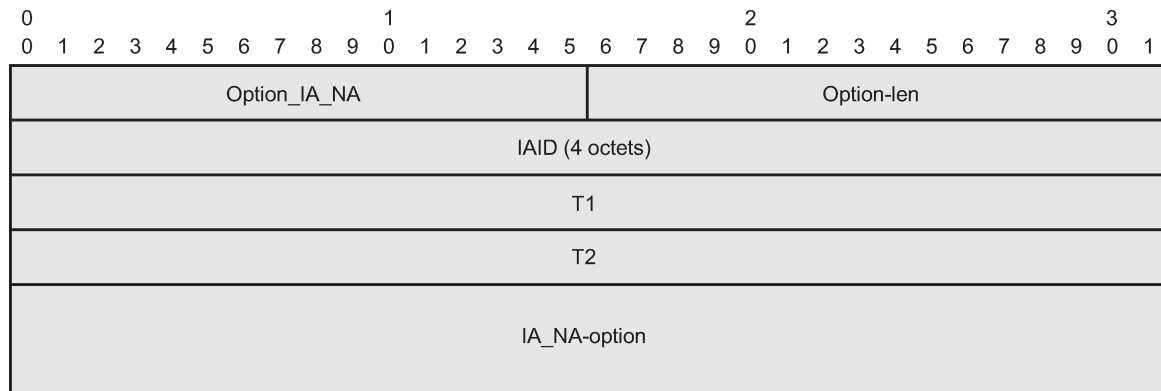


By default the local DHCP server does **not** inject Renew (T1) and Rebind (T2) times so these two timers can still be supplied via RADIUS without duplication by the local DHCP server.

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 234: Format of the IA-NA Option](#) shows the IA-NA option that carries the T1/T2 parameters.

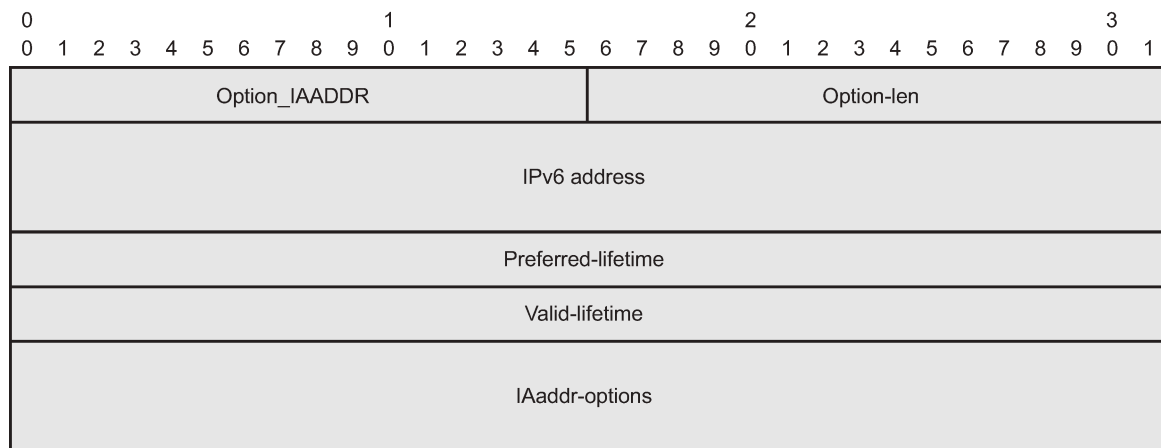
Figure 234: Format of the IA-NA Option



al_0615

The format of the IA address option is shown in [Figure 235: Format of the IA Address Option](#). This option carries preferred and valid lifetimes.

Figure 235: 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 DHCPv6 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 236: Topology](#).

Figure 236: 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 I PoE 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 "int1-1" 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
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 # DHCPv4 option passed on to the client.
      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
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
  egress
    qos 2
  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\xf4', #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 = "g1-1"
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 += 0xf0b616c636174656c2e636f6d
Alc-ToClient-Dhcp-Options += 0xe612637573746f6d2074657374206f7074696f6e
Alc-ToClient-Dhcp-Options += 0x3a040000012c
Alc-ToClient-Dhcp-Options += 0x3b0400000190
Alc-ToClient-Dhcp6-Options +=
0x0011001a0000197f00e60012637573746f6d2074657374206f7074696f6e
Alc-ToClient-Dhcp6-Options +=
0x0017002020010db800010001000100010001000120010db8000100010001000100010001
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 48: 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 += 0xf0b616c636174656c2e636f6d domain-name (15) = alcatel.com |
| Alc-ToClient-Dhcp-Options += 0xe612637573746f6d2074657374206f7074696f6e |

| |
|---|
| <pre>custom -option (230) = "custom test option"</pre> |
| <pre>Alc-ToClient-Dhcp-Options += 0x3a040000012c renewal time T1 (58) = 300s (5min)</pre> |
| <pre>Alc-ToClient-Dhcp-Options += 0x3b0400000190 rebind time T2 (59) = 400s (6min 40sec)</pre> |
| <pre>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"</pre> |
| <pre>Alc-ToClient-Dhcp6-Options += 0x0017002020010db800010001000100010001000120010db800010001000100010002 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</pre> |
| <pre>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)].</pre> |

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:BNG1# configure log
*A:BNG1>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:BNG1>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 47: 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 45: 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 45: RADIUS inserted raw options](#) and [Table 47: 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: 637573746f6d2074657374206f707469666e
  ❑ 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 47: CLI Inserted DHCP Options](#)) followed by the RADIUS supplied options inserted by the DHCPv6 relay ([Table 45: RADIUS inserted raw options](#)).

There are two DNS options [23] since they are supplied via two sources (DHCPv6 server and RADIUS Alclient-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 g1-1
  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: 0000000000000258000003840005001820010db800030001...
    IAID: 00000000
    T1: 600
    T2: 900
  ▣ IA Address
    Option: IA Address (5)
    Length: 24
    Value: 20010db8000300010000000000000001000004b0000004b0
    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:BNGL1# 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  : g1-1
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     : g1-1
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', {}]]}]

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\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\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\x0f\xa0\x00\x19\x00)\x00\x00\x00\x00\x00\x07\xd0\x00\x00
\x0b\xb8\x00\x1a\x00\x19\x00\x00\x0f\xa0\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](#)
- [Summary](#)
- [Overview](#)
- [Configuration](#)
- [Conclusion](#)

Applicability

The information and configuration in this chapter are based on SR OS 11.0.R4.

Summary

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 SAP's, 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*.

Overview

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.

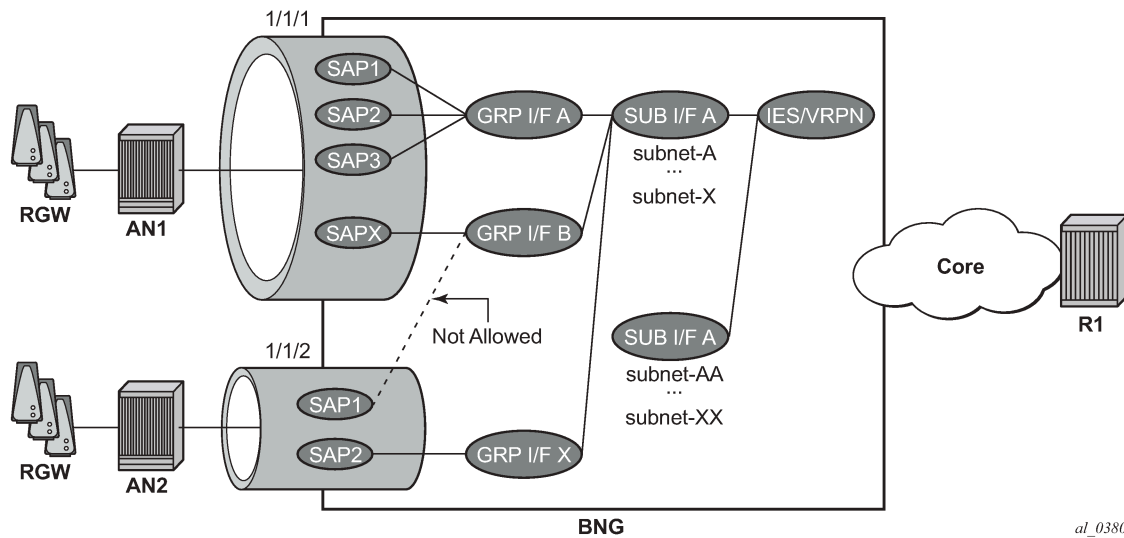
- VLAN/subscriber model.
- VLAN/service/subscriber model.
- VLAN/access node model.
- Per group-interface load balancing in multi-chassis redundancy configurations. Redundancy is out of the scope of this example.

The components needed in the Routed CO model are depicted in [Figure 237: 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 237: Components of the Routed CO Model



al_0380

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
  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
=====

The corresponding IPv4 routing table looks as follows.

```
*A:BN# 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:BN#
```

The corresponding IPv6 routing table looks 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  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:BN#

```

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:BNG# 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:BNG#
```

The IPv4 routing table includes the subnets configured on the subscriber-interfaces.

```
*A:BNG# show router "Base" route-table ipv4
=====
Route Table (Router: Base)
=====
Dest Prefix[Flags]  Type  Proto  Age      Pref
Next Hop[Interface Name]  Metric
-----
10.1.1.0/24         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:BNG#
```

For IPv6, the interface table looks as follows.

```
*A:BNG# show router "Base" interface ipv6
=====
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
                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:BNG# #
```

The IPv6 routing table includes the prefixes configured on the subscriber interfaces.

```
*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  14d04h08m 0
toDHCP-1                0
2001:DB8::75/128       Local Local  14d04h08m 0
system                  0
2001:DB8:101::/48      Local Local  14d03h10m 0
sub-int-1                0
2001:DB8:102::/48      Local Local  14d03h10m 0
sub-int-1                0
2001:DB8:201::/48      Local Local  14d03h08m 0
sub-int-2                0
2001:DB8:202::/48      Local Local  14d03h08m 0
sub-int-2                0
2001:DB8:F101::/48     Local Local  14d03h10m 0
sub-int-1                0
2001:DB8:F102::/48     Local Local  14d03h10m 0
sub-int-1                0
2001:DB8:F201::/48     Local Local  14d03h08m 0
```




```

sub-int-2
2001:DB8:F202::/48
sub-int-2
-----
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 238: 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 238: 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
      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
    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
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
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:BN#

```

The list of active subscribers can be displayed as follows.

```

A:BN# 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 238: 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:BNB# 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:BNB#
```

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: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   02h25m15s    0
  sub-int-1                  0
10.1.1.11/32                Remote Sub Mgmt 02h25m10s    0
  [grp-int-1-1]              0
10.1.1.12/32                Remote Sub Mgmt 00h49m52s    0
  [grp-int-1-1]              0
10.1.1.13/32                Remote Sub Mgmt 00h47m40s    0
  [grp-int-1-2]              0
10.1.2.0/24                 Local  Local   02h25m15s    0
  sub-int-1                  0
10.2.1.0/24                 Local  Local   02h25m15s    0
  sub-int-2                  0
10.2.2.0/24                 Local  Local   02h25m15s    0
  sub-int-2                  0
10.2.2.11/32                Remote Sub Mgmt 02h25m10s    0
  [grp-int-2-2]              0
10.11.11.1/32               Local  Local   02h25m33s    0
  toDHCP-1                  0
192.0.2.75/32               Local  Local   02h25m33s    0
  system                     0
192.0.2.76/32               Remote ISIS    02h24m43s   15
```



```

192.168.12.2
192.168.12.0/24          Local  Local  02h25m15s  0
toR1                    0
192.168.202.0/24       Local  Local  02h25m15s  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: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                    0
2001:DB8::75/128           Local  Local  02h27m24s  0
system                       0
2001:DB8::76/128           Remote  ISIS   02h26m32s  15
FE80::E84C:FFFF:FE00:0-"toR1"  10
2001:DB8:101::/48          Local  Local  02h27m06s  0
sub-int-1                     0
2001:DB8:101:1::1/128      Remote  Sub Mgmt 01h13m43s  0
[grp-int-1-1]                 0
2001:DB8:101:2::1/128      Remote  Sub Mgmt 01h13m25s  0
[grp-int-1-2]                 0
2001:DB8:102::/48          Local  Local  02h27m06s  0
sub-int-1                     0
2001:DB8:102:200::/56      Remote  Sub Mgmt 01h13m43s  0
[grp-int-1-1]                 0
2001:DB8:102:300::/56      Remote  Sub Mgmt 01h13m25s  0
[grp-int-1-2]                 0
2001:DB8:201::/48          Local  Local  02h27m06s  0
sub-int-2                     0
2001:DB8:202::/48          Local  Local  02h27m06s  0
sub-int-2                     0
2001:DEAD::/64             Local  Local  02h27m05s  0
toR1                           0
-----
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:BN#

```

The corresponding IPv6 FIB on card 1 is as follows.

```

*A:BN# 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:BN#

```

The addresses of the individual subscriber hosts show up in the RIB but they do not show up in the FIB.

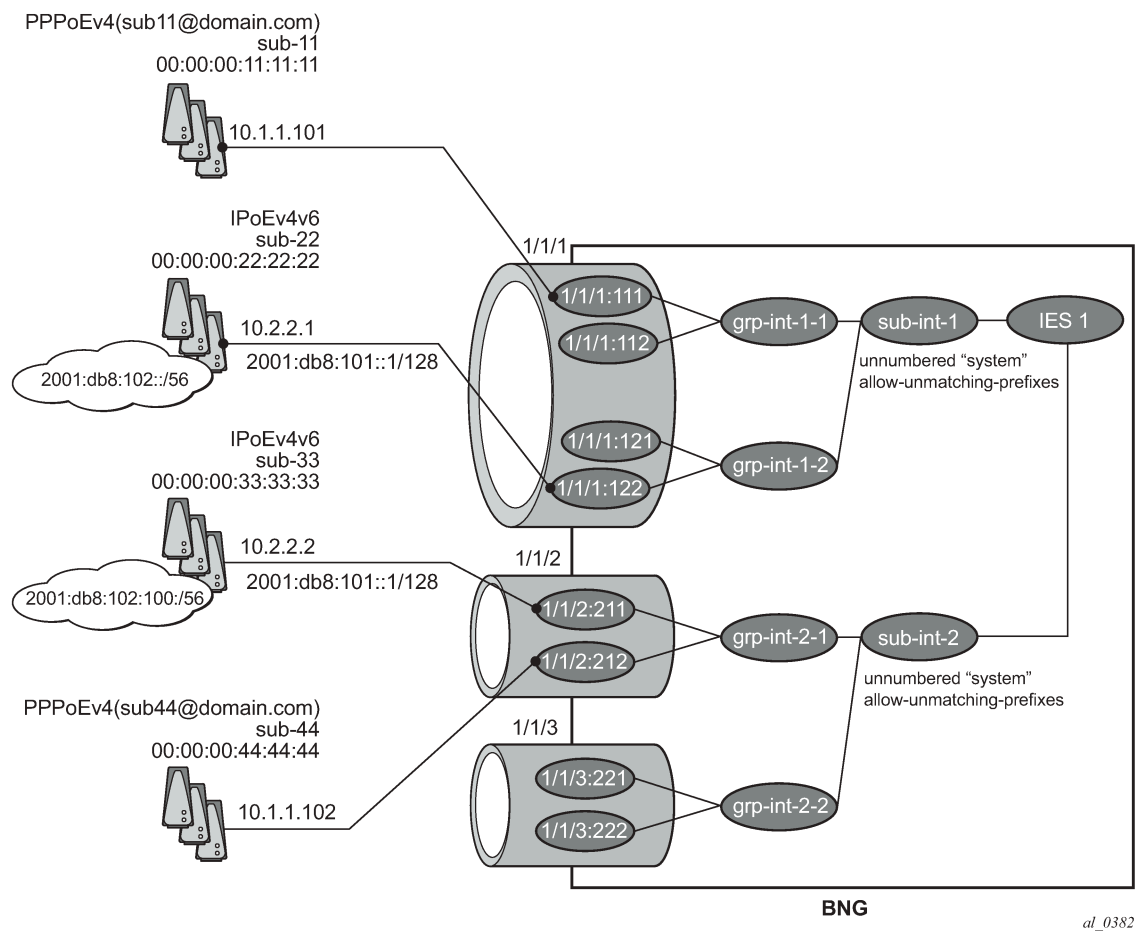
- /32 for IPv4-hosts.
- /DPL (Delegated Prefix Length) for IPv6 DP hosts, /56 in this example.
- /128 or /64 for IPv6 wan host.

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 239: 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 239: 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
        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
  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
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      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
lb-pool4-1          Up        Up/Down      Network  loopback
10.1.1.254/24      n/a
lb-pool4-2          Up        Up/Down      Network  loopback
10.1.2.254/24      n/a
lb-pool4-3          Up        Up/Down      Network  loopback
10.2.1.254/24      n/a
lb-pool4-4          Up        Up/Down      Network  loopback
10.2.2.254/24      n/a
sub-int-1          Up        Up/Up       IES Sub  subscriber
Unnumbered If[system]
FE80::EA:4B:FF/64  n/a
PREFERRED
sub-int-2          Up        Up/Up       IES Sub  subscriber
Unnumbered If[system]
FE80::EA:4B:FF/64  n/a
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/Down      Network  1/1/12
192.168.12.1/24   n/a
toRADIUS-1          Up        Up/Down      Network  1/1/10
192.168.202.75/24 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      Subscriber
IP Address
MAC Address      PPPoE-SID Origin      Fwding State
-----
1/1/1:111      sub-11
10.1.1.101
00:00:00:11:11:11      1      IPCP      Fwding
1/1/1:122      sub-22
10.2.2.1
00:00:00:22:22:22      N/A      DHCP      Fwding
1/1/1:122      sub-22
2001:DB8:101::1/128
00:00:00:22:22:22      N/A      IPoE-DHCP6      Fwding
1/1/1:122      sub-22
2001:DB8:102::/56
00:00:00:22:22:22      N/A      IPoE-DHCP6      Fwding
1/1/2:211      sub-33
10.2.2.2
00:00:00:33:33:33      N/A      DHCP      Fwding
1/1/2:211      sub-33
2001:DB8:101:1::1/128
00:00:00:33:33:33      N/A      IPoE-DHCP6      Fwding
1/1/2:211      sub-33
2001:DB8:102:100::/56
00:00:00:33:33:33      N/A      IPoE-DHCP6      Fwding
1/1/2:212      sub-44
10.1.1.102
00:00:00:44:44:44      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
|   |
|   |-- 10.2.2.1
```



```

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 10
192.168.12.0/24 Local Local 00h49m24s 0
toR1
192.168.202.0/24 Local Local 00h49m24s 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: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]
Next Hop[Interface Name] Type Proto Age Pref
Metric
-----
2001:DB8::11/128 Local Local 01h03m27s 0
toDHCP-1
2001:DB8::75/128 Local Local 00h06m34s 0
system
2001:DB8:101::1/128 Remote Sub Mgmt 00h36m12s 0
[grp-int-1-2]
2001:DB8:101:1::1/128 Remote Sub Mgmt 00h35m58s 0
[grp-int-2-1]
2001:DB8:102::/56 Remote Sub Mgmt 00h36m12s 0
[grp-int-1-2]
2001:DB8:102:100::/56 Remote Sub Mgmt 00h35m58s 0
[grp-int-2-1]
-----
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
-----
=====
A:BN#

```

The corresponding IPv6 FIB on card 1 looks as follows:

```

A:BN# 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
-----
=====
A:BN#

```

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

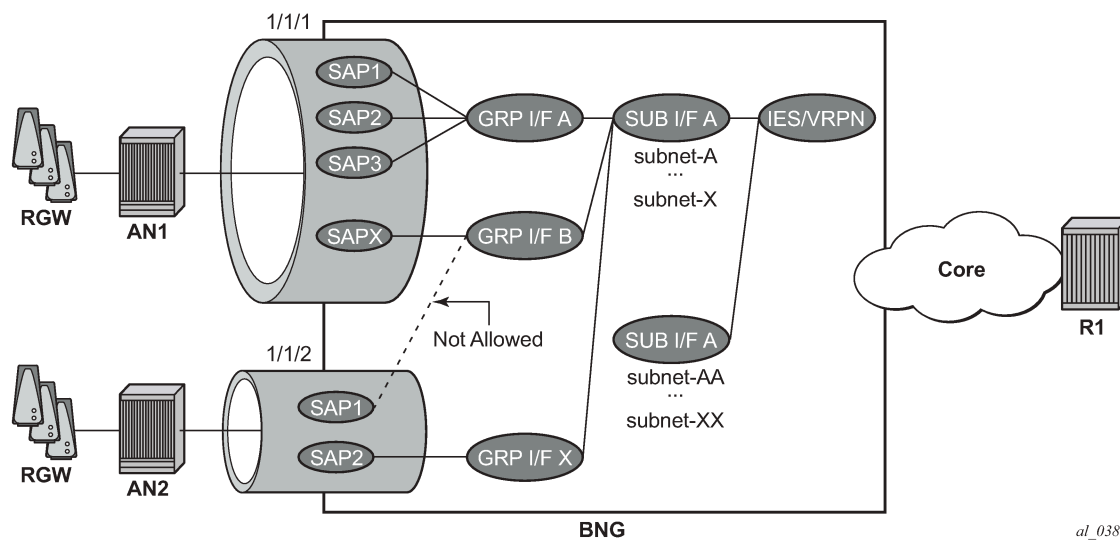
An alternative to the scenarios described above does exist in the form of a mixed numbered/unnumbered (hybrid) scenario as depicted in [Figure 240: Hybrid Configuration](#).

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 240: 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 237: 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]
  Next Hop[Interface Name]
Type   Proto   Age           Pref
Metric
-----
10.1.1.0/24
  192.168.12.1      Remote  ISIS      00h14m11s  15
                    20
10.1.2.0/24
  192.168.12.1      Remote  ISIS      00h14m11s  15
                    20
10.2.1.0/24
  192.168.12.1      Remote  ISIS      00h14m05s  15
                    20
10.2.2.0/24
  192.168.12.1      Remote  ISIS      00h14m05s  15
                    20
192.0.2.75/32
  192.168.12.1      Remote  ISIS      00h14m17s  15
                    10
192.0.2.76/32
  system            Local   Local     62d21h22m  0
                    0
192.168.12.0/24
  toBNG             Local   Local     05d04h43m  0
                    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]
Next Hop[Interface Name]          Type   Proto   Age           Pref
Metric
-----
2001:DB8::75/128
FE80::E84B:FFFF:FE00:0-"toBNG"    Remote ISIS   00h04m25s  15
10
2001:DB8::76/128
system                             Local  Local   02d05h54m   0
0
2001:DB8:101::/48
FE80::E84B:FFFF:FE00:0-"toBNG"    Remote ISIS   00h04m25s  15
20
2001:DB8:102::/48
FE80::E84B:FFFF:FE00:0-"toBNG"    Remote ISIS   00h04m25s  15
20
2001:DB8:201::/48
FE80::E84B:FFFF:FE00:0-"toBNG"    Remote ISIS   00h04m25s  15
20
2001:DB8:202::/48
FE80::E84B:FFFF:FE00:0-"toBNG"    Remote ISIS   00h04m25s  15
20
2001:DEAD::/64
toBNG                               Local  Local   01h42m18s   0
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                               LocalPref  MED
      Nexthop                               Path-Id     Label
      As-Path
-----
i     10.1.1.0/24                             100        None
      192.0.2.75                             None        -
      No As-Path
i     10.1.2.0/24                             100        None
      192.0.2.75                             None        -
      No As-Path
i     10.2.1.0/24                             100        None
      192.0.2.75                             None        -
      No As-Path
i     10.2.2.0/24                             100        None
      192.0.2.75                             None        -
      No As-Path
i     10.11.11.1/32                          100        None
      192.0.2.75                             None        -
      No As-Path
i     192.0.2.75/32                          100        None
      192.0.2.75                             None        -
      No As-Path
i     192.168.12.0/24                        100        None
      192.0.2.75                             None        -
      No As-Path
i     192.168.202.0/24                       100        None
      192.0.2.75                             None        -
      No As-Path
-----
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                100      0
      None                        2
?     2001:DB8:101:1::1/128
      ::FFFF:C000:24B
      No As-Path                100      0
      None                        2
?     2001:DB8:102::/56
      ::FFFF:C000:24B
      No As-Path                100      0
      None                        2
?     2001:DB8:102:100::/56
      ::FFFF:C000:24B
      No As-Path                100      0
      None                        2
-----
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]  Path  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                  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 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                 0
2001:DB8:101::1/128    Remote BGP     00h36m40s    170
192.0.2.75 (tunneled) 0
2001:DB8:101:1::1/128 Remote BGP     00h36m40s    170
192.0.2.75 (tunneled) 0
2001:DB8:102::/56      Remote BGP     00h36m40s    170
192.0.2.75 (tunneled) 0
2001:DB8:102:100::/56 Remote BGP     00h36m40s    170
192.0.2.75 (tunneled) 0
2001:DEAD::/64         Local  Local   01h21m47s    0
toBNG                  0
-----
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](#)
- [Summary](#)
- [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.

Summary

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 etc.) 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 master-slave 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.


Overview

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 241: Network Redundancy Components for ESM Routed CO](#).

Figure 241: Network Redundancy Components for ESM Routed CO



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/VP RN 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 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 master (for example, fails to receive SRRP messages from the master), 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 master; 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 master.

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 master, 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 master, 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/VPRN 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 master) will be forwarded by the backup SRRP router to the master 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
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
    exit
  no shutdown
  exit
exit
exit
```

At this point, any subscriber ARPing for the gw-ip-address will receive a response from the SRRP master 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 master.

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 master.

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
```

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 master 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 master if three consecutive messages are not received.

At all times the keep-alive-interval of the master 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 master/backup election, based on the priority (1 to 254) configured in the SRRP instance. The device with the highest priority will become master.

```
*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 master). The master/backup status is per SRRP instance (not per IP address). For example, the master 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 master 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 master.

A basic form of load distribution can be achieved by having the master SRRP for some group-interfaces on one peer and the master for other group interfaces on the other peer. Clearly, a failure may cause all masters 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 master 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 master, 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
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 master 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
```

```
        exit
    no shutdown
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 master of the related group-interface.

When a subscriber is created on the master 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 a dummy group interface with the **oper-up-when-empty** command.

```
# on BNG-2
configure
  service
    vprn 1 customer 1 create
      subscriber-interface "sub-int-1" create
      group-interface "group-int-1" create
      oper-up-while-empty
      sap 1/1/3:1 create
      ---snip---
    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 master.

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 master.

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

```

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 master 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]
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:BNB-2#
```

The same subscriber on the backup peer will have a host route in the FIB pointing to the redundant interface.

```
*A:BNB-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:BNB-1#
```

Subscriber Related Information

To check the subscriber information:

```
*A:BNB-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:BNB-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: 021BFF000000015CA701E9
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
Omcrc Alarmed Entries:        0
Omcrc 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      DLen  Flags      timeStamp
SyncTag     deleteReason code and description      #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
Omcrc Alarmed Entries:        0
Omcrc 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: 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:

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:BNB-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 master 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
Msg 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 master 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 master the operational state is always master – 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 master) 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:BNG-2#
```

BFD-Related Information

To check the BFD session state.

```
*A:BNG-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 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:BNG-2#
```

To check the MAC addresses of the SRRP, this can be done by checking the MACs table in the BSAN.

```
*A:DSLAM# show service fdb-mac
```

```
=====
Service Forwarding Database
=====
ServId      MAC                Source-Identifier  Type   Last Change
          Transport:Tnl-Id
-----
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=Oam  P=Protected-MAC  C=Conditional  S=Static  Lf=Leaf
=====
*A:DSLAM#
```

SRRP Debug Commands

There are debug commands 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
```

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 master 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 and 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/VRPN 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 242: 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 242: 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 243: Explicit BRG authentication](#). The first phase is the BRG authentication phase, the second phase is the device authentication phase. The example in [Figure 243: Explicit BRG authentication](#) uses IPv4, but also works with IPv6.

Figure 243: 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 IPoE sessions.

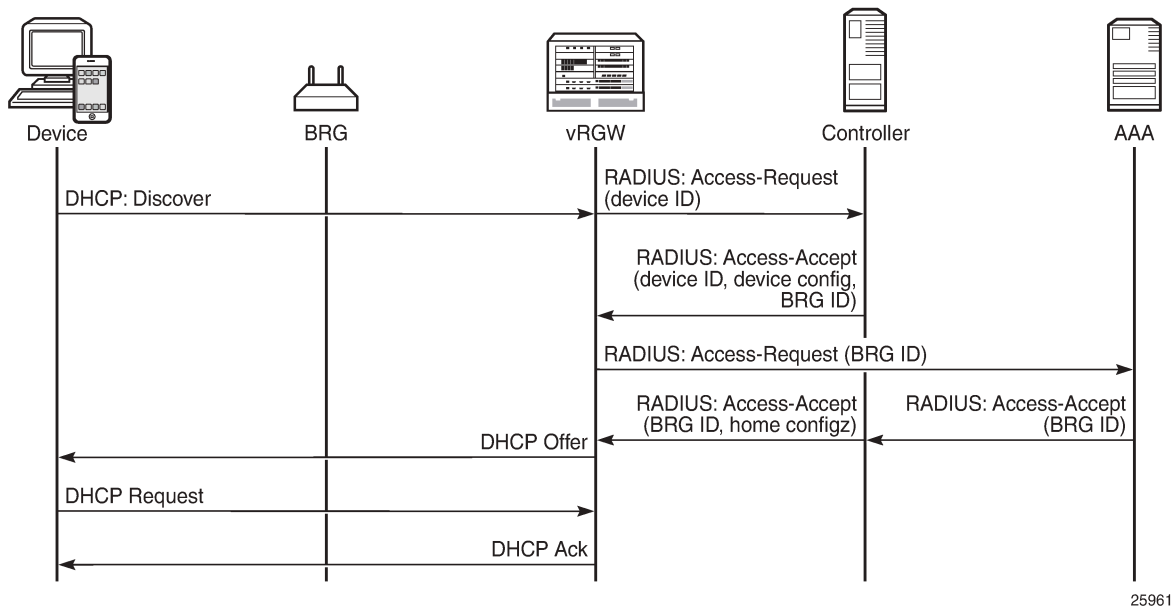
If the home-specific data includes Alc-Reserved-Addresses attributes defining static devices, these devices are authenticated automatically as soon as at least one dynamic device is connected.

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 244: Implicit BRG authentication](#) for an example.

Figure 244: Implicit BRG authentication



In [Figure 244: 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 IPoE 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 245: 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 245: 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
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
    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
    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

```

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.TR6SdCXXMO/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:BN#

```

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:BN# 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:BN#
```

Even when the last dynamic host disconnects, the static hosts remain, as shown by the following command.

```
*A:BN# 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:BN#
```

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:BN# clear subscriber-mgmt brg gateway brg-id "00:0c:29:00:00:10"

*A:BN# show subscriber-mgmt brg gateway brg-id "00:0c:29:00:00:10" hosts
No entries found.

*A:BN# show subscriber-mgmt brg gateways
No entries found.
*A:BN#
```

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 TcHWkAooYb5TcPJ9KPR3M.
    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:BN#
```

Releasing both dynamic devices results in all BRG hosts being deleted, including the static hosts.

```
*A:BN# show subscriber-mgmt brg gateway brg-id "00:0c:29:00:00:10" hosts
No entries found.
*A:BN#
```

Also, the BRG is deleted automatically.

```
*A:BN# show subscriber-mgmt brg gateways
No entries found.
*A:BN#
```

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 246: vRGW-HLE](#).

Figure 246: vRGW-HLE



27632

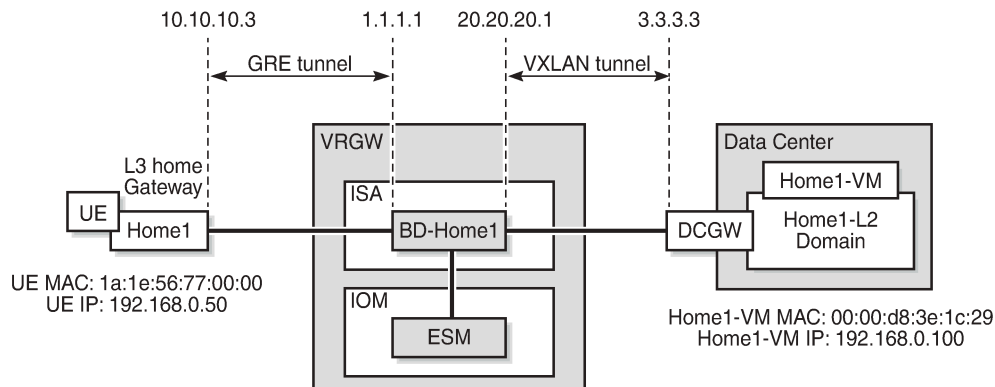
Bridge Domain Characteristics

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 247: BD Connections in the Data Plane shows the BD connections in the data plane.

Figure 247: 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
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
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

```

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
```

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: SVCMGR #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: SVCMGR #8452 Not allowed when lanext is no shut

*A:WLANGW>grp-if>wlan-gw>ranges>range>vrgw>lanext>access# multi-access
MINOR: SVCMGR #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.

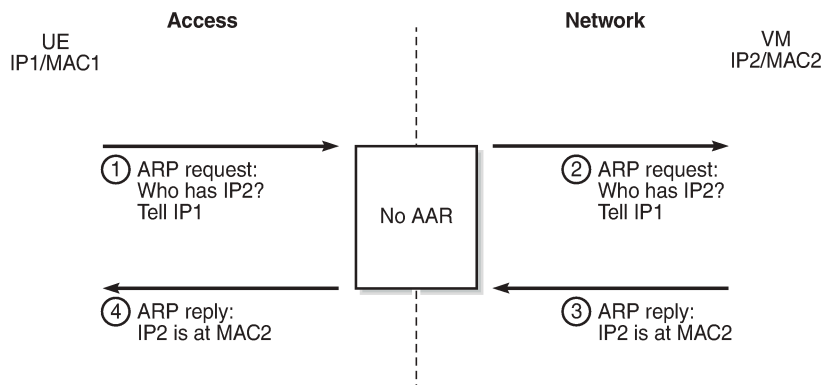
With AAR enabled, the system populates the ARP and neighbor tables with the learned ARP and neighbor entries via either:

- 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 248: 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 248: ARP Requests Flooded with AAR Disabled



27634

AAR Enabled

To enable AAR, **assistive-address-resolution** is configured in the **lanext** context, as follows:

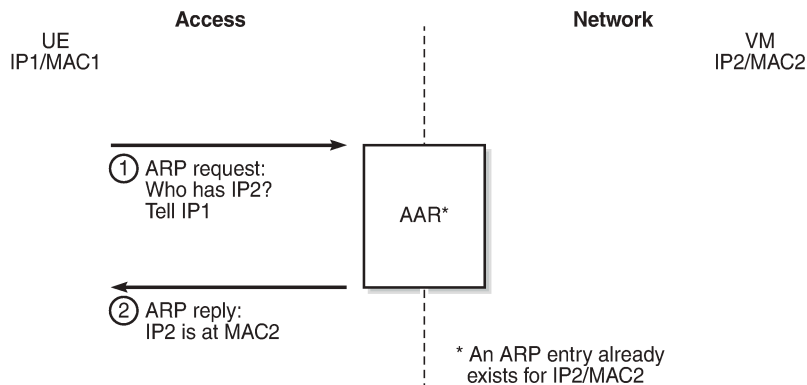
```
*A:SR7-CMPT-vRGW>grp-if>wlan-gw>ranges>range>vrgw>lanext#
[no] assistive-addr* - Enable/disable assistive address resolution in bridged domains
```

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: SVCMGR #8452 Not allowed when lanext is no shut
```

In [Figure 249: 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 249: No ARP Request Flooding with AAR Enabled



27635

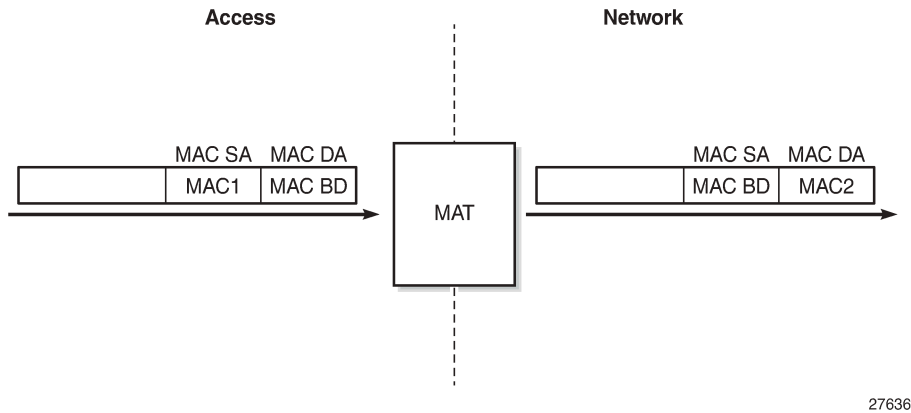
MAC Address Translation

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:

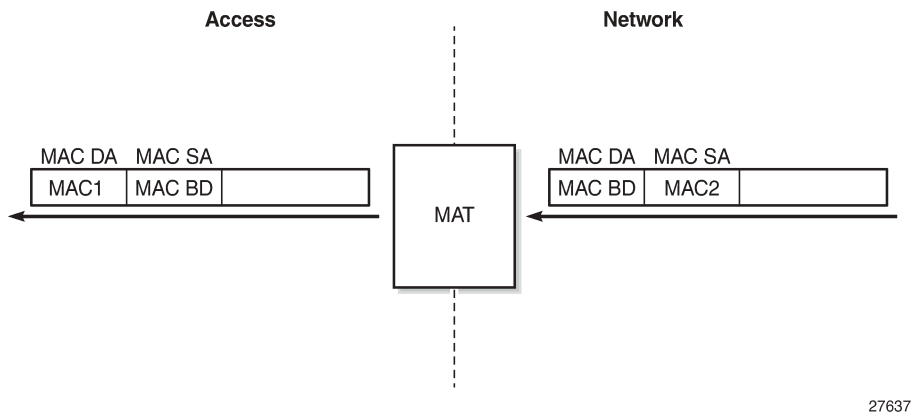
- traffic from access to network - the system changes the source MAC address to BD MAC address, as shown in [Figure 250: MAT - Access to Network Direction](#).

Figure 250: 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 251: MAT - Network to Access Direction](#).

Figure 251: 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: SVCNMR #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: SVCMGR #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: SVCMGR #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 250: MAT - Access to Network Direction](#) and [Figure 251: 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 49: ARP/MAT/BD-MAC-Prefix Possible Combinations](#) shows possible combinations for ARP/MAT/BD-MAC-Prefix, where Combo1 is the default combination.

Table 49: 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  Router     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  Dst-IP
Anchor
-----
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  Class      Type      Vlan  Src-IP
Port    Class      Encap     Vlan  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 252: 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.

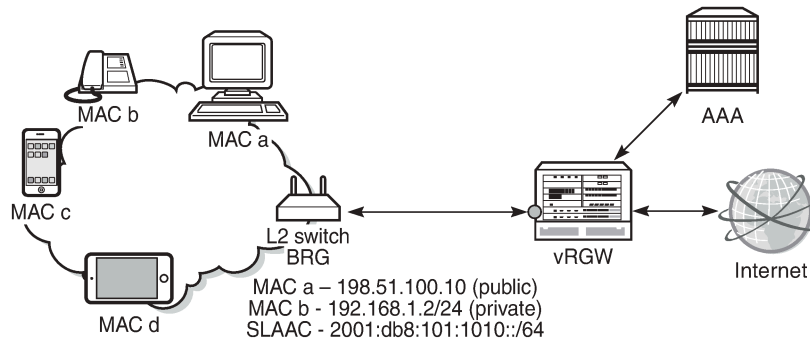
For the vRGW to offer IPv6 connectivity to the devices in the home network, the vRGW provides the following features:

- 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 252: 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.
- Static IPv6 addresses are not supported. However, a static IPv4 device can get an IPv6 SLAAC prefix if IPoE-linking is enabled.
- No per-home pool is supported for IA_NA, but DHCPv6 can be relayed to an external or local DHCPv6 server.

The vRGW only operates in IPoE bridged mode, so that multiple hosts on the same BRG share the same SLAAC/64 prefix. Only one SLAAC prefix per BRG is supported.

Prefix delegation (PD) is not supported in the vRGW.

Demilitarized Zone (DMZ)

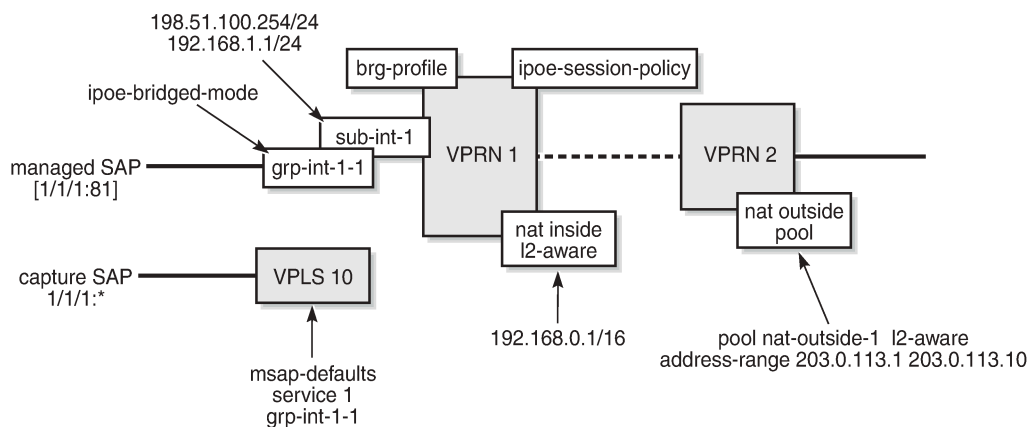
When a DMZ host is specified within a home, all downstream traffic not matching an existing flow or port forwarding rule is sent to that host. Without a DMZ host, this traffic is dropped. DMZ requires a single port-reservation block to be defined on the outside L2-aware NAT range.

Configuration

The services configuration overview shown in [Figure 253: Services Configuration Overview](#) applies to the examples used throughout this chapter. For L2-aware NAT, two routed services are configured. In this example, 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. Addresses from the documentation range as defined in RFC 5737 are used for the public addresses.

The examples in this chapter use capture and managed SAPs. The capture SAP is created manually on VPLS-10, and the managed SAPs are created dynamically on VPRN-1. Also, static SAPs can be used to support BRGs, but in that case the configuration file can grow rapidly when many SAPs are required.

Figure 253: Services Configuration Overview



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
    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
      exit
      no shutdown
    exit
  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

```

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
```

```
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
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
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
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:BN#
```

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:BN# 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:ddd:1::1
DNS 2 v4            : 1.1.2.2
DNS 2 v6            : 2001:db8:ddd: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:BN#
```

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:BN#

```

The following command shows the subscriber hosts on VRPN-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:BN# 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:BN#

```

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:BN# show service id 1 dhcp lease-state

```



```

=====
DHCP lease state table, service 1
=====
IP Address      Mac Address      Sap/Sdp Id      Remaining Lease   MC
                :                :                : LeaseTime : Origin  : 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:BN#
    
```

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:BN# show router 1 route-table

=====
Route Table (Service: 1)
=====
Dest Prefix[Flags]      Type  Proto  Age      Pref
  Next Hop[Interface Name]      Metric
-----
10.11.11.1/32           Local  Local  11h48m56s  0
   int-DHCP              0
192.168.1.0/24          Local  Local  11h48m56s  0
   sub-int-1             0
198.51.100.0/24         Local  Local  11h48m56s  0
   sub-int-1             0
198.51.100.110/32       Remote Sub Mgmt 11h48m41s  0
   [grp-int-1-1]        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:BN#
    
```

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:BN# 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           Mac Address      Up Time      MC-Stdbby
Subscriber-Id
[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:BN#

```

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.

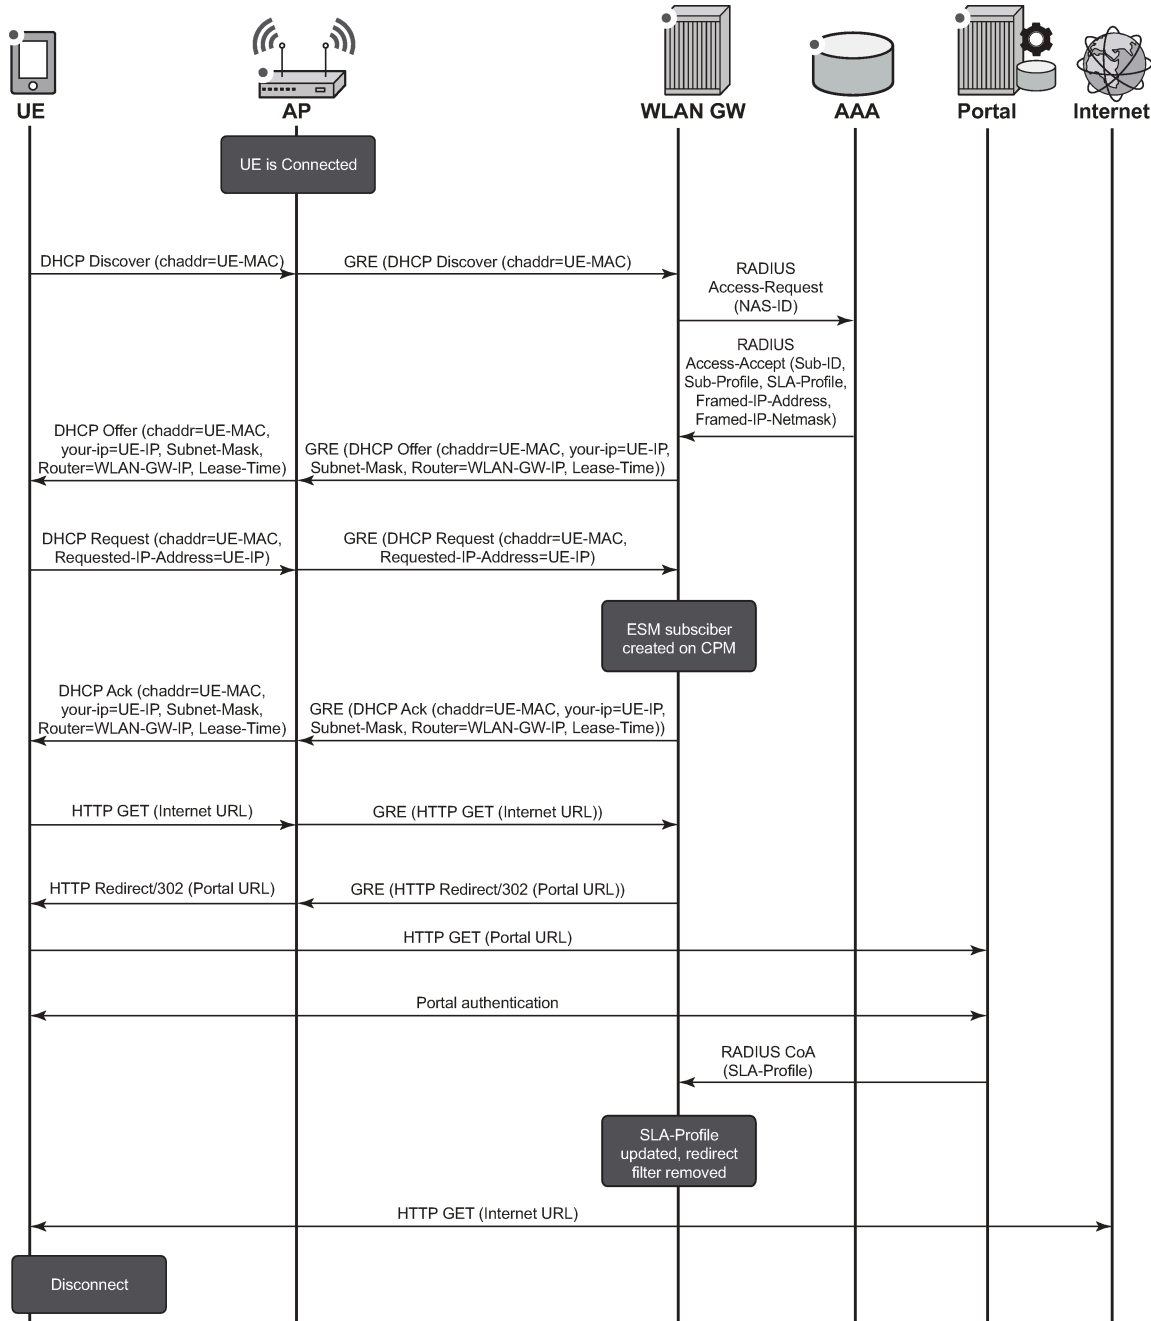
Besides authentication, a major consideration is the method used to achieve Internet connectivity. Will the users require public addresses or are private addresses sufficient? In case few public IP addresses are available, private IP addresses can be assigned to UEs and the WLAN-GW can perform a Network 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 254: 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

Figure 254: Call Flow for Open SSID



al_0831

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 255: 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 255: 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 255: 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
    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 endpoint 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
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
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 256: 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 256: 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 "7USmr6f7JkxaGnDDq1uqwEAJKGbhZr5i" hash2
              no shutdown
            exit
          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
  -----
```

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
-----
```

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 "nUeorYjgFZtuAqIwoUOL0DFxF43rhSf/" hash2
            send-accounting-response
            wlan-gw
                address 192.11.11.11
            exit
            no shutdown
        exit
    exit
    no shutdown
-----
```

VPRN 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 "S11" create
  address 192.168.1.1/24 populate-host-routes
  group-interface "G11" 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
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
=====

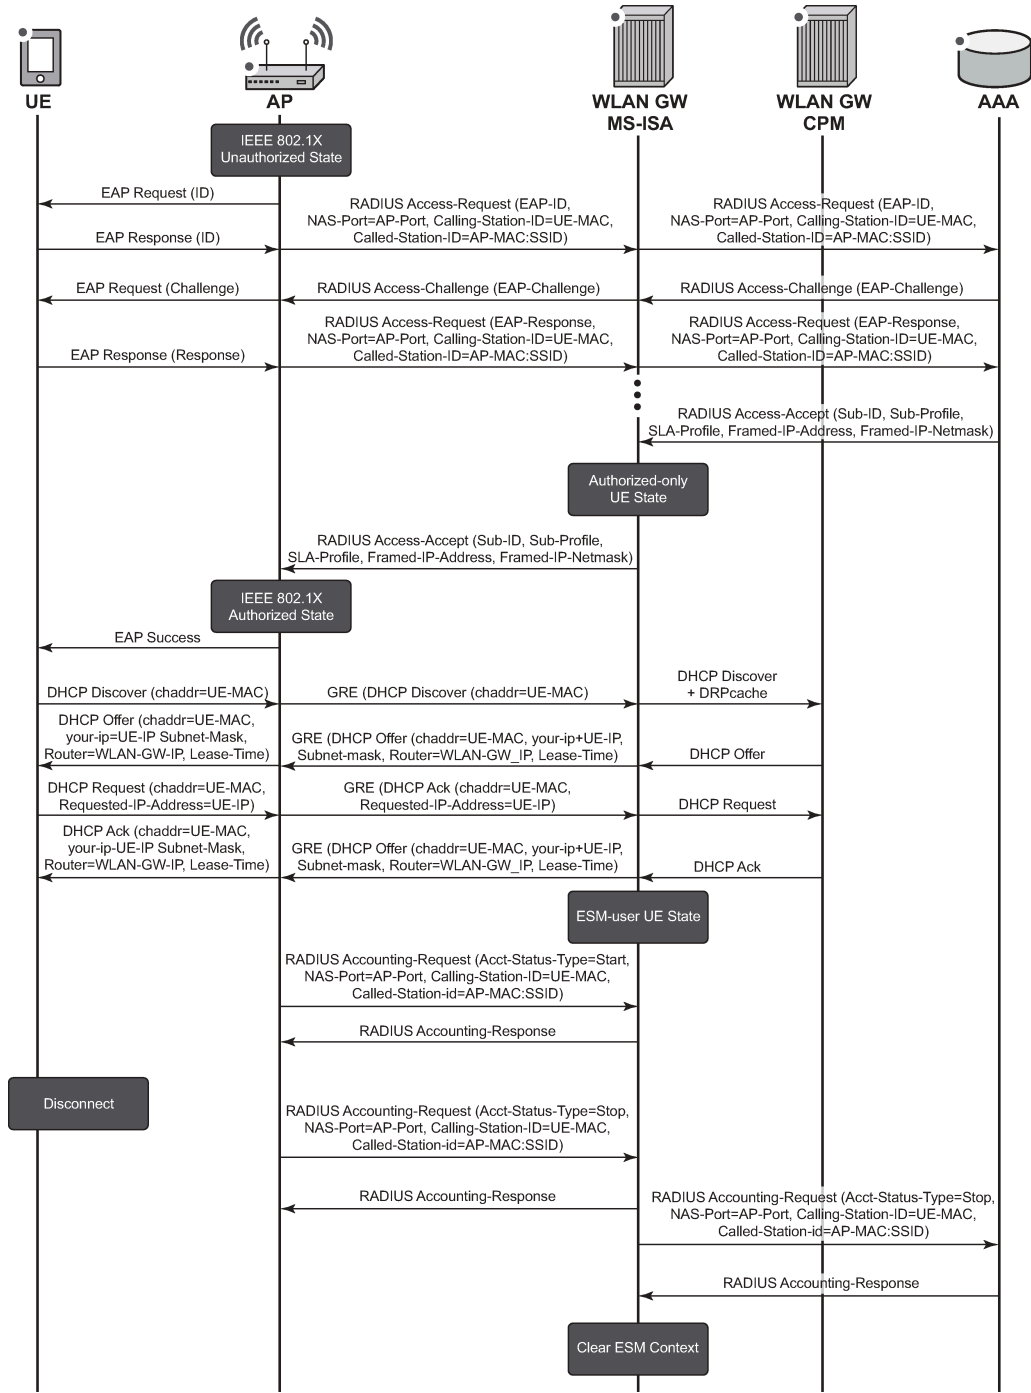
```

Call Flow

Figure 257: 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 257: 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 0xcbe7fb9182312534ea50ecdffc8ed59874401515968ae276
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.

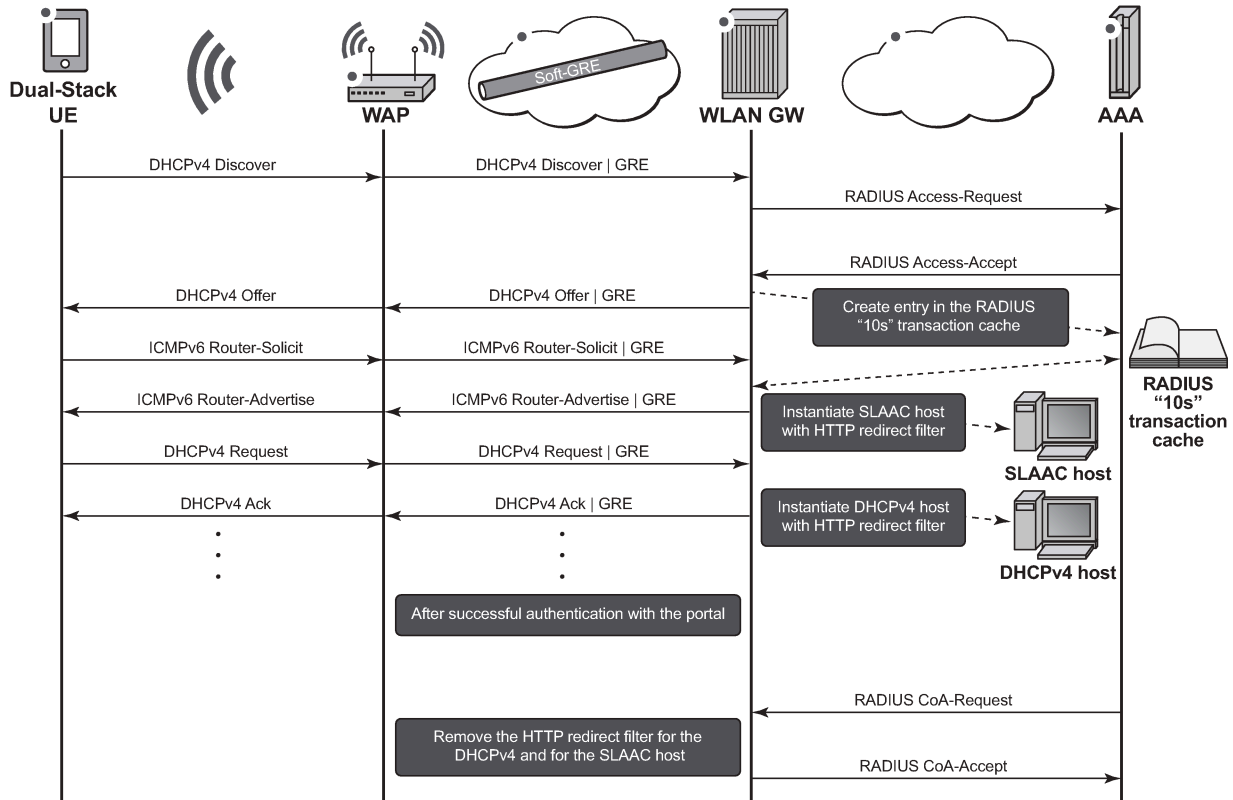
Depending on the type of UE, it may be allocated an IPv4 address through DHCPv4 and an IPv6 address through Stateless Address Auto-Configuration (SLAAC), or DHCPv6, or both (not all UEs have support for DHCPv6). Therefore, a UE can instantiate up to three IPoE hosts: a DHCPv4 host, a SLAAC host, and 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 258: 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.

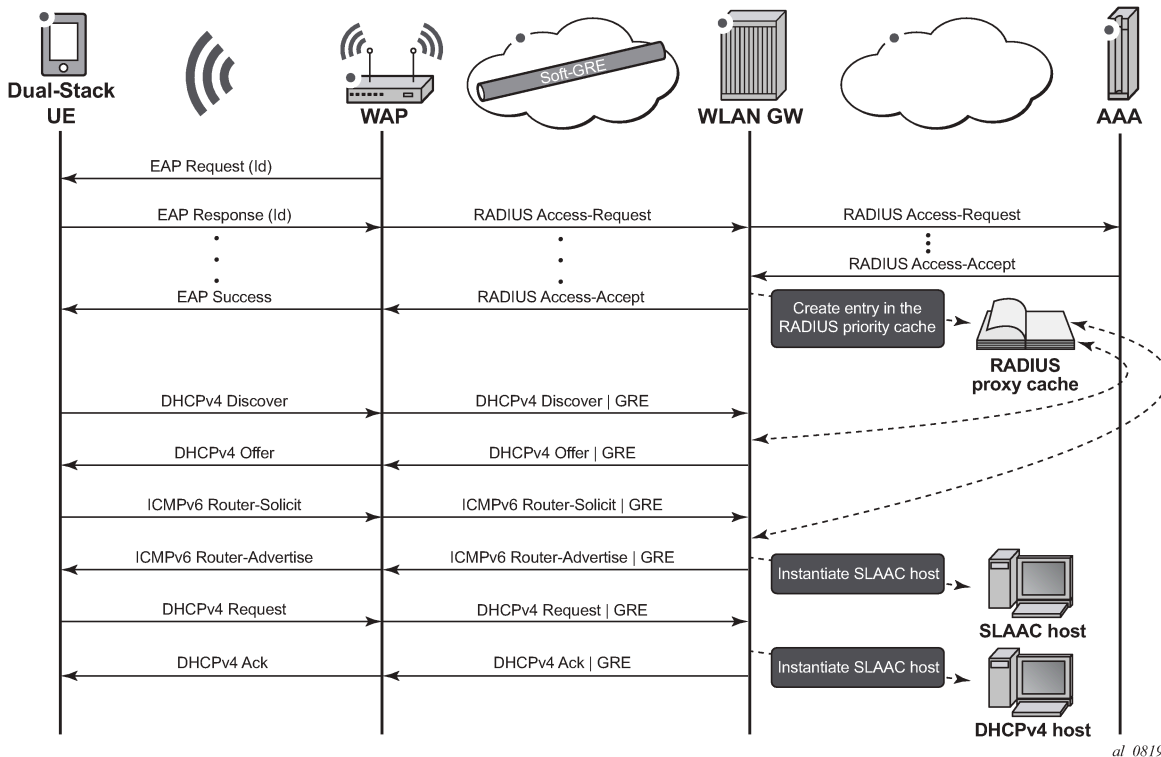
The following three IPv4/v6 dual-stack UE IP address assignment models are available:

- 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 259: 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 260: DHCPv4 + SLAAC/64 with DHCPv4 linking model — closed SSID

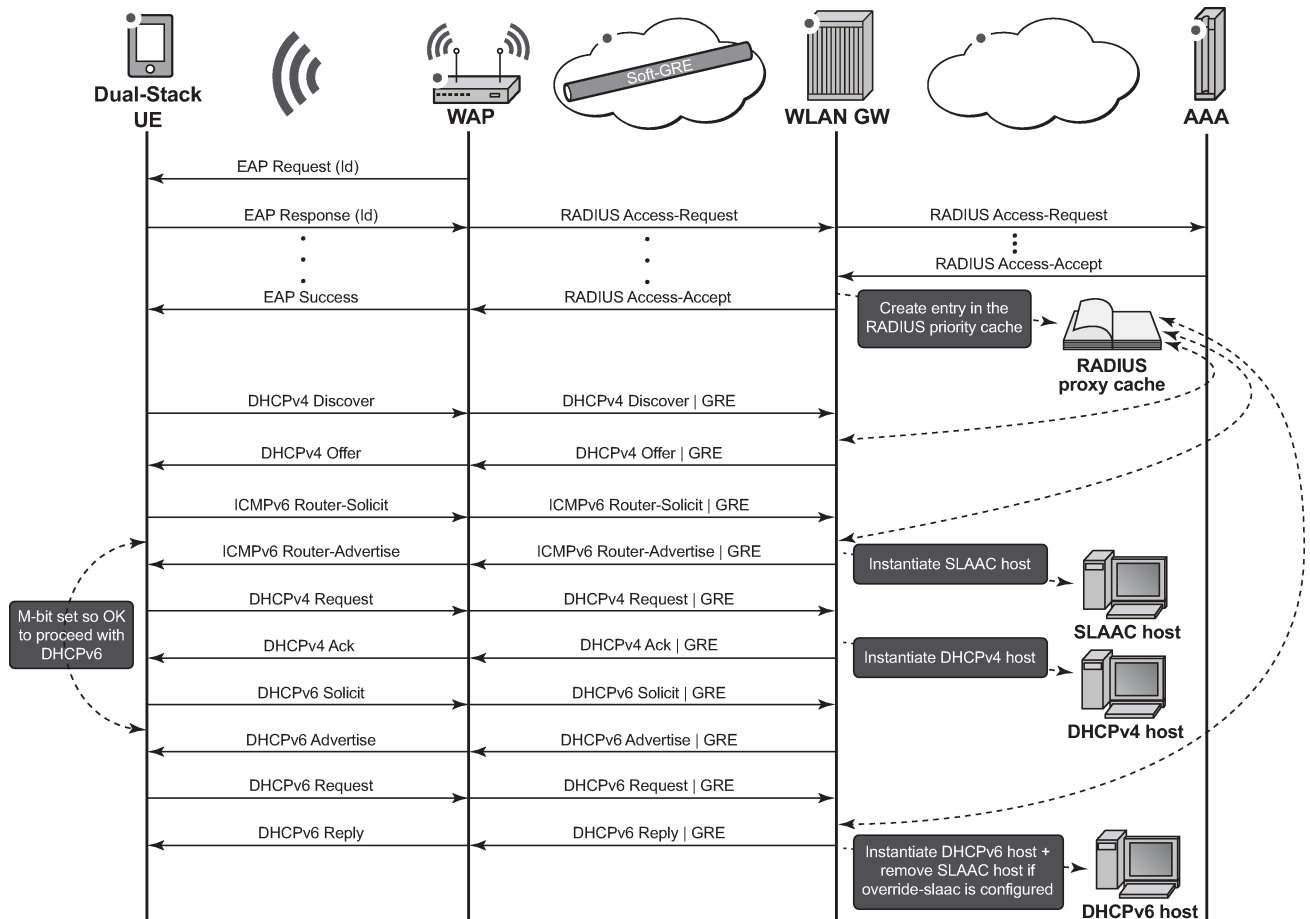


al_0820

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 261: DHCPv4 + DHCPv6/128 IA_NA model — closed SSID



al_0821


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 262: DHCPv4 + SLAAC/64 with DHCPv4 linking model — DTA



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
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
-----
10.255.0.1
2001:db8:ffff::/64
-----
MAC Address
-----
b0:9f:ba:b9:40:f8 N/A
b0:9f:ba:b9:40:f8 N/A
-----
PPPoE-SID Origin
-----
DHCP
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
    
```

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: SVCMGR #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: SVCMGR #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  : 192.0.2.2        GRE Dst IP   : 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.

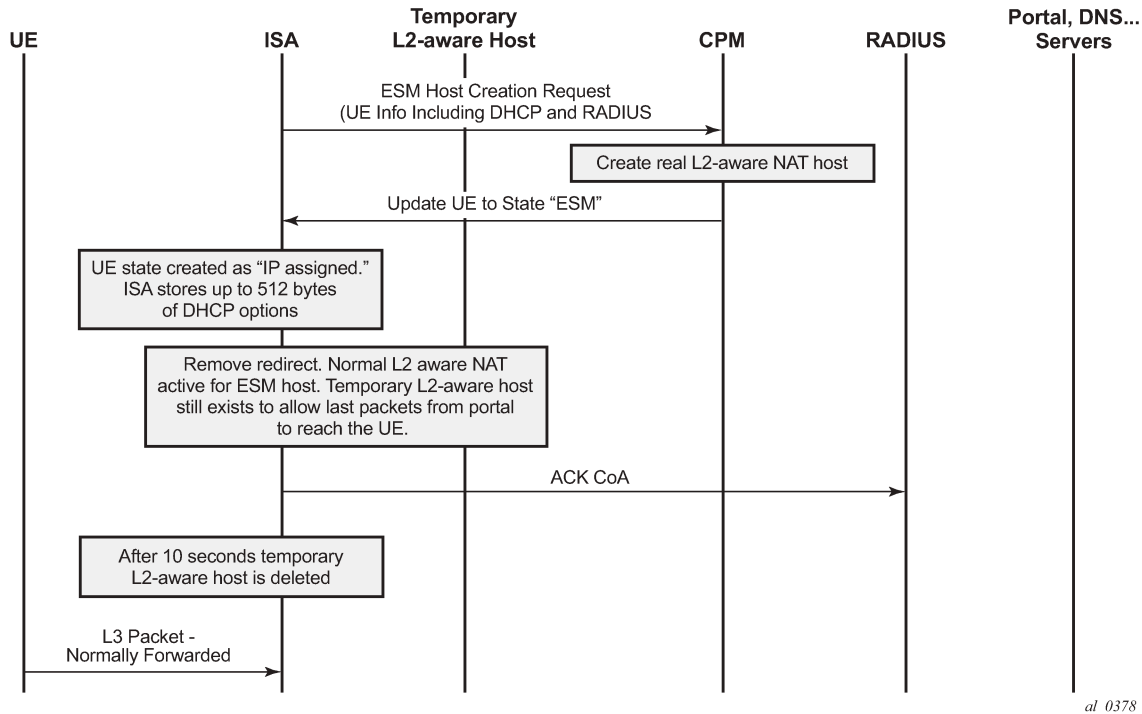
The overall sequence of events to take a UE from migrant to authenticated state, where the forwarding of UE traffic is not restricted, is shown in [Figure 263: Sequence of events to establish and authenticate a migrant User](#).

Figure 263: Sequence of events to establish and authenticate a migrant User



al_0377

Figure 264: 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 "E40PedK6agrEIPr2DEoJyVR8PQ3XkFF7" 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-attribs
```

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

```

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

exit
no shutdown
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](#)
- [Summary](#)
- [Overview](#)
- [Configuration](#)
- [Conclusion](#)

Applicability

The information and configuration in this chapter are based on SR OS Release 12.0.R4.

Summary

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.

Overview

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 265: 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 265: WiFi Offload Scenario with Open SSID, DSM and LI



al_0814

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](#)